QUESTION WORDS IN SCALAR ENRICHMENT:

A COMPUTATIONAL APPROACH

by

SAMUEL H. FISHER

(Under the direction of Don McCreary)

ABSTRACT

In this thesis, I investigate the feasibility of an interactive computational model for improving the lexical semantic inferencing of the user. The model assumes lexical semantic inferencing to be a function of lexical knowledge and lexical semantic entailment scales enabling implicature. During processes of user-computer question answering, implicature scales are targeted for enrichment, and their component lexical items are added to the question-answer exchange. Particular emphasis is placed on the role of question words in a question-answer exchange and their relation to scalar implicature.

INDEX WORDS: Dialogue processing, question answering, conversational implicature, attentional state, anaphora resolution, lexical processing

Question Words in Scalar Enrichment:
A Computational Approach

by

Samuel H. Fisher

B.A., Emory University, 1997

A Thesis Submitted to the Graduate Faculty
of The University of Georgia in Partial Fulfillment
of the
Requirements for the Degree

Master of Arts

Athens, Georgia

2007

Question Words in Scalar Enrichment:
A Computational Approach

by

Samuel H. Fisher

Approved:

Major Professor: Don McCreary

Committee: Sarah Blackwell
Keith Langston

Electronic Version Approved:

Maureen Grasso
Dean of the Graduate School
The University of Georgia
May 2007

Acknowledgements

I wish to thank my major professor, Don McCreary, for his guidance and attention during the design and writing of this thesis. I wish also to thank my committee members, Sarah Blackwell and Keith Langston. I extend additional thanks to Joe McFall for his continued friendship and intellectual comradery over the last decade; to William Kretzschmar, Jr. for his help in deciding on a tag set for use in computing research; to Paula Schwanenflugel for her concision as a researcher and teacher; and to the host of Linguistics faculty and students with whom I had the pleasure of learning.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

In this thesis, I investigate the feasibility of an interactive computational model for improving the lexical semantic inferencing of the user. The model assumes lexical semantic inferencing to be a function of lexical knowledge and lexical semantic entailment scales enabling implicature. During processes of user-computer question answering, implicature scales are targeted for enrichment, and their component lexical items are added to the question-answer exchange. Particular emphasis is placed on the role of question words in a question-answer exchange and their relation to scalar implicature. For example, the word *which* with the word *car* (e.g. *which car*) may be used in a question to extend the speaker's awareness of a more specific meaning of *car* which applies to the context of the dialogue and to the intention and knowledge of the addressee.

The interpretation of question words here assumes interrogative pronouns to be long-distance anaphors, or intersentential dialogical cataphors. Long-range anaphora are defined as exophoric, which means their referents lie outside the immediate linguistic discourse context (Bhat 2004). Interrogative pronouns refer to entities, attitudes, events, states, times, processes, conditions, and attributes (i.e. anything that can be referred to), which are not known by the speaker, and therefore require the cooperation of an interlocutor or external information source for reference resolution.

A second assumption is that question words are instrumental in creating lexical sets which form the basis for conversational implicature. For example, consider that a learner of English as a Foreign Language (EFL) points to a helicopter and says "What call that airplane?" The learner is seeking some lexical item that discriminates a more specific type of aircraft. When the question word is resolved to "helicopter" (as by someone answering the question), then the speaker may be able to use the contrast between *airplane* and *helicopter* more effectively to understand and achieve implicature in future discourse. Thus, in dialogue, interrogative pronouns are used to attend to the intentional structure of the discourse by a speaker. Quoting Asher (1998: 7), "[the speaker] asked the question because s/he doesn't have enough information to achieve the intention."

I propose that by using and tracking question words (in a model of attentional state) during processes of human-computer question answering, lexical semantic sets which license implicature are identifiable and their component lexical items can be automatically introduced to the interlocutors through subsequent Q/A exchanges. This "mining out" of implicature scales through Q/A sequences leads, over time, to the contextualization and ordering of lexical sets in the user's mental lexicon. The process of introducing implicature scales to a user is called *scalar enrichment*.

Scalar enrichment is a process of linguistic interaction in which individuals augment their lexical knowledge with certain lexemes that may be unknown or unstructured with regard to other lexemes which are already known. When a language user acquires the ability to recognize the contrasts between lexemes relating to the same set of semantic entailment, then the potential for pragmatic enrichment emerges. Scalar enrichment, as pursued here, is not to be confused with *pragmatic enrichment*. According to Gundel (2003),

> "pragmatics is construed as an account of the inferential processes which take as their input the result of linguistic decoding and 'enrich' that input by way of pragmatic inferences for those aspects of a speaker's intended meaning that are left underspecified by linguistic form, e.g. reference and ambiguity resolution and conversational implicature."

In contrast, scalar enrichment applies to linguistic decoding, whereby linguistic forms either do not exist for a speaker prior to a discourse or are erroneously decoded because they are not structured appropriately for that context. The knowledge following *scalar* enrichment is necessary for understanding and producing conversational implicature, for enabling pragmatic enrichment. For example, the question word *which* enables the enrichment of semantic subsets that are necessary for inferring and eliciting a more robust set of intentions. If an interlocutor uses the word *car*, then the intention is that the meaning of *car* does not necessarily entail the meaning of *minivan*. It is through the use of *which* in combination with *car* (i.e. *which car*) in contextualized linguistic interaction that an interocutor tests and possibly extends the intentionality of an utterance to cover some subordinate lexeme of *which car* (e.g. *minivan*). The process could prove useful in a number of domains including ESL instruction, where increased exposure to and interaction with contextualized vocabulary improves the learner's systems of lexical inference.

1.1 Thesis Organization

In what follows, I explore how a computer program may be created to facilitate an awareness of lexical use, and subsequently of implicature. Lexical awareness relies on conversational interactivity, context, discourse coherence, and lexical knowledge. First, a review of the literature provides the framework from which the current investigation emerges. It includes theoretical and applied works on scalar implicature and discourse processing.

Following the literature review, I present an overview of the computer system that was developed. It serves as the starting point for understanding the larger operations in which scalar enrichment is proposed through question answering.

In chapter 4, the basic components of wh-questions and the functional aspects of interrogative proforms are reviewed. This includes justification for viewing interrogative proforms as anaphora, as well as analysis of how interrogative pronouns factor through the attentional structure of dialogue to manifest themselves in the cognitive apparatus where lexical knowledge and implicatural contrast reside.

In chapter 5, I present a model that has been implemented in the computer program. This proves to explicate how anaphora function in a model of attentional state. An algorithm for resolving the reference of pronouns and interrogative pronouns is traced.

Chapter 6 presents how the lexical knowledge and scalar contrast that is recognized by an interlocutor can be augmented through the operation of question words in the attentional state model of dialogue. In doing so, I characterize the representation of lexical knowledge and scalar contrast in the computer's memory and show how this structure interacts with the attentional model through computer-based resources. Moreover, I present a data structure and procedure for linking the attentional state model to the computer's lexical knowledge base.

Chapter 2

Literature Review

The literature review aims to provide a summary of contributions central to the basic assumption that question words can be implemented programmatically in a dialogue agent to strengthen the contrast between lexical items so they may later be used by interlocutors in producing and understanding conversational implicature. In what follows, I examine contributions from the literature on scalar implicature, Centering Theory in attentional state, the cognitive status of referring expressions used in dialogue, and the pragmatic interpretation of question words as definite or indefinite proforms. This will set the groundwork for understanding what has been implemented in the software and what improvements need to be made to it in order to tie into an interactive procedure for scalar enrichment.

2.1 Scalar Implicature

Scalar implicature is identified with Grice's (1975) maxim of quantity, which states that a speaker be (1) as informative as required and (2) not more informative than required (for the addressee to infer the speaker's intention along with the conventional meaning of the linguistic forms expressed in the utterance). Assuming interlocutors are cooperative in dialogue (Grice's Cooperative Principle), the maxim suggests that interlocutors realize a contrast between what is said in a particular context of reference and what a speaker could have said in a possibly related context of reference. This includes choosing among words to describe what a speaker believes to be true and what a speaker cannot confirm as true. At minimum, lexical items of a particular class must be drawn into contrast in order to select the more (or less) informative one. Horn (1984) elaborates on this idea by proposing scales of opposition for canonical sets (now known as "Horn Scales") including items denoting degrees of quantity (e.g. <all, some>), and logical connectors (e.g. <and, or>), among others.

Lexical semantic entailment is a logical process of reasoning through Horn Scales wherein "an ordered n-tuple of expression alternates $<x_1, x_1, \ldots, x_n>$ such that where S is an arbitrary simplex sentence-frame and $x_i > x_j$, S(x_i) unilaterally entails S(x_j) [and S(x_j) does not necessarily entail S(x_i)]" (Levinson 2000: 79). Along this interpretation of entailment, Horn Scales represent an increasing or decreasing degree of necessary entailment depending on how they are used by a speaker. Horn illustrates the contrast between lexical items in a scalar set through squares of opposition. An example square for the two sets <all, some> and <none, not all> appears in Figure 2.1, taken from Levinson (2000: 65).

Figure 2.1: The traditional square of opposition

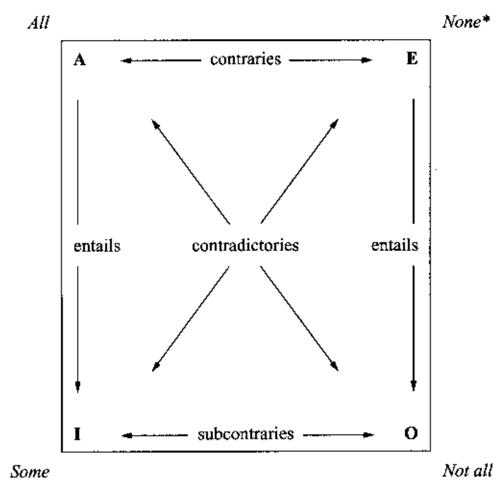

The square of opposition above delineates entailment, contrary, contradictory, and subcontrary relations between lexical items from two Horn Scales. When a lexical item *entails* another lexical item and the entailing item is expressed as true, then the entailed item (the latter) must also be true. For the scale of <*all, some*>, the expression of "all" as true means that "some" is also true. However, the expression of "some" as true means that "all" is false and "not all" is true. The relation of *subcontaries* is captured by the relation between "some" and "not all". Both have to be true.

Unlike subcontrary relations, lexical items which are *contraries* of each other cannot both be true. For example, “all” and “none” cannot both be true in specifying the quantity for the same set of things; they are contraries of each other. They can, however, both be false. *Contradictory* relationships between the meanings of two words characterize when the words cannot both be true or false at the same time. Thus, when “all” is false, “not all” has to be true; when “all” is true, “not all” has to be false. The same can be said of “some” and “none”; when “some” is true, then “none” has to be false, and vice versa.

Hirschberg (1985), building on Horn’s scales of opposition, proposes a theory of scalar implicature that combines scales of contrast with sets of alternate or non-hierarchical words. These scales constitute partially ordered sets, or *posets*, of lexical items. For example, in the poset <{granny smith, golden delicious, honey crisp}, apple>, the subset {*granny smith, golden delicious, honey crisp*} constitutes an alternate set where any of the words in the alternate set are equally more specific (and thus stronger) than the other subset containing just *apple*. Moreover, by use of a word in an alternate set, all other words in the alternate set are cancelled. She further espouses that lexical posets enable the contrast of nearly any two items, provided there is some context in which they may be ordered. For instance, assuming the ordering of *eat–for–lunch–today*, the set <*spaghetti*, *fried chicken*, *sushi*> may emerge. The values in the scale are entirely dependent on the available options presented by the context (i.e. location, stores that sell some food item, whether certain stores are open, how much money is available to spend on food, etc.)

Levinson (1983, 2000) argues for two ways in which scales of opposition may be organized in dialogue. One he attributes to generalized conversational implicature (GCI) and the other to particularized conversational implicature (PCI). PCI’s, he maintains, are less conventional in their application and constitute a context-dependent establishment of relational orderings, i.e. some particular context. Specifically, “an implicature i from utterance U is *particularized* iff U implicates i only by virtue of specific contextual assumptions that would not invariably or even normally obtain” (Levinson 2000:16). The relations used to make GCI’s on the other hand, are wide-spread within a community of speakers. Of GCI's he claims: "an implicature i is *generalized* iff U implicates i *unless* there are unusual specific contextual assumptions that defeat it" (Levinson 2000:16). Levinson provides examples distinguishing GCI's from PCI's, reproduced below:

(1) A: "What time is it?"

B: "Some of the guests are already leaving."

PCI: 'It must be late.'

GCI:''Not all of the guests are already leaving.'

The PCI in the example above attends to the "goals or plans" of interlocutors and may be thought to subsume or include the GCI interpretation that 'Not all of the guests are already leaving' (Levinson 2000:17). GCI's are, in effect, conversational implicatures made at the lexical semantic level of contrast and entailment. The contrast is ubiquitous in the mental lexicons of a community of speakers:

> The GCI theorist is simply claiming that speakers carry their lexicons on their backs as it were, from context to context, and it is *mutual knowledge*[1] of this fact that elevates the Q-heuristic [i.e. quantity scalar opposition] to a default mode of inference (Levinson 2000: 108).

Levinson (2000) further elucidates how anaphoric expressions such as pronouns and the ways articles may be used constitute GCI scales of opposition with regard to their levels of specificity in making reference. This is important, since anaphors help establish coherence in dialogue, bringing interlocutors into mutual understanding of discourse topic and focus.[2] For instance, Levinson (2000: 289) identifies the scale < <he, John>, the man > to show that "he" draws a coreference to its antecedent (e.g. "John"), [3] while "the man" makes a disjoint

[1] My italics

[2] A discourse topic is the general thing being talked about while the focus is the lexical item(s) specifying the topic in one or more ways.

[3] Along the same line of reasoning, "he" entails "John" due to "he" being more specific within the discourse about John. Had (2a) been chosen, then "John" might be interpreted as a different person named "John" (i.e. $John_3$).

reference to some other man who is not "John." Consider the following example sentences from Levinson (2000: 287).

(2) a. "*John*$_1$ told *her*$_2$ that *John*$_1$ gave *her*$_2$ a valentine."

b. "*John*$_1$ told *her*$_2$ that *he*$_1$ gave *her*$_2$ a valentine."

c. "*John*$_1$ told *her*$_2$ that *the man*$_3$ gave *her*$_2$ a valentine."

In terms of semantic entailment, use of *the man* in (2c) means that the speaker cannot truthfully state the more specific reference to "John," which is stated in (2b). From this example, one can see how anaphoric expressions may also constitute GCI scales of opposition. With coreference, discourse coherence is maintained; however, disjoint reference marks a loosening of coherence since a new discourse entity is introduced or referenced, marking a topic change.

2.2 Attentional State

Centering theory emerged from the efforts of Grosz (1977), Sidner (1981, 1983), Grosz and Sidner (1986), Joshi and Weinstein (1981), and Grosz, et al. (1995) to develop a general theory of discourse. Particular to attentional state, Sidner (1983) proposed the notion of discourse focus. Focusing on local coherence in discourse segments, Grosz, et al. (1995) championed the notion that discourse requires coherence, and that by definition a discourse segment is most coherent when it contains one and only one center. For example, compare the following utterance sequences from Grosz, et al. (1995: 206).

(3) a. John went to his favorite music store to buy a piano.

b. He had frequented the store for many years.

c. He was excited that he could finally buy a piano.

d. He arrived just as the store was closing for the day.

(4) a. John went to his favorite music store to buy a piano.

b. It was a store John had frequented for many years.

c. He was excited that he could finally buy a piano.

d. It was closing just as John arrived.

In (3), the discourse is more coherent because it centers around only one discourse entity, i.e. "John." In (4), there is more than one entity or topic being focused on (i.e. "John", "store"). In (3), no topic change occurs despite the presentation of other discourse entities. "John" remains front and center, so to speak, throughout all four utterances. In (4), however, the topic changes three times (at 4b, 4c, and 4d), oscillating between "John" and "store" with "he" and "it." For this reason, example (3) is more coherent than example (4).

The center is the discourse entity most in focus by the interlocutors' attentional states. Degrees of coherence were speculated to affect the inferential load of the hearer (Joshi and Weinstein, 1981). It was also observed that when a single discourse entity is kept in focus, then that discourse entity is most likely to be pronominalized.[4] Psycholinguistic studies (Hudson-D'Zmura and Tanenhaus, 1998) confirm these intuitions about coherence and inferential load. From these findings, centering algorithms became useful for anaphora resolution, especially for pronouns. Therefore anaphora resolution is considered a side effect of the overall process of topic tracking in centering theory (Tetreault, 2001). The degree of continuity about the center from utterance to utterance is also correlated with ease of comprehension in reading and speech (Hudson–D'Zmura and Tanenhaus 1998). Centering chiefly differs from other models and algorithms for anaphora resolution in two regards: (1) discourse is marked by coherence, (2) discourse segments containing precisely one center of attention (or topic) across several utterances, have a greater degree of continuity and as such are more coherent (Grosz, et al. 1995).

[4]Example (4) at (4b), does not express this tendency. "John's favorite music store" is pronominalized to "it" and elevated to the discourse center, while "John" is not pronominalzed at all. This is an atypical example of pronominalization because it eschews a system of lesser coherence. Example (3) illustrates the point completely. "John" is immediately pronominalized at (4b) and remains so throughout the example.

2.3 Models for Anaphora Resolution

From studies of the (semantic) propositional content of linguistic expressions, Discourse Representation Theory (DRT) emerged (see Kamp 1981 and Kamp and Reyle 1993). DRT tracks discourse and discourse segments by abstracting sentences into logical statements, or conceptual representations intended to account for various phenomena including, but not limited to, quantification. It differs from centering theory in that multiple discourse entities may share equal significance throughout a local discourse segment. Associated with the DRT model due to its promise of intrasentential anaphora resolution are syntax-based algorithms for matching pronouns with their antecedents. Hobbs (1978) and Lappin and Leass (1994) present two such syntax-based algorithms for pronoun resolution. Their approaches derive from linguistic work on syntactic structures, particularly government and binding theory. While the Hobbs' and Lappin and Leass' algorithms perform reasonably well for intrasentential resolution (Hobbs 76% accuracy, Lappin and Leass obtain 86% accuracy; Lappin and Leass 1994), neither algorithm fairs well in resolving pronouns to antecedents intersententially. In evaluation of their studies which included a centering algorithm proposed by Brennan, et al. (1987), Lappin and Leass found that "attentional state plays a significant role in pronominal anaphora resolution and that even a simple model of attentional state can be quite effective" (Lappin and Leass 1994:552).

2.3.1 Centering Theory

From the studies of anaphora resolution and the deficiencies of the DRT framework, as well as the relevance imputed to models of attentional state in drawing reference in dialogue, centering theory appears to be a good choice for implementation. Centering theory maintains that each discourse entity in an utterance is a potential center in the discourse segment. However, only one entity may be the center for a single utterance. By ranking the salience of each entity in an utterance relative to the salience of entities in previous utterances, a center is found.

For every utterance U, there is a backward–looking center (Cb), a set of forward-looking centers (Cf) which comprise a Cf-list, and a preferred center (Cp). Given a sequence of utterances $\{U_i, .., U_{i+(n-1)}, U_{i+n}\}$, the Cb is the topic (or most salient entity) in the previous utterance (e.g. U_i given U_{i+1}). There can only be one Cb in any utterance. All entities in the previous utterance, including that which is tagged as the Cb, are listed as Cfs in the Cf-list. Depending on salience ranking functions for the Cf-list, each Cf is allotted a relative salience. After each Cf entity has been assigned the (newly) computed salience, the Cf-list is ordered highest to lowest. The Cf with the highest salience is then considered to be the preferred center (Cp) for the following utterance U_{i+1}.[5] The particulars to the basic centering algorithm are laid out in Table 2.1. below (from Grosz, et al. 1995, cited in Tetreault 2001:509). Consider, for example, the statement "Margaret eats dinner. She likes spaghetti." In the first sentence of the statement, the set {Margaret, dinner} constitutes the Cf-list with "Margaret" being the Cb. For the second sentence, {she, spaghetti} constitutes the Cf-list with "she" as the Cb. However, "she" also equals the Cp with respect to the first sentence. Here, the Cp of the second sentence is equal to the Cb of the first sentence.

Table 2.1: Constraints, Rules and Transitions of the basic centering algorithm

Constraints
1. There is exactly one backward-looking center (Cb).
2. Every element in the Cf-list for U_i must be realized in U_i.
3. The center of an utterance (Cb) is the highest ranked element of the Cf-list in the previous utterance that is realized in the present utterance.

Rules
1. If some element of Cf(U_{i-1}) is realized as a pronoun in U_i, then so is Cb(U_i).
2. Transition states are ordered such that a sequence of *Continues* is preferred over a sequence of *Retains*, which are preferred over sequences of *Shifts*.

[5] Preference is interpreted as a function of coherence.

According to most models of centering, the salience weight for an entity in the Cf-list is calculated by the grammatical role it performs in the utterance. The following relations among grammatical roles provides a general understanding for the order of salience based on syntactic structure: subject > direct object > indirect object > others.[6] Syntax-based algorithms of anaphora resolution (e.g. Hobbs 1978, Lappin and Leass 1994) demonstrate the significance of grammatical roles in determining the referent of an anaphor. The underlying reason for elevating grammatical role in salience ranking in centering comes from previous successes of syntax-based models of anaphora resolution.

Sometimes the Cb in the previous utterance U_{i-1} is not equal to the Cp in the present utternace U_i. Such cases mark a weakening of local coherence – a topic shift – and are referred to as transitions. Three transition states – *Continue*, *Retain*, and *Shift* – are proposed in the core model of Grosz et al. (1995). Transition states are defined in terms of the relationship between the Cb and Cp of the previous utterance and the Cb of the present utterance. Strube and Hahn (1998: 314) show the relationships governing local coherence from utterance to utterance. The *Continue* transition type is exemplified in the utterance sequence, *Margaret eats dinner. She likes spaghetti*, mentioned earlier. Here, 'Margaret' is the Cb of utterance U_{i-1} and 'she' is both the Cb and the Cp of utterance U_i. Strube and Hahn's alignment of transition types is reproduced in Table 2.2. Given our recognition of the Cb in U_{i-1} and of the Cb and Cp in U_i, the *Continue* transition (i.e. no transition) may be inferred.

Table 2.2: Centering Transition Types

	$Cb(U_i) = Cb(U_{i-1})$	$Cb(U_i) \neq Cb(U_{i-1})$
$Cb(U_i) = Cp(U_i)$	CONTINUE	SHIFT
$Cb(U_i) \neq Cp(U_i)$	RETAIN	

[6]It should be noted that the salience of relations among grammatical roles is "typical"; however, objects may be attributed greater salience through less typical structuring, such as by stress in speech (Hirschberg 1985) or italics in writing.

2.3.2 Variations on the Centering Model

Variations on the basic model of centering (Grosz, et al. 1986, 1995) have been performed in studies by Brennan, et al. (1987), Strube (1998), Strube and Hahn (1999), Tetreault (2001), Poesio, et al. (2004), and Kibble and Power (2004), among others. The first transformation of the basic centering approach appears in Brennan, et al. (1987; henceforth, "BFP"). In their centering approach, the *Shift* transition state of the centering model is subcategorized into *Rough Shift* and *Smooth Shift*, and that Retains are preferred over Smooth Shifts which are preferred over Rough Shifts. The BFP algorithm consists of three steps: (1) generate all possible Cb and Cf combinations, (2) filter combinations by contraindices and centering rules, and (3) rank remaining combinations by transitions. The BFP algorithm did not obtain best results for intrasentential pronoun resolution (Lappin and Leass 1994), and mildly underperformed when compared with other tests performed by Strube and Hahn (1999), Tetreault (2001), and Poesio et al. (2004).

2.3.2.1 The S-List

Strube (1998) proposed his S-List (i.e. a hierarchically ranked "salience list") to replace the need for backward-looking centers, believing that an ordering of the Cf-list on *functional* grounds (i.e. discourse function, rather than on purely syntactic, grammatical salience) provides better performance. Strube proposed this idea primarily to satisfy conditions posed by languages such as German, in which topicalization allows for movement of an object into subject position. The S-List is, theoretically, the most radical departure of centering since Grosz, et al. (1995). Essential to Strube's proposal is the categorization of entities in the S-List according to (1) *hearer-old*, (2) *hearer-new*, and (3) *mediated* discourse entities[7]

[7] It must be stressed that we are talking about entities, referring expressions, and anaphora, which manifest linguistically during a dialogue or discourse. These designations are only to be interpreted from the purview of communication between speaker and hearer in the attentional structure. On the other hand, Hearer-old/new etc, in the epistemological sense, ascribes these designations in terms of whether the individual interlocutor has a lexical semantic representation and prior understanding of the entity and (some of) its relations, maintained in the semantic module of his/her intentional

(this idea borrows directly from Prince's (1981) Familiarity Scale). For maintaining local coherence, the categories in the S-List are ranked. Old entities are preferred over mediated entities, which are preferred over new entities.

Hearer-old entities are classified as either "evoked" or "unused". Evoked entities constitute pronouns, definite referring expressions, and previously mentioned proper names. Unused entities are proper names that are known prior to the discourse. *Hearer-new* entities are classified as "brand-new," to the discourse and possibly previously unknown to the hearer. *Mediated* entities are comprised of "inferable", "containing inferable", and "anchored brand-new." An inferable is an entity that is linked, associated, or related to a hearer-old entity, but in which the "hearer is not expected to have in his/her head [during the present time in the dialogue] the entity in question" (Prince 1992:305). Ultimately, an inference is required to identify the antecedent of an inferable. In the statement "We saw a boat on the water. They waved to us," the addressee can infer that 'they' refers to some people on the boat, because 'boat', with its meaning as something containing people when driven on water, generates the possible referents. Containing inferables are indirectly inferable, involving an even less direct connection to the referent. Anchored brand-new entities occur in situations where a noun phrase contains a hearer-old modifier (anchor) for a brand-new head, as in the case of possessive pronouns.

This is important to the system proposed in this thesis because every possible implicature scale that emerges in a question–answer exchange needs to be identified by the system. Such identification is needed further for guiding inference procedures (i.e. processes of natural language understanding) as well as for content planning by the system in generating output from the results of inference. To exemplify this arrangement of the S-List model, consider the following sequence of utterances:

---

structure. Naturally, an entity – in the epistemological sense and represented in the intentional structure – originated at some time during a dialogue (or discourse, as with written language) where it was introduced to the hearer through some linguistic expression – a focusing mechanism – in the attentional state.

(5) a. Paula hates Texas.

b. She is a professor.

c. Her class provides a great learning experience.

In (5a), 'Paula' is a previously unused noun phrase. In the predicate nominal (5b), 'she' is the NP characterizing 'Paula' as evoked, and 'professor' is an NP that gets classified up to this point as unused due (in part) to the indefinite article that specifies it. In (5c), 'her class' is the NP, where 'her' is the evoked hearer-old anchor modifying 'class', which acts as the head of the NP. By utterance (5c), 'Paula' is no longer the center of the discourse segment, having been replaced by 'class'. The semantic association of professor in (5b) acted as a mediating entity – a containing inferable – which led 'class' to be inferred. Because of the evoked status of 'Paula' to which the possessive 'her' refers, 'class' gets projected to the top ranked entity in the S-List. From this example, the S-List approach theoretically works well for ranking possessive pronouns below personal pronouns as well as for inference making. [8]

Other centering approaches are unable to account for indirect realizations of discourse entities. Because of the attention the S-List algorithm pays to referring expressions, or entities, that are indirectly realized in a discourse segment (i.e. "functional anaphora" (Hahn, et al. 1996); inferables and containing inferables (Prince 1981)), it is particularly appealing. In order to enable a dialogue agent to understand utterances as input, it must be able to map the relationship a word can take in addition to those it does take in the particular context of discourse.

2.4 The Givenness Hierarchy of Cognitive Status

Taking a different tack from Centering Theory in the understanding and analysis of attentional state and reference resolution, Gundel, Hedberg, and Zacharski (1993) propose that

[8] In the case of inference making, the system should be able to understand the relationship between 'professor' and 'class' in order to hone the semantic frame of discourse segment, to conjure up meaningful situation based on these utterances.

the degree of reference (definite to indefinite) may be identified along a unidirectional implicational scale directly related to the cognitive statuses of an interlocutor. For them, an entity in discourse may be defined according to six cognitive statuses in the scale, reflected in Table 2.3 below (reproduced from Gundel, et al. 1993):

Table 2.3: Givenness Hierarchy

in focus	<	activated	<	familiar	<	uniquely identifiable	<	referential	<	type idenitifiable
it		*this that this* N		*that* N		*the* N		indefinite *this* N		*a* N

Each status of a concept represents the definiteness of it within an interlocutors memory structure as triggered by use of a referring expression for that concept in a dialogue. Depending on the linguistic form in which it manifests in an utterance, a referring expression can trigger the concept at any of the levels in the hierarchy. A status to the left entails all statuses to the right. However, any status to the right implicates not consistent with statuses to its left. For example, consider the status distinction between *he* and *the man*. If a referring expression falls under the uniquely identifiable status (e.g. *the man*), then it does not necessarily entail the in–focus, activated, and familiar statuses that may be attributed to *he*. So, while a concept that is familiar to an interlocutor with regard to the dialogue, it is also uniquely identifiable, referential, and type-identifiable to the interlocutor; however, that concept is not necessarily activated or in-focus in the dialogue, and therefore not necessarily activated or in-focus in the interlocutors memory.

Linguistic forms signaling each status are located under their respective statuses in Table 2.3. It is important to observe that interrogative proforms, forms which also denote a cognitive status of definite-indefinite reference, do not appear in the Givenness Hierarchy. Central to this thesis is the idea that they be incorporated into the hierarchy.

Gundel (1998), in an effort to unify the Givenness Hierarchy with centering theory, proposes to do away with the backward-looking center, or Cb. The primary issue of the Cb on her account involves resolving a pronoun to its antecedent which is not the Cb, i.e. is not most "in-focus," but is "activated" from what has been said. An entity could have an activated status by appearing in a subordinate clause or phrase.[9] The contention is that such an entity may also be appropriately referred to with a pronoun in a subsequent utterance, thereby evoking the entity and bringing it to an "in-focus" status in the dialogue (and into the Cb, or left-most position in the list). This can be seen in (6), where 'restaurant' is activated but not most salient in (6a). In (6b), it becomes most salient and is pronominalized as 'it'.

(6) a. $Margaret_i$ likes eating $dinner_j$ in that $restaurant_k$.

b. It_k offers great quality at a reasonable cost.

The centering algorithms, including the S-list, are not able to connect a pronoun occurring in such a subordinate phrasal position to that out-of-focus yet activated entity. Rather, the centering algorithms would incorrectly choose to resolve it to the most focused or "in-focus" Cb.[10] Thus, 'it' would resolve to either 'Margaret', or when assuming gender constraints, to 'dinner.' 'It' would not resolve to 'restaurant.'

Gundel, et al. (1993: 293) observe when this problem emerges in the course of a dialogue, "that the addressee accommodates and is able to associate the correct referent with the form in spite of the fact that it was used inappropriately." Without mentioning the exact form with which the accommodation (or repair[11]) occurs, it is obvious from their example[12] that the addressee fixes the reference by asking a question:

(7) K.1: Barb got it.

N.2: Catmopolitan?

K.3: Yeah.

N.4: Catmopolitan.

[9]The present S-list implementation strips out all relative clauses before sending the referring expressions contained in a sentence to the S-list algorithm.

[10]For Strube (1998)'s account, the same goes.

[11]See Levinson (1983).

[12]From the Frederickson tapes, reported in Gundel, et al. (1993).

In this dialogue, "Catmopolitan" is already activated in N's mind, and would likely exist already in the attentional model (e.g. S-list). This signals one way in which QA exchanges may be modulated. Because "Catmopolitan" already exists (and has an activated status) in the attentional model of the dialogue, a yes/no question seems like an appropriate method of repair. If we change around the utterance at K1, replacing "Barb got it" with "Barb likes games", we might expect a wh-question as a form of accommodation at N2 (e.g. "which games?").

In summary, work on scalar implicature by Grice (1975), Horn (1984), Levinson(1983, 2000), and Hirschberg (1985) provide an understanding of lexical semantic contrast and implicature scales evoked during dialogue. The works of Gundel, et al. (1993) and Gundel (1998, 2003) are significant because they provide a referential framework from which wh-questions and interrogative pronouns can be espoused to *activate* lexical semantic contrast within implicature scales. As such, their work lends support to the implicational properties of interrogative pronouns. Works on centering theory, particularly that of Strube's (1998) S-list offer a means of implementing attentional state, and hence cognitive activation from expressions occurring in dialogue to representations in an interlocutor's memory.

Chapter 3

System Overview

Dialogue management systems[1] vary widely depending on the intended application and audience. The system proposed here attempts to enrich a user's lexical knowledge with respect to semantic hierarchies that enable implicature. This should be accomplished by retrieving lexical scalar sets from a lexical semantic resource (e.g. WordNet) and introducing them in the dialogue. Target competencies of the system include: 1) tracking discourse topics throughout dialogue, 2) resolving anaphors to their antecedents both intrasententially and intersententially, 3) resolving interrogative pronouns to their answers (or referents), 4) conceiving new topics for conversation, 5) planning propositions around such topics, 6) planning sentences for those propositions, and 7) generating text accordingly. With these competencies, the Q/A system may be considered fairly autonomous in its conversational interactions. Processing competencies (1), (2), and (3) can be accounted for using the centering algorithms in concert with computational models of spreading activation and lexical decision.

3.1 Purpose of the Study

Motivation for this thesis stems from overlapping interests in computational dialogue processing, question answering systems, interactive computer-assisted language learning (ICALL), second language acquisition, and theoretical and applied pragmatics. These interests coincide with the desire to make a robust, automated, question answering system that promotes English language acquisition by non-native English speakers. Several components

[1]Jurafsky and Martin (2001) present an overview of dialogue managers and some example applications.

of the system were developed and stand as a starting point for the design and development of a more complete system. While this endeavor has demonstrated the difficulties of dialogue processing, it has also highlighted the many possibilities for computational modeling of dialogue and the human-computer interface that realizes such dialogue. In particular, pragmatic modeling and analyses of intentionality in utterances seem tantamount to any effective performance of such a system. In the rest of this section, I provide an overview of the system and briefly discuss its various components, which will prove useful later in establishing a frame of reference and means of understanding the ideas presented throughout the thesis. The thesis then seeks to illuminate potential improvements, espousing theoretical revision to some of the components in the system and proffering additional components.

The Q/A system developed thus far enables users to pose one or more English sentences as input, which are accepted, processed, and responded to. Numerous classes and procedures constitute the internal processing of a user's input. First, the input is cut into sentences, which are designated by the usual English sentence boundary punctuations: ".", "!", "?". Each sentence is passed through the system one at a time for analysis.

The first stage[2] of sentence analysis is tagging, a process by which words are assigned to the various parts of speech (e.g. noun, verb, etc.). Parts of speech are represented in a tag set. The system employs an amalgamation of the UCREL C7 tags, Biber tags, and some homemade tags that were created in order to handle specific phenomena encountered through the development process. A formal tagger is not employed, per se, but a list of lexical items are stored locally alongside their possible tags (ranked in order of likelihood of use) and the lemma form of the lexical item. This list is then referenced during tagging and subsequent

[2]Actually, the first stage should involve a spell-checker as well as, perhaps, morphological analysis. While these are important to any CALL system, they are considered negligible for the central focus of the thesis: scalar enrichment.

procedures. For words not contained in the list, the tagging procedure simply assigns the ordered list {*nn, ..., in*} to the word and takes the word itself to be the lemma.[3]

The parser[4] then performs a syntactic analysis of the sentence and derives a tree structure characterizing the sentence. The syntactic analysis is crucially dependent on a grammar, which delineates the set of rules by which a sentence can be parsed. The grammar of this system is an underspecified, probabilistic context-free grammar (PCFG). It is underspecified because the target user is assumed to lack the necessary competence for specific uses of the full-range of English grammar. Specifically, the grammar is underspecified for feature agreement, since the user may commonly be expected to enter utterances with feature disagreement (e.g. "I eats cake").

After parsing, sentences are passed to a feature assessor which checks to see if the agreement criteria for subject/verb, singular/plural, and subject/object pronouns are maintained. If one or more agreement criteria are not met, then the system fixes the agreement errors, logs a note of what kind of error was encountered, and prepares a feedback string (which is only added to the final output after further processing occurs).

The next process involves fact extraction, or what is sometimes called text simplification (Siddharthan, 2002). Fact extraction is simply the process in which sentence structures are converted to logical form. So, for example, assume that the sentence "The kid the story describes likes games" is entered. At the fact extraction phase, this sentence is converted to the following predicate-argument structure:

Surface Form:	"The kid the story describes like games."
Logical Form:	`describe(simplepres1, story_nn, kid_nn),`
	`like(simplepres1, kid_nn, story_nn).`

[3] More sophisticated (and accurate) taggers exist including MontyLingua (Liu, 2004) and the Stanford POS Tagger (Toutanova and Manning 2000; Toutanova, et al. 2003). These work with the Penn Treebank tag set.

[4] The specific parser implementation used here is the chart parser described in Russell and Norvig (2003).

During fact extraction, referring expressions are also extracted and sent to the centering procedure, which maintains and updates information regarding the attentional states of the human-computer discourse. The attentional state procedure relies on Strube's (1999) S-list algorithm. Via the S-list algorithm, pronouns are resolved to their antecedents. The S-List is crucial to what comes later in the thesis because it functions as the departure and destination points for the process of spreading activation.

The term, spreading activation, is common to psycholinguistic literature and characterizes the cognitive process wherein a network of semantic relations between lexical items (lexemes) is primed or activated and subsequently used to search for associative links to other potentially relevant lexemes. There are generally two stages in spreading activation: automatic spreading activation and selective activation.

The spreading activation module implemented in the current system is most basic. For a term ranked as most salient in the S-list (and therefore in the discourse segment), the term's n^{th} orbit of lexical semantic associates is collected from the WordNet lexical semantic database (e.g. given "girl", then the hypernym list {*woman, person, human, mammal*} would be returned) . Finally, a random lexical item is selected from the list returned from spreading activation and is held for later inclusion in the response text.

Following centering and spreading activation, the logical form of a sentence (LF) is submitted to the knowledge base (KB). In the case of a declarative sentence, its LF is inserted into the KB unless the KB already holds this predicate structure. In the case of an interrogative sentence, the KB is queried with the LF where either (1) arguments in the predicate are supplied by a question word (i.e. a *wh-question*) or (2) the structure is an interrogative that does not contain any question words (i.e. *yes-no questions*). [5] The response of a query results in (a) affirmation, (b) denial, or (c) a declaration of ignorance (e.g."I don't know"). Note that in the case of interrogative procedure, unification of arguments in answer predicates with variables in query predicates constitutes a type of spreading activation. In effect,

[5] These are marked only by the "?" punctuation; no question words accompany the "?" in the input.

the spreading targets a more precise sentence-level activation rather than just term-level activation. In the cases of (a) and (b), upon unifying the arguments of a predicate query, spreading activation on the arguments within the predicate may be extended. In the case of (c), spreading activation can occur prior to accessing the knowledge base in order to expand the possible predicate structures that might map to the predicate query.

Finally, results from the KB are submitted to a natural language generation (NLG) procedure. The NLG procedure converts any LF to surface form. Then it takes any grammar checking feedback and spreading activation lexeme and outputs all three or any subset thereof as the response to the input.

From this description of the current system under development, I draw particular attention to the modules governing (a) attentional state, (b) spreading activation, and (c) the knowledge base. For these modules lie at the heart of the thesis. The attentional state procedure tracks discourse topics in question–answer exchanges. The values or entities contained in the attentional state model are expanded through spreading activation. Both of these are necessary for anaphora resolution. They tie together entities expressed in dialogue with lexical knowledge contained in an individual's mind (or the computer's lexical semantic database, i.e. WordNet). [6] Meanwhile, the knowledge base serves to store sentence-level knowledge that consists of the combination of different term-level arguments.

In the current implementation, the system has no means of initiating questions. It merely attempts to map an answer to a question posed by the user. Therefore, additional modules that need to be designed and implemented include (d) an intitiative grammar or plan reasoning module, (e) a nonmonotonic, lexical semantic reasoner, and (f) a user model. These are surmised in the thesis conclusion.

[6]This is described in details in chapters 5 and 6.

Chapter 4

Question Words in Dialogue

Question words are deictic, which means that their reference functions indexically according to the contexts in which they appear. Interrogative pronouns are the most generalized form of reference; they are indefinite descriptions (Bhat 2004) and may be applied (and resolved) in much the same way that anaphora/pronoun resolution occurs. The difference however is that the attachment of reference comes from some speaker-external knowledge base (like the internet, another interlocutor, a news broadcast, etc). For example, if a speaker asks "Who is the tallest person in the world?," then 'who' must be resolved to some referent, a person, who is not yet uniquely identified. The speaker expects to obtain the resolution from an external source, hopefully, from the addressee. Although question words may not typically be thought of as anaphora or pronouns, in fact they are extremely anaphoric; they represent the longest-range form of coreference in the language; they must traverse an interactive exchange in order to be identified with their referents.

Question words fall into two categories: pronominal and adnominal. Pronominal question words (henceforth, interrogative pronominals) include *who*, *what*, *when*, and *where*, which syntactically function like a phrase, and most commonly a noun phrase.[1] An adnominal is any word which modifies a noun, such as *green* in *green car*. Ad-jectives, for instance, are adnominals. Possessive pronouns are also adnominals (e.g. *her* car). The set of interrogative pronouns which constitute interrogative adnominals are *which*, *what*, and *whose*.

[1]Note that derivational forms of verbs, such as gerunds, can function as nouns by being interpreted as events.

4.1 Interrogative and Relative Pronouns

Some interrogative pronouns are isomorphic to some relative pronouns. These include *who*, *which*, *where*, and *when*. The implicational use of these words are commonly use in two types of construction: (1) appositive constructions, and (2) interrogatives or questions. Appositives are essentially relative clauses or subordinate clauses. They introduce some added content to the interpretation of a discourse entity. Take, for example, (8):

(8) The kid who told the story likes games.

In (8), "who told the story" is an appositive construction; it expresses additional contextual (background) information for the way in which "the kid" is to be interpreted. From a functional perspective, the word "who" acts anaphorically for its antecedent "kid". The example above can be broken down into two basic statements:

(8)′ The kid told a story. *He[who]* likes games.

The revised version (8)′ aims to demonstrate that "who" in appositive constructions functions identically to that of personal pronouns. Appositive use of question words function more simply like coreferential pronouns. - i.e. they mark coreference.

4.2 Interrogative Proforms

On the other hand, interrogative proforms in questions make disjoint reference to entities not yet in the discourse but corefer to more general entities in the addressee's memory. Interrogative pronouns function much like indefinite articles do in declarative utterances. Consider the following:

(9) Who likes games?

(9)′ Q: Who told the story?

A: The kid [who likes games].

4.2.1 Cataphora

Interrogative pronouns are cataphora. A *cataphor* is an expression which corefers to something expressed later. For example in "Eventhough he ate three meals today, Sam was hungry again at 8pm." Here, 'he' corefers to 'Sam', which appears after 'he.' Interrogative pronouns operate in almost exactly the same way. The only difference is that they are intersentential and intra-discoursal. In a model of attentional state operating over a dialogue, what was once a disjoint reference in $(9)'$ Q (viz. U_{i-1}) becomes a point of coreference in A (viz. U_i)).

4.2.2 Anaphora

In contrast to cataphoric reference, anaphora exist to make coreference to entities previously mentioned in a dialogue. However, the case for treating interrogative pronouns as anaphora is supported along several lines of evidence. Bhat (2004) lends credence to this notion by showing cross-linguistically how questions words and question phrases mirror indefinite pronouns. He frames his argument in terms of a reference continuum where anaphora are situated toward endophoric or exophoric poles. Although not explicitly stating that interrogative pronouns may be treated anaphorically, Strube (1998) includes the class of relative pronouns among the coreferential forms of discourse entities.

This interpretation is further supported given the type-identifiable and referential cognitive statuses of the Givenness Hierarchy. For at least the interrogative proform can be intermediately resolved by the addressee as referential to a specific type (e.g. *who* = person, *where* = place, *when* = time, *what* = thing). Consequently, the functional aspect of interrogative proforms may be established. It is important to note that this theoretical interpretation be made because it supports the underlying procedures by which the system resolves question words to their referents and uses degrees of definiteness in bounding the scope of lexical scales (Hirschberg 1985) that are targeted for enrichment.

4.3 Functional Aspects of Interrogative Proforms

When interrogative pronouns function as nominal specifiers, potential implicature scales emerge for the word that is being specified (e.g. *which game*). Question words extend scales of quantity (viz. Horn Scales) and of quality (viz. Hirschberg's alternate, non-scalar sets). In mutual comparison, question words comprise a set of alternates which specify for person, place, thing, or time (i.e. quality). This may be represented in the alternate set {*who, where, what, when*}, where the set as a whole functions to qualify the future quantity scales of interest (that which is in-focus) in the dialogue. For example, "who" qualifies "person(s)" instead of "time", which is qualified by "when." From each in this scale of question word alternates, scales of quality are not extended all at once per se, but each induces an anchor on the weak end of a scale which qualifies the scalar orderings (i.e. quantity) of its possible answers in the least as type-identifiable. In the case of adnominal uses, they can range on the scale from referential to in-focus. For example, "who" delineates a scale beginning with "person" at its weakest end and some more specific type or subset of people toward its stronger end:

<Principal Jones, school principal, principal, administrator, {*which, what*} *person*, *who*>

By this interpretation, interrogative proforms coincide with the Givenness Hierarchy of cognitive status for referring expressions at the level of type-identifiable. In adnominal uses, it is important to note that interrogative proforms function more elaborately - they contract or reshape the cognitive status scale in accordance with the entities to which they modify. For while an entity may be familiar and activated to an addressee, as in the context of *those cats there that you and I can both see and are consciously discussing*, interrogative proforms may be used appropriately to further specify and bring into focus only one of the *cats* for a particular description.

The functional aspect of interrogative proforms leads, in effect, to the traversal of implicature scales of reference. The scales below represent how an oscillation occurs through a series of questions in an exchange, where the interrogative pronominal anchors the scale at type-identifiable status and the adnominals suffice to move down it toward the stronger end of uniquely identifiable.

<{*what,which*} time, {*what,which*} day, . . . , {*what,which*} year, *when*>

<John's car, {*whose*} car, that car, {*which, what*} car, a car,
{*which, what*} vehicle, a vehicle, *what*>

In this chapter, I have described the two kinds of interrogative proforms, pronominal and adnominal. I provided reasons why these forms should be interpreted as anaphora. Then I provided insights into the functional aspects of interrogative proforms, highlighting the mechanism for scalar traversal. In the next chapter, I will discuss a model of attentional state that is necessary for coordinating scalar traversal with question words.

Chapter 5

Modeling Question Words in Attentional State

Comparison of centering approaches and non-centering approaches to anaphora resolution and topic tracking in general have motivated an implementation of centering in the software proposed by this thesis. In this chapter, I describe components of the system developed for modeling attentional state in question answering exchanges. I begin by demonstrating how the model operates on declarative constructions and how functional anaphora in those constructions are referentially resolved. Then I enhance the model to accommodate another form of anaphora – interrogative proforms. An operational description of the system at the attentional level will then prove instrumental in espousing a system design for scalar enrichment.

5.1 Implementing Pronoun Resolution

The pronoun resolution algorithm implemented in the current system is drawn from Strube's (1998) S-list algorithm. The input to the algorithm is a list of lists, called the referring expressions list (RefExpr). Each sub-list is a 4-tuple data structure that holds (1) a referring expression (the head of a phrase) extracted from the user's input, (2) the lemmatized form of the referring expression, (3) the part of speech tag associated with the referring expression's use in the user's utterance, (4) the word order position of the expression in the user's input, and (5) the entire phrase in which the referring expression was encountered. The referring expressions list is passed to the S-list algorithm and its constituents are included in the attentional state model.

The basic S-list algorithm is reproduced below in Table 5.1. In step one, the algorithm attempts to resolve pronouns before adding them to the list of salient discourse entities. To do so, it looks at each entity in the S-list from left to right in accordance with the S-list ranking criteria for coherence (or givenness), and compares that entity with the pronoun. Resolving the coreference of the pronoun requires that it examine feature criteria of the entity and match it to the features of the pronoun. These features include gender, number, and person (binding constraints are only relevant if the pronoun is reflexive). For example, consider "John went to the store to buy a mop, because he had spilt the milk on his rug." Here, {John, the store, a mop, he, the milk, his rug} constitutes the set of discourse entities on which the S-List will operate. 'John' is a proper noun not preceded by a determiner and is therefore marked unused. 'The store' is marked by a definite – not an indefinite – article and therefore is also marked unused. 'A mop' is marked by a indefinite article and is therefore brand-new. 'He' is a pronoun so becomes evoked. 'The milk' is not marked by an indefinite article and so is unused. 'His rug' is marked by the possessive pronoun 'his' and therefore 'rug' is a mediated entity classified as anchored brand-new.

Table 5.1: The S-List algorithm

For each utterance ($U_1 \ldots U_n$): for each entity E_x within U_i:

1. If E_x is a pronoun, then find a referent by looking through the S-List left to right for one that matches in gender, number, person, and binding constraints. Mark entity E_x as EVOKED.
2. If E_x is preceded by an indefinite article, then mark E_x as BRAND-NEW
3. If E_x is not preceded by a determiner, then mark U_i as UNUSED
4. Else mark E_x as ANCHORED BRAND-NEW
5. Insert E_x into the S-list given the ranking described above.

To satisfy gender constraints, lists of the 300 most common English male and female names[1] are stored locally on the computer. These are searched to determine the gender of the referring expression currently under inspection by the algorithm[2]. This is necessary because the part-of-speech tag set (UCREL C7; see Jurasfsky and Martin, 2000) lacks this gender detail. For entities where a proper name will not suffice, as in the case of "president", then the system performs a look up of the term in another lexical semantic resource, namely WordNet, to determine if it represents a person. If the non-personal pronoun, "it", falls under resolution, then the S-list entity searches WordNet to contrast the entity with words that maintain the "person" distinction. If none is found, then "it" defaults to expletive "it" and is discarded from the algorithm. Possessive pronouns follow the same resolution procedure of their personal and non-personal counterparts. "I" and "you" are most easily resolved as "user" and "computer" (or, if you like, "Chewbacca" and "C3-PO"). Because plurality is encoded in the tag set, plural pronouns are compared with the tags of S-list entities for resolution.

Let's look at a simple example to illustrate what is going on in the model of attentional state. In (10) below, two declarative statements are asserted in sequence. The first is "John bought a car." The second is "He likes it." The referring expressions extracted from the first sentence are "john" and "car."[3] These are represented in the referring expressions list at line 2. Next, the information status of each referring expression is determined. Because "john" is a proper noun, it is classified as U, or unused. "Car," on the other hand, has an indefinite article modifying it, which denotes a discourse-new entity. Thus, "car" is deemed BN, or

---

[1]http://names.mongabay.com/

[2]Admittedly, this process could be speeded up by encoding the gender specification directly into the S-list data structure, but it is not so slow as to keep th euser waiting for a response.

[3]The inevitable destination of these discourse entities is the Prolog knowledge base. Prolog interprets capitalized forms as variables. To get around this, all entities which refer to objects in the world are adjusted to lowercase. We'll see how these circumstances change with regard to interrogative ponouns in the next section.

brand-new. It is important to note that this system does not specify hearer-old except by reference to the current discourse.[4]

(10) 1 *John bought a car. He likes it.*

2 RefExpr: [[john, john, np, 0, john], [car, car, nn, 3, a car]]

3 S-LIST: [[U, john, john, np, 0, 0, John], [BN, car, car, nn, 4, 0, a car]]

4 LF: `buy(simplepast1,john_np,car_nn)`

5 RefExpr: [[he, he, pp3a, 0, he], [it, it, pp3it, 2, it]]

6 S-LIST′: [[E, john, john, np, 0, 0, he], [E, car, car, nn, 3, 0, it],
[U, john, john, np, 0, 0, John], [BN, car, car, nn, 4, 0, a car]]

7 LF: `like(simplepres1,john_np,car_nn)`

Now it is time to add the entities to the S-list. "John" is added to the S-list and then "car." This throws "car" directly on top of "john!"[5] To be more precise, it lands "car" directly to the right of "john." The relative positions of "john" and "car" in the S-list are irrelevant at this point, because after entities are inserted into the list, the list is sorted (Strube 1999). First, entities are sorted by information status per the hierarchy of Old>Med>New (via Hahn 1996 and Prince 1981). Next, they are sorted by the utterance in which they appear, ranking all entities in the most recent utterance U_n (or 'he' and 'it' in "He likes it.") higher than those in previous utterances and positioning all entities within it on the left pole of the list. It renders the most distant entities utterance U_i last and on the right pole. For example, U_{n-1} (or 'John' and 'car' in "John bought a car."). Finally, entities are sorted by their relative positions within an utterance, with earlier words in an utterance ranking higher than later words. The S-list at line 3 reflects that the S-list update has occurred. Line 4 displays the first sentence in logical form as it is presented to the knowledge base.

[4]In fact, if an entity is hearer-old, then it would be represented in the knowledge base. It is also constructive to consider that series of discourse models between an interlocutor and the computer could be stored and used in its own knowledge base for picking up where their last dialogue left off. I will return to this idea at the end of this thesis where I consider future directions.

[5]The list structure functions as a stack, which is usually conceived of with vertical spatial orientation.

The second sentence begins to be processed at line 5 where the referring expressions “he” and “it” are processed. “He” is recognized as a singular personal pronoun of male gender. So the algorithm starts looking through the S-list left to right in search of an entity that matches the criteria. Because “john” is tagged as np (i.e. proper noun), the system identifies a partial match with the criteria of person. Also, because it is a name (i.e. proper noun), the system looks at the lists of 300 male and 300 female names to determine gender. The entity “john” in the S-list is matched with the occurrence of “John” in the male names list. Therefore, there is no need to search any further, “john” has been matched with “he”. The system then creates another entry for “john” in the S-list. Since “john” in this entry emerged from the resolution of a pronoun, the information status is defined as E, or evoked. The rest of the information contained in the referring expressions sub-list for “he” is added, rendering [E, john, john, np, 0, 0, he]. The same procedure follows for “it” except that the lexical resource is a list of non-human entities. Finally, the S-list is updated based on the *sortal* criteria of discourse coherence mentioned earlier in the section.

Although it would be prudent here to begin another trace, one of possessive pronouns, space does not permit. It should be mentioned however, that possessive pronouns, as adnominals, determine an evoked information status for the entity they modify. This topic is covered in detail in section 5.2.2.

5.2 Interrogative Proform Resolution

Shifting the focus to interrogative pronoun resolution, recall that interrogative pronouns occur as interrogative pronominals and adnominals. In the subsequent two sections, I explain how these two classes of interrogative forms are engaged by the attentional model of dialogue through their implementation in the S-list.

5.2.1 Interrogative Pronominals in the S-List

Discussed in chapter 4, interrogative pronominals are distinguished at the linguistic level by the forms *who*, *what*, *when*, and *where*. As their classification suggests, they are pronouns that refer to indefinite conceptualizations of entities including people, non-human things, places, and times. One distinction between them and their declarative counterparts is that interrogative pronominals corefer and make disjoint-reference, embodying a conjunction of disjunctions. Definite pronouns, unlike interrogative pronominals, can only corefer, and specifically, only to items already in the current discourse.[6] In the formulation presented here, this does not alter the necessary information status of the interrogative pronominal with regard to the discourse model. It does, however, (a) require that the discourse model representation be modified to accommodate them, and (b) reverse the direction from which they are resolved from backward looking (as with definite pronouns) to forward looking.

(11) *Who drives the car?*

S-List: [[E, Who, who, whp, 0, 1, who], [U, car, car, nn, 3, 1, the car]]

LF: `drive(simplepres1, Who, car_nn)`

Example (11) illustrates these changes made to the original model in order to handle interrogative pronominals. The example provides the sentence for the model to process, the S-list entity description lists (in ranked order), and the logical form (LF) of the sentence. In (11) the S-List algorithm takes in the referring expressions "who" and "car" from the sentence. "Who" is identified by the system as an interrogative pronominal, and therefore no attempt whatsoever is made to resolve "who", (viz, [E, Who, who, whp, 0, 1, who]). The interrogative pronominal is treated as an entity, in and of itself. A declarative pronoun, in contrast, never appears in the second position in the entity description list - it is always resolved or else discarded. Upon resolution, however, a declarative pronoun does upgrade the information status of its antecedent to evoked (if not already evoked) in the givenness hierarchy of discourse. Interrogative pronominals are always evoked. While this can be explained by a

[6] This can also be described as a dialogue session.

thorough examination of the intentional nature of questions in dialogue (e.g. see chapter 4), instead, I defer to Strube's (1998) specification that relative pronouns and appositives always be evoked.

It is difficult to get a full sense for how the algorithm resolves interrogative pronominals without drawing into focus their referents - i.e. their answers. Table 5.1 below provides all of the necessary information for a proper trace of the algorithm. In Table 5.1, the human asks the computer, "Who is the president of America?," to which the computer replies, "George Bush is the president of America." The processing steps described in example (11) as well as those described in the previous section on pronouns are performed to bring the human's utterance (the question) into the attentional state. Thus, the referring expressions 'who', 'president' and 'America' are extracted (see *RefExpr* in Table 5.1). Next the information status of each is assessed, and then each entity is added to the S-list. The last step in the algorithm sorts the S-list. What happens next is crucial. The utterance processing leaves the attentional structure of the dialogue, and enters the intentional structure where all things known to the system are stored and can be reasoned through (i.e. the computer's memory and inference system). This happens by passing the logical form (LF) of the question,[7] `be(simplepres1, Who_who, president_npt)`, to the logic engine as a "query" to determine if an answer is known.

5.2.1.1 Conferring the Knowledge Base

The logic engine is represented as a knowledge base (KB) of Prolog rules (predicates) and clauses (entailment conditions). The rules and clauses are searched over in an attempt to unify the question word with another item that matches the present tense specification for the predicate, `be(simplepres1,Who_who,president_npt)`, or more simply, is(Who, president). (It is important to note that the system creates a duplicate LF in cases of predicate nominals, such as in (2). The arguments in the second and third positions of the duplicate

[7]Harabagiu (2006) terms this the "QLF"

Figure 5.1: QA exchange trace

H: *Who is the president of America?* (question)

RefExp: [[who, who, whp, 0, who],
[america, america, np, 5, America],
[president, president, npt , 3, the president]]

LF1: `be(simplepres1, Who_who, president_npt)`
LF2: `be(simplepres1, president_npt, Who_who)`

S-LIST: [[E, who, who, whp, 0, 1, who],
[U, president, president, npt, 3, 1, the president of America],
[Ic, america, america, np, 5, 1, of America]]

C: *George Bush is the president of America.* (answer)

LF: `be(simplepres1, 'george bush_np',president_npt)`

RefExp: [[george bush, george bush, np, 0, who],
[america, america, np, 6, America],
[president, president, npt, 4, the president]]

S-LIST$_i$: [[E, who, who, whp, 0, 1, who],
[E, president, president, npt, 3, 1, the president of America],
[Ic, america, america, np, 5, 1, of America]

S-LIST$_j$: [[E, george bush, george bush, np, 0, 1, who],
[E, president, president, npt, 3, 1, the president of America],
[Ic, america, america, np, 5, 1, of America]

are then reversed in order. The reasons for doing so will soon be apparent). Because the system has the lexeme, George Bush, as an argument in a predicate, `be(simplepres1, president_npt, 'george bush_np')` and `be(simplepres1, president_npt, 'george bush_np')`, the system is able to unify the question word "Who" with the argument 'george bush_np' for the matched predicate. [8]The logic engine then provides a response or answer in LF to another procedure that extracts the referring expressions from the LF (e.g. RefExpr under C: in Table 5.1) and sends them to the S-List. Another procedure then transforms the LF answer response into English for output.

Returning to the attentional structure, the referring expressions of the system's utterance (or, sentence) are sent to the S-list.[9] The algorithm is already primed for an answer and is looking for a match between the incoming expressions and its expected answer type, or EAT (Harabagiu 2001, 2006). In this case, "who" is the anaphor in question, and its EAT specifies for a singular or plural noun of personhood. The first referring expression in the RefExp list, ["George Bush", np, 0, George Bush"], comes into play. It has been tagged as a proper noun (i.e. "np"), which the algorithm recognizes as one of the part-of-speech tags that either (a) defintiely denote personhood or (b) possibly denote personhood. Thus, with agreement criteria met through (a), "George Bush" replaces "who" at the top of the S-list.

Before moving on to resolving interrogative adnominals, the story behind possible denotation and definite denotation needs further clarification. Part-of-speech tags are typically set up to make syntactic or surface-level distinctions between words. No tag sets to my knowledge encode gender, for instance. With who, personhood is a requirement in order for the interrogative pronominal to coreference the entity that is focused in the answer. The tags that definitely denote personhood include proper names, titular proper nouns (i.e. titles), and the plurals of both. Tags possibly denoting personhood include general noun tags, such as "nn" and "nns" (singular and plural, respectively). When the system runs into one of

[8]The process of inference described here is extremely simplistic.

[9]The before and after states of this are illustrated in Table 5.1 under the computer's response.

these before it has encountered any of the tags in the set of definite denotation, it has to search other lexical resources, such as WordNet.

For example, suppose that instead of “George Bush”, the system knows only that “a Christian fundamentalist” is president. The noun, “fundamentalist”, would have to be recognized as either a person or not a person. After finding the index in WordNet (3.0) for “fundamentalist”, the algorithm needs to climb up the classification hierarchy four levels to find that, indeed, it is a person.[10] This same analysis and look-up process is performed for each kind of interrogative pronominal.

5.2.2 Interrogative Adnominals in the S-List

An adnominal is any word which modifies a noun, such as ‘green’ in *green car*. Question words that function as interrogative adnominals in English are the words *what*, *which*, and *whose*[11]. Adnominals further specify the identity and context of the nouns they modify, e.g. *what* day, *which* president, and *whose* car. Representing interrogative adnominals in dialogue is essential to an effective dialogued-based question-answer exchange.[12]

Interrogative adnominals appear where (a) a discourse already exists and inside the discourse, (b) entities are already being attended to. Consider the following exchange:

(12) A: The kid likes games.
B: Which games does he like?
A: (i) Video games.
(ii) He likes video games.

[10]To help the process, a morphological analyzer might be implemented to recognize that any w ord suffixed with “-ist” is a person (holding some established belief). But a morphological analyzer would not work for the word “kid”.

[11]How could be interpreted or translated into “in what way” or “by what process” . Similarly, why could become “for what reason”. These question words are not addressed here, but are mentioned in the thesis Conclusion.

[12]Later we will also see their greater impact on the intentional structure of discourse, but for now, let’s turn to their realization in the S-list.

In this dialogue exchange, A tells B that a kid – whom B already knows[13] – like games. B asks A to tell him a more specific description regarding the games that the kid likes. A answers B with the more specific description (either the elliptical response (i) or the more expressive (ii)). Now look at the interaction with the S-list traced after each turn, including the logical form that results from the resolution algorithm.

Figure 5.2: Tracing: supposition, question, answer

A: *The kid likes games.* (supposition)
(LF) like(simplepres1,kid_nn,games_nns)
S-LIST: `[[E, kid, kid, nn, 0, 2, the kid],`
[U, games, game, nns, 0, 3, games]]
B: *Which games does he like?* (question)
(LF) `like(simplepres1,kid_nn,kind_of(Which,games_nn))`
S-LIST: [[E, which, which, which, 0, 0, which games],
[E, games, game, nns, 0, 2, which games],
[E, kid, kid, nn, 0, 3, he],
[U, games, game, nns, 1, 3, games]]
A: (a) *Video games.* (answer)
(LF) `like(simplepres1,kid_nn,‘video games_nns’)`
S-LIST: [[E, video games, video game, nn, 0, 0, which games]
[E, games, game, nn, 0, 2, which games],
[E, kid, kid, nn, 0, 3, he],
[U, games, game, nns, 1, 3, games]]
(b) *He likes video games.*
(LF) `like(simplepres1,kid_nn,‘video games_nns’)`
S-LIST: [[E, video games, video game, nn, 0, 0, which games]
[E, games, game, nn, 0, 2, which games],
[E, kid, kid, nn, 0, 3, he],
[U, games, game, nns, 1, 3, games]]

Immediately following the first turn taken by A, the S-list reflects that “kid” is the most salient item in the discourse. However, B responds with a question that seeks more

[13]Due to the definite article “the”, then “kid” is hearer-old via Prince (1981).

information regarding an item in the list, namely "games". This reflects a transition in the discourse structure, prompted by B's preference for more information.[14]

Of particular interest is the question presented by B and the responding answer provided by A. In B's question, "which" is separated from "games" and both are given an evoked status. As with declarative possessive pronouns and relative pronouns in appositive constructions (viz. "the store, which is located on the corner"), interrogative adnominals take an evoked status. Strube (1998) classifies possessive pronouns as evoked and so should interrogative adnominals be evoked. However, recall the example associated with the S-List algorithm of Table 5.1, *John went to the store to buy a mop, because he spilt the milk on his rug.* Here, "rug" takes on the anchored brand-new, mediated status. Therefore, one might think that 'games' in *which games* should be anchored-brand new, too. 'Games' is justifiably evoked when B's question enters the S-List because 'games'already exists in the S-List as unused. By referencing again with the interrogative adnominal, it becomes elevated in focus to an evoked status. Where A maintains the focus of attentional state on 'kid', B shifts it to 'games'. Finally, when the question is posed by B to A and the attentional state has been updated according to B's question, then the implicature of B's question (which is conventional) is passed to A's intentional structure for interpretation.

What makes B's question interesting is that its logical form involves an embedded predicate. The embedded predicate denotes a "kind-of" relation which holds between the lexemes serving as the predicate's arguments. Any interrogative adnominal strives to be bound to some more specific type of the noun it modifies. The noun it modifies delineates the expected answer type with a more definite description of its own class of concepts.

The system can go about finding the "kind-of" relation for the lexeme in a variety of ways, but the ones proposed here include (1) processing through the knowledge base for a match to the "kind-of" predicate conjoined to its container predicate like, or (2) searching WordNet for the argument, "games" (in its lemmatized form, "game"), and returning a

[14] Which will enable B to better cooperate in the current dialogue with A and possibly future dialogues with others.

more specific lexeme of the concept "game". In the case of (1), the system already has represented the concept: `like(kid, video games), kind_of([computer game, video game], game).` So it can immediately deduce that video game is a kind of game that the kid likes. On the other hand, strategy (2) mines WordNet for the "kind-of" relation and selects one for response. Strategy (2) seems very infelicitous though. Imagine asking someone where a grocery store is, and – just to save face in the immediate context – s/he tells you it is along some (randomly chosen) path. That wouldn't necessarily serve the goal you intended by starting the Q/A exchange. So while the first strategy seems reasonable, albeit tedious, the second does not seem useful.

Hirschberg (1985), for her system QUASI, proposed to use both strategies in concert. She suggested mining a lexical classification database (e.g. WordNet, although it wasn't developed yet) for specific lexical semantic relations and filling up the KB during the process. Later, at runtime, all of the necessary semantic knowledge would be available without any look-up.

Let's assume that strategy (1) is undertaken by the system for processing the logical form of B's question in example 12. A match is found, and the predicate `like(simplepres1, kid, 'video game_nn')` is passed back to the attentional structure. Then the referring expressions 'kid' and 'video game' get sent to the S-list algorithm.

Inside the S-list procedure, the interrogative adnominal 'which' is waiting to match up the expected answer type criteria of the 'game' with one of the referring expressions. Next, 'kid' comes into play. A look up in WordNet of "kid" under 'game' is rejected. The procedure moves to the next word, 'video game.' Again, a look up is performed for 'video game' under 'game', and is successful. So 'video game' replaces 'which' in the S-list, taking on its evoked status. The S-list is sorted and 'video game' comes up as the focus of discourse.[15] This is important because is it illustrates how interrogative proforms can be used you traverse lexical semantic scales of the same concept – the basic procedure for enriching pragmatic knowledge at the level of lexemes.

[15]Even though 'kid' is already active and evoked in the attentional state.

In this chapter, I presented centering theory and studies motivating its use as a computational model of attentional state. By comparing its instantiation in the BFP and the S-list algorithms, I determined that the S-list appeared right for implementation in the present system, largely due to its emphasis on resolving functional anaphora. Then I examined the inner-workings of the S-list as it resolves declarative pronouns (the base case), interrogative pronominals, and interrogative adnominals. In the next chapter, I describe how this computational implementation of attentional state can be extended and used to augment the lexical semantic structures of a user's pragmatic knowledge.

Chapter 6

Augmenting Intentional Structure with Question Words

In this chapter, I examine how the model of attentional state described in chapter 5 interfaces and interacts with the intentional structure of the software application. This intentional structure consists of networks of lexical meaning and constitutes the knowledge base of the application. The latter structure can be broken down into two components: the initiative or goal-oriented component and the knowledge component. In section 6.1.1., I define the system goals as they related to dialogue. In section 6.1.2., I describe how lexical knowledge[1] is represented in the knowledge base of the question answering system. This includes the WordNet hierarchy. In section 6.2. I describe how lexical knowledge structures interact with the S-list attentional model through a sequence of contextualized lexical activation. In section 3.3., I characterize how question words are used by the application in concert with its goal-specification to enrich the contextualized lexical knowledge of the user.

6.1 Components of Intentional Structure

Intentional structure consists of lexical knowledge, the ability for inference, various cognitive statuses, and a system of planning and initiative in communication. In this section, I discuss system goals and the representation of lexical knowldge in the system.

6.1.1 System Goals

Important and inherent to any interlocutor is the sense of preference that s/he can assert in an environment - be it in dialogue, or in narrative. Underlying preferences are several

[1] Terminological note: lexical knowledge refers to the representative existence of a lexeme and its relation to other lexemes

motivating factors, initiatives, or goals. As Hirschberg notes, "speakers' goals can determine what is relevant or salient in a discourse" (Hirschberg, 1985: 164).

At least some goal gets communicated when a person utters a statement or a question in dialogue. It is what drives the model of attentional state and lies at the heart of task-oriented dialogue and of collaborative agents. For example, *Where was Franz Kafka born?*[2] conveys the act of information-seeking and the goal of resolving the question word to its referent. It also conveys the broader goal of understanding the name of a place in a contextualized relation to the words *Franz Kafka* and *born*. Knowing a priori only that the word, *Prague*, exists in relation to say, *Czech Republic* and *city*, leaves the goal unattended, particularly if the goal is to continue the dialogue by being cooperative with (or helpful to) the interlocutor. This also entails being relatively cognizant of his/her goals, as by recognizing that a question seeks an answer. It is important because an unanswered question may be viewed as uncooperative or irrelevant by the questioner unless an inventive, cooperative reply is made to substitute for the answer (Levinson 1983: 293). An answer of silence may be interpreted by the questioner similarly as uncooperative or as motivation to continue, or take the ground again in the exchange.

Only two goals (although reactions may be a better description) are specified for the agent: (1) answer questions based on the informational content in its knowledge base, and (2) introduce lexical items that are metonymically entailed in WordNet by focal entities in the discourse model. The first represents a reaction specification. The second, on the other hand, is an initiating action, an introduction of focus content. The second goal is functionally aimed at generating contextualized relations between lexical items in a discourse while concurrently strengthening the scalar opposition between metonymically entailed entities. This may be accomplished by having the system simply use an answer response pattern of the form word_2 is a kind of word_1, as in example (9) from the last chapter in which, assuming the computer agent is speaker A, the computer replies to "Which games does he like?" with the answer,

[2]From the TREC 2004 QA track data set (TREC-13), Question 22.1

"A *video game* is a kind of *game*." In effect, the agent's initiative is designed to simulate a process, which Levinson (2000: 238) describes as, *metonymic particularized transfer*, or MPT.

MPT takes those lexical items which have fairly standardized, conventional, or generalized[3] relational scales of opposition (Horn (1984); e.g. hot/cold, administrator/person, minivan/car, etc.) and makes them particular to the dialogue context and other lexical items with which they occur. The simulation is based on the idea that a series of questions about the same topic can hone in on the generalized scale, and thereby enrich the user's vocabulary relating to those scales. As these scales become more apparent to the user, then the user will acquire a better competence for using them to convey generalized conversational implicatures. A side effect of the process could be that the system transfers the metonymic scales from WordNet (i.e. generalized) into its knowledge base as both generalized and particularized, provided it interacts with a trainer first. This would, in effect, prime the system for certain scales and propositons, allowing for a more robust abductive reasoning component.[4]

6.1.2 Lexical and Scalar Representation in WordNet

The lexical semantic database, WordNet, serves as the primary system for representing the conceptual relations of lexical items in the agent. In WordNet, lexical items are organized hierarchically to reflect gradable relations between words of the same, more general concept. This organization allows for anaphora resolution, as shown in the previous chapter. It also enables dialogue agents and question answering systems to compute intensions and implicature (Hirschberg, 1985; Harabagiu, 2006).

WordNet maps lexemes to one another through metonymy relations. Metonymies are logical relations that hold between the concepts that lexemes signify. The full list of relations is available in the WordNet Glossary,[5] but the major players include synonym, antonym,

[3]To use Levinson's preference of terms

[4]The abductive, or pragmatic, reasoning component will be speculated in the Conclusion.

[5]http://wordnet.princeton.edu/gloss

hypernym, hyponym, meronym, holonym, attribute, troponym, and instance-of relations. *Synonyms* are lexemes in symmetrical relationship to each other. Together, they form an equivalence class. The words, ‘conveyance’ and ‘transport’, are synonyms. They mean the same thing in some context. Synonyms function as a alternate set (Hirschberg, 1985).[6]

Antonyms, in contrast, are lexemes denoting opposites. For example ‘hot’ and ‘cold’ are antonyms. Antonyms form the polar ends of a scalar set (e.g. <*hot*, warm, tepid, cool, *cold*>, Hirschberg, 1985).

Hyponyms and *hypernyms* represent two opposing directions of the specific-general scale. A ‘computer game’ is a hyponym of ‘game,’ and ‘game’ is a hypernym of ‘computer game.’ Thus, the scale <computer game, game>.

Meronyms and *holonyms* are espoused to be part of X, member of X, or substance of X relation types. These correspond with Hirschberg’s (1985) part/whole, set/member (or set/subset), and type/subtype relations. For example, ‘ear’ has the part-meronym ‘ear drum’, while ‘ear’ is a part-holonym of ‘auditory system.’ ‘Team’ has the member meronym ‘stringer’[7] and ‘stringer’ has the member holonym ‘team’. ‘Tobacco’ has the substance meronym ‘nicotine’, and ‘nicotine’ has the substance holonym ‘tobacco.’

An *attribute* is defined as “a noun for which adjectives express values. [For example] The noun weight is an attribute, for which the adjectives light and heavy express values” (WordNet Glossary, http://wordnet.princeton.edu/gloss).

Troponym relations are specific to verbs of the kind *doing X entails doing Y in a specific manner*, as in whispering entails speaking very quietly and with limited vocalization. As such, we can represent the scalar set <whispering, speaking> on a scale of <strong, weak> or <specific, general>.

The *instance-of* relation characterizes a definite-indefinite description scale at the strongest end of the specific-general scale. For instance, Jimmy Carter is an instance-of

[6]However, they can be made into a hierarchical scale by way of frequency of occurrence in some context. This will be elaborated later.

[7]A member of a squad on a team. e.g. ‘a first stringer’; ‘a second stringer’ (from WordNet 3.0)

the concept, president. This relation can be represented as < Jimmy Carter, president>, or on an extended scale of specific-general, <Jimmy Carter, president, person>.

Finally, *glosses* constitute another means of calculating implicature scales. For the glosses of words in WordNet help not only in answering definition question (e.g. "What is an atom?"), but also in identifying important lexemes used in creating the definition. This may be necessary for some scalar sets since instances of spurious metonymic relations do exist in WordNet (Priss, 1998) and some relations that are not explicitly denoted through one of WordNet's metonymy structure.

Although the current system does not explicitly make use of glosses in mining implicature scales, it could be easily made to do so, especially by using Extended WordNet[8] (Harabagiu 2001). Since lexical knowledge can be thought of in terms of semantic structure. For example, the definitions of a lexeme, or the lexeme's gloss, needs to be converted into logical form. Harabagiu (2001) has provided extended WordNet to do so.

6.2 Extending Attentional State through Spreading Activation

As illustrated in the previous chapter, when a user submits a declarative or interrogative utterance and inherently expresses context, salience, and various referents, the computer performs an attentional state update, by extracting the referents in the utterance and incorporating them into the S-list as discourse entities. Recall also that in order to resolve pronouns such as *he* to antecedents like *the kid*, the S-list algorithm consults lexical semantic resources for the type-identifiable description of the target term (antecedent) that matches the feature criteria for the pronoun (e.g. person). The same is done for resolving interrogative pronouns that appear in the S-list. For who, the computer looks in WordNet and the proper names files to see if its referent, or answer term, is type-identifiable as a person. The details of WordNet's organization were laid out in the previous section. What this section aims to show is how this internal organization of lexemes in WordNet becomes connected

[8]http://xwn.hlt.utdallas.edu/

to the system goals and dialogue through the model attentional state. The first step in this process is to identify sets of lexical items, which together and under some ordering, make up a gradable concept scale (e.g. <specific, general>). This is accomplished through spreading activation.

6.2.1 Spreading Activation

Spreading activation denotes the cognitive process in which networks of semantic associations become active under certain stimuli, or priming actions. In language interaction, it characterizes the process by which one word or set of words enters the cognitive apparatus of the hearer, where the semantic concept for that word is located, and then propagates an association pattern to concepts (and their representative lexemes) that are pragmatically related to that concept. In psycholinguistic studies looking at the processibility constraints of particular kinds of linguistic forms, the lexical decision task (LDT) is often used to gauge the categorical and conceptual distance between words. This distance is marked by a decision on-set delay characterized by the time interval it takes a respondent to select a word, given a conceptual relation to map and a lexical prime. The spreading activation of an LDT is typically construed along two phases: an automatic phase and a selectional phase. Both phases are simulated in the agent.

Extending the attentional state model of the agent to the intentional structure and knowledge base invokes spreading activation. This is realized when the S-list algorithm performs a WordNet look-up in order to match type-idenitifable criteria in anaphora resolution. However, with a more robust model, each word in the S-list becomes activated (Gundel, et al. 1993; Gundel, 2003; cf. Givenness Hierarchy). Spreading activation is what actually follows from an "activated" entity in the attentional state.[9]

Figure 6.1 shows the first orbit of spreading activation for the word "story" in WordNet 2.0. It propagates out from "story" to show all of the words (and concepts) with which

[9]It should be noted that here I address the automatic phase of spreading activation. The selectional phase will be described in the following section, Identifying Basic Level Categories.

"story" is immediately associated. Along one of the hypernym paths, we see that fiction becomes activated. Along another hypernym path, news is activated. And along yet another hypernym path, the synonyms, report and account, are co-activated. So it follows for the rest of the metonymic relations immediately associated with the word "story". Empty sets are also shown in order to present where no path for activation exists. For "story", this includes antonym relations.

Figure 6.1: 1st orbit of spreading activation for *story* in WordNet 2.0

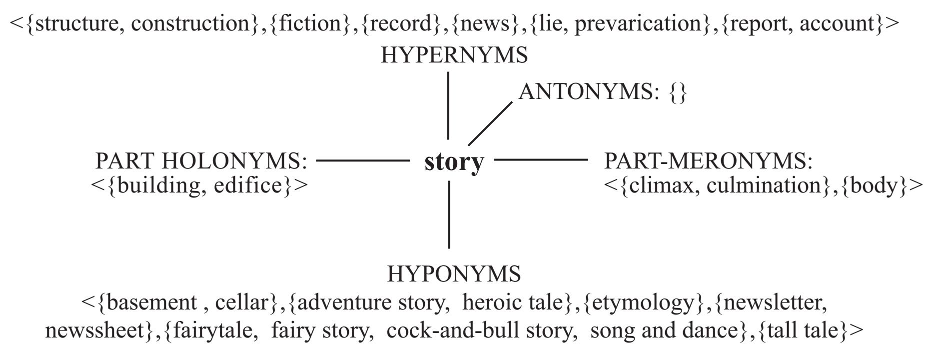

Spreading activation is a seminal aspect of the agent's ability to translate intention, for it connects entities in the attentional state model to the knowledge and intentional structures. The procedure for spreading activation is fairly simple: take the salient items in the S-list and expand them to activate the hearer-old (or, computer-known) items[10], i.e. metonymically related lexemes stored in WordNet. This process has been performed by a number of computational linguists including Hahn et al. (1996) and Harabagiu (2006). For example, in *The kid the story describes likes games*, the word 'story' activates out to its hyponyms, including the kinds of stories 'fairy tale' and 'adventure story' as well as its part-meronyms such as 'climax' and 'culmination.'

Previous researchers have made a point to rank the orbits of activation (Hahn et al, 1996; Harabagiu, 2006). Hahn, et al. propose various criteria for determining the usefulness

[10]Not necessarily discourse-old, but hearer-old in terms of internal memory

of metonymic paths in resolving issues of textual ellipsis. They characterize WordNet relations as either plausible, metonymic, or implausible. The first orbit denotes plausible paths. Extended orbits are considered metonymic. While cyclic paths in spreading activation are deemed implausible.

Strube (1998) via Hahn, et al. (1996) report that meronym and hypernym/hyponym relations logically entail, which was an initial reason for choosing Strube's S-list for attentional state management in the current system. However, it seems that other relations provided in WordNet also function in a logical entailment. For synonyms map the relation of reflexivity while antonyms map antisymmetric relations. Holonyms may be interpreted as whole/parts relations, and so are subsumed under what Strube's definition of meronym. Note that beyond "logical entailment", Resnik(1998: 260) points to "plausible entailment" as "those properties that any reasonable speaker would infer of something by virtue of its membership in that class." He utilizes glosses for obtaining such additional plausible entailments. He gives the example of {*plank, board*} having the plausible entailment of 'flat' eventhough it is not linked through the attribute relation. The synset for {*plank,board*} carries no metonymic path to 'flat'; however, another synset of "board" entails the word through its gloss. It seems that a re-construal of the synonym relation for 'board' (as a homonym, deemed viable through a probabilistic calculation from frequency data).[11] Resnik's approach points to glosses as another way for salient entities in discourse to activate other words.

Harabagiu (2006) also uses glosses to extend lexical activation in automatically answering complex questions. For example, her system analyzes the information request, *Biographical information needed for Elizardo Sanchez, Cuban dissident*, by breaking down the word 'biographical' into into its nominal form 'biography.' She then uses the term, 'life,' found in the gloss of 'biography,' as the target concept for extrapolating the individual, simple questions

[11]It may be constructive to mention that each word encountered in WordNet spreading activation should be weighted based on frequency of exposure and use, as well as on the frequencies of collocations from any genre or any broad level description reflected in a corpus. More on this to come.

contained in the example complex question. In the case of the target concept, 'life,' the constituent concepts include 'birth,' 'death', and 'origin' among others. These lexemes are then correlated with their interrogative proforms 'when' and 'where' followed by distinct analyses for each in relation to the other terms in the query, such 'Elizardo Sanchez.' Through this example, one can see that glosses serve a significant purpose in selectional spreading activation.

6.2.2 Identifying Scalar Sets through Spreading Activation

With the automatic phase of spreading activation triggered, a more selectional phase must be instituted in order to identify scalar sets. Levinson (2000: 107) contends that if a scalar set emerges in a particularized manner, then it may range indefinitely. Take for instance the scalar sets involved in process stages (Hirschberg, 1985), where each process stage is a value in the scalar set. These may incur an infinitely continuous set when particularized. Consider the concept of reincarnation, where there is birth, life, death, birth, life, death, and so on (so-called helical relations of time). Now consider that each instance of life carries a more specific value such as *cat* and later *human* and still later *ant*. Holding down the range in which the life subset has some bound appears difficult. Levinson remarks that Hirschberg's (1985) particularized theory of scalar implicature carries this effect. When the use of a set is particularized, it can instantiate to any two concepts, and therefore lexemes, whether or not the words have a generalized ordering, say from metonymic relations. This, of course, enables language speakers to implicate any of the vast number of creative oppositions which commonly can be found in absurdist humor or surrealistic literature.[12] In particularized contexts, the concepts are realigned, in effect, pulled away from the general, standard scales enabling English implicature. By argument of the discrete nature of lexical storage in computer memory as well as the discrete temporal reality of dialogue sessions, my claim is that dialogue constrains the range a scalar set can take. Hirschberg's (1985) analysis supports

[12]Attempts to models such oppositions are beyond the scope of the current project, but may be taken up in the future.

this claim. For the first argument, the metonymic relations in WordNet *are the scales* with which the broader community of English speakers use to produce generalized conversational implicature. For these relations have been derived from enormous balanced corpora representative of the generalized contexts in which lexemes appear and contrast with each other. WordNet only holds a certain discrete number of items. The computational complexity it large, but not infinite with regard to the number of combinatorial permutations any lexemes can take. To the second and stronger argument, the generalized nature of the scales is either reinforced or mutated[13] over time through their use in particular discourse and dialogue at a particular time.

This still leaves, however, the issue of how to fill up the generalized scales for use a particular dialogue. Consider that, given the word, “car”, the range of hypernyms alone extends through 11 levels and 19 words. A host of other terms are identified as well for each word’s sister terms[14] (e.g. “thing” is a sister of “entity”; “motor vehicle” is a sister of “car”) and for all of the other relations that accompany the spreading activation of “car”. It is computationally expensive to retrieve every word and to delineate it in an implicatural scale. Furthermore, it is ineffective if the terms retrieved for scalar enrichment are of little use to the user or if such words themselves will adversely mark the user’s speech[15]. Thus, the system must prefer certain words in the semantic hierarchy over others - scales must be pruned and bounded. Hirschberg (1985) proposes that this can be accomplished by preferring basic level categories on the one hand, and utilizing terms currently held in the discourse model on the other. These two lexical items constitute the poles and directionality of a scalar set.

[13]Mutated maintains Levinson’s interpretation of *particularized*

[14]Sister terms are terms that are “immediate hyponyms of the same superordinate”; http://wordnet.princeton.edu/gloss

[15]Less useful words activated from “car” include “instrumentality, instrumentation”, “conveyance”, and WordNet’s top-most general semantic concept, “entity”.

6.2.2.1 Identifying Basic Level Categories

Research on classification hierarchies conducted by Berlin and Kay (1969), Rosch (1973), and Lakoff (1987), among others, led to an understanding of how the brain processes words in relation to the experiential salience with which people encounter the words and, more importantly, the concepts that the words describe. These studies also show how cultures raise the salience of certain categories into perceivably more useful contextualizations, and hence lexicalization. The most salient categories in a lexicon are commonly described as *basic level categories.*

> BASIC categories have been defined by Rosch (1973, 1976) as those categories in classification hierarchies which carry the most information; are most differentiated from other members of a hierarchy at their level; are the most inclusive categories whose members have a significant number of attributes in common; have similar motor programs and similar shapes; and can be easily identified from the average shapes of members of the class" (Hirschberg 1985: 157).

By the basic level category schema, lexical items are stratified across four hierarchical levels: superordinate, intermediate, basic, and subordinate. These levels provide insight into how a dialogue agent can select certain lexemes over others for inclusion in a scalar set.

For instance, "minivan" has the hypernymic structure appearing is Figure 6.2. It contains no hyponyms, which means that it is the subordinate level category. And its hypernyms range more generally to 'entity.' This is useful for a scale of <specific, general>, in which the subordinate term fills the slot in the scale denoting the most specific term. Filling in the rest of the scale involves only a little more work. The hypernym "car" is the next level up. Most native English speakers will agree that the word 'car' is a fairly basic level category. It describes a specific object in the world which can be distinguished on its own from other objects due to its relative shape, motor program, etc. Next in line is 'motor vehicle,' which seems, without and further elaboration to lexemes further up the hypernym structure, to

Figure 6.2: WordNet 3.0 hypernym structure of *minivan*

minivan
car
motor vehicle, automotive vehicle
self-propelled vehicle
wheeled vehicle
vehicle container
conveyance, transport
instrumentality, instrumentation
artifact, artefact
whole, unit
object, physical object
physical entity
entity, thing

denote a basic category. However, proceeding further still, 'vehicle' appears. And so 'vehicle' also may be assumed basic. This presents a problem for identifying basic level categories.

Fortunately, Green (2006) presents a two-phase methodology for obtaining the basic level categories from WordNet. The second phase, which I wish to adopt, describes several heuristics that act as scoring criteria for lexemes in WordNet. Following these heuristics, the software can be made to accommodate Hirschberg's use of basic level categories.

The first heuristic is based on the length of a lexeme. It states that "if a word is longer than 15 characters, it is unlikely that it names a basic level category." We can see from Figure 6.2 that 'self-propelled vehicle' fails this criteria. So we exclude it as a basic category. Second, "if a lexical unit is a phrase [including] two or more words," then it is not a basic term. This eliminates 'motor vehicle,''automotive vehicle,' 'self-propelled vehicle,' 'wheeled vehicle,' and 'physical object.' Third, "if the name of a concept is included within the name of a more specific concept," then the name shared by both the more specific and more general concept probably signals a basic category. Therefore, 'vehicle' is weighted more for selection than are 'motor vehicle,' etc. Frequency of occurrence defines the fourth heuristic, where a word contained in a word frequency count over a large set corpora is weighted more than its direct hypernyms and its direct hyponyms, then the more frequent word is weighted higher and the less frequent words are eliminated. Also, if the concept expressed by a lexeme contains two or more part-meronym relations, then it might be a basic level category. Moreover, lexemes can be further weeded out by identifying superodrinate and subordinate level categories. If there are no hyponym relations for a lexeme's synset, then it is probably representative of the subordinate level. In the case of the 'minivan' example, no hyponyms emerge. So 'minivan' can be considered a subordinate category. Identifying superordinate lexemes may be accomplished by seeing whether the lexeme has "more than four [hyponym] levels beneath it" (Green, 2006). In the 'minivan' example, any hypernym level above 'wheeled vehicle' could thereby be a superordinate. Using these heuristics, it may be deduced that 'car' is basic, 'minivan' is subordinate, 'vehicle' may be basic or superordinate (i.e. intermediate), and any of the set {conveyance, transport, instrumentality, instrumentation, artifact, artefact, whole, unit, object, entity, thing} may be superordinate.

6.2.2.2 Bounding the Scope of Scalar Sets through Discourse Salience

In addition to using basic level categories in the identification of scalar sets, evoked[16] discourse entities may be used to bound the scope of a scalar set to either side of a basic level cateogory (Hirschberg 1985: 160-164). This constitutes Hirschberg's entry level for set bounding. she proposes that

> scalar implicatures will not be licensed 'above' (at a more specific level than) the basic level or above a more specific evoked level (Hirschberg 1985: 160).

This is critical because, as noted by Levinson, the range of a scalar set could go on indefinitely. However, a salient discourse entity, drawn from the S-List, can provide a set bound. In so doing, it also provides a degree of contextualization to the process of identifying the lexical

[16] Hirschberg finds that salient discourse entities define "entry level" lexical items: "the entry level of a classification hierarchy establishes its 'appropriate level of detail' (Hirschberg, 1985: 161)".

items in the scalar set. In terms of the Givenness Hierarchy (Gundel, et al, 1993), the cognitive status of the referring expression that emerges in the S-list may be assessed. Moreover, the directionality of the scale will be effected.

6.3 Storing Scalar Sets in the Attentional Model

Following the S-list update, the agent sends the discourse entities into WordNet for spreading activation. The spreading activation and lexical decision processes then select specific lexemes for potential addition to the discourse. But where should those lexical items be held and under what ordering should they be made available to the system? It is in dealing with these lexical items that the S-list fails, for it cannot project or impose future discourse topics. It merely keeps track of those that have already been expressed. A structure is needed to house the items the system has preferred and maintains the initiative to introduce. Hirschberg (1985: 190) recognizes the utility of storing salient scalar sets in relation to dialogue, but does not implement it in QUASI since there is no initiative or plan established.

6.3.1 The Intentional Centers List

I propose a data structure called the Ci-list, or intentional centers list, to store returned scalar sets of activated, system-preferred items. The Ci-list data structure is a 3-tuple, or set of three fields, consisting of (1) the position in the S-list where the Ci-list was triggered, (2) the name of the semantic relation (e.g. general-to-specific or hyponymic) which orders (3) a list of lexemes delineating the scalar set of the trigger term including synonyms of the trigger term. This structure is illustrated in Figure 6.3.

The first field in the Ci-list, $Index_S$, simply points to the location in the attentional state where the Ci-list was triggered. For example, given the S-List *[[U, kid, . . .], [U, game, . . .]* where 'kid' is in position 0, and 'game' is in position 1, then $Index_0$ will link the Ci-list for 'kid' to its S-List representation and $Index_1$ will link the Ci-list for 'game' to its S-List representation. The second field in the Ci-list denotes the metonymic relation that was used to order the lexical items in the scalar set. It may take any of the WordNet metonymic relations except for synonymy.[17] Finally, the last field in the Ci-list (the list of terms related to the trigger term) is further subdivided into a 2-tuple consisting of a frequency count for a lexeme (if available) and the lexeme itself. This list of 2-tuples is important because it carries the ordered scalar set derived from basic level category heuristics. The Ci-list as a whole represents the priming effect for the system in raising scalar sets to the level of potential focus dialogue. Incorporating items into the Ci-list during dialogue is the first step in structuring a procedure for introducing those items and there scalar inter-relations to the user. With further procedures (including a content planner), then lexical acquisition, lexical contexualization, and thus, potential scalar enrichment may be offered to the user.

Figure 6.3: C_i-List

$$[\text{Index}_S,\ \text{Relation}_r,\ [[\text{Freq}_i,\ \text{Lex}_i],\ [\text{Freq}_j,\ \text{Lex}_j],\ \ldots,\ [\text{Freq}_n,\ \text{Lex}_n]]]$$

6.4 Traversing Scalar Sets with Question Words

In chapter 4, I presented the notion that question words function to traverse lexical items in implicature scales. The dialogue agent will use question words in order to traverse implicature scales and introduce values in those scales to the user. In this section, I illustrate how a possible software implementation of this traversal may be obtained through the S-List and Ci-lists. To demonstrate, consider again example (12) from chapter 5 about which games the kid likes. In doing so, imagine that speaker A is the computer and speaker B is the user.

(12) A: *The kid likes games.*
B: *Which games does he like?*
A: (i) *Video games.*
(ii) *He likes video games.*

[17] Synonyms are all stored together with their trigger term.

The computer first expresses the content "The kid likes games." This presents the suppositional information, 'the kid' about which speaker B should already have some activated memory. Following the statement, the S-List fills in with 'kid' and 'games.' The Ci-lists for each are also filled which includes each word's hypernym scale. For 'kid,' the Ci-list fills with the scale <kid, juvenile, person>. The 'games' Ci-list fills with <game, activity, act, event>.[18]

Table 6.1: S-List and C_i-List supposition convergence

Computer: *The kid likes games.*

S-list		C_i-List
index	lexeme	
0	*kid*	[hypero,[[53,[kid]], [_,[juvenile, juvenile person]] [6833, person, individual, someone,...]]
1	*game*	[hypero,[[[53,[game]], [43,[activity]], [24,[act,...]], [62,[event]]]

Speaker B (the human) responds to the statement by asking "which games" the kid likes. At this juncture, 'which,' 'games,' and 'kid' enter the S-List. The interrogative phrase 'which games' is separated and 'which' is entered with an evoked status. Because 'games' is the nominal specified by 'which,' it gets raised to an evoked status. Similarly, 'kid' enters as the antecedent of 'he' and so becomes evoked as well. Due to the S-List's favoring of earlier sentence positions in ordering discourse entities, 'which' becomes the most salient item in the list followed by 'games' and 'kid.' Table 6.2 illustrates the order of salience in the S-List after speaker B asks the question. It is important to note that the Ci-list for 'which' is empty, represented as [**?**,**?**].

Speaker B's question prompts the computer to extend the referential scale along the specific-general relation (because it is a hypernym-hyponym relation) associated with 'game.'

[18] Also, [[53,[game]], [8,[contest, ...]],[1,[social event]],[62,[event]]] emerges, but due to space limitations, this list is not contained in Table 6.1.

Table 6.2: S-List and C_i-List question convergence

Human: *Which games does he like?*

S-list		C_i-List
index	lexeme	
0	*Which*	[hypero,[[**?**,**?**]], [53,[game]]],
1	*game*	[hypero,[[[53,[game]], [43,[activity]], [24,[act,. . .]], [62,[event]]]
2	*kid*	[hypero,[[53,[kid]], [_,[juvenile, juvenile person]] [6833, person, individual, someone,. . .]]

Since 'which' AND 'game' are both elevated to an evoked status in the attentional state due to the character of interrogative adnominals, 'which' carries with its activation the nominally specified, 'game.' In effect, the interrogative adnominal triggers the next most salient word in the S-list . Both words provide a bounding mark for the scale. 'Which' represents the entry level while 'game' represents the basic level category bound.[19] Together, they carry out the instruction for finding a more specific lexical item in the hyponym path of the basic level category, 'game.' As mentioned, Table 6.2 illustrates that the Ci-list mapped to 'which' holds the empty list, [**?**,**?**]. The empty list for 'which' remains this way until its referent is matched.

The computer then finds the set {video game, computer game} and selects the one with a higher frequency of occurrence given a corpus.[20] The selected word 'video game' then replaces 'which' in the S-list. This is demonstrated in the first row of Table 6.3. Also note

[19]In cases of interrogative proniminals such 'where,' the basic level is 'place' or 'location.'

[20]The frequency count for this synset is not available in WordNet. Therefore, frequency data from some corpora would need to be assessed. Moreover, in situations such as this one, where a much larger set of subordinates are found, frequency data, again, determines which lexemes are selected for inclusion in the list.

Table 6.3: S-List and C_i-List answer convergence

Computer: (a) *Video games.*

S-list index	lexeme	C_i-List
0	*video game*	[hypero,[[_,**[computer game, video game**]], [53,[game]]], [43,[activity]], [24,[act,. . .]], [62,[event]]]
1	*game*	 [43,[activity]], [24,[act,. . .]], [62,[event]]]
2	*kid*	[hypero,[[53,[kid]], [_,[juvenile, juvenile person]] [6833, person, individual, someone,. . .]]

that the entire scale for 'video game' is represented. By mapping the entire scale for 'video game,' the scale itself may be considered active. Finally, the computer returns 'video games' to answer the human's question.

This thesis proposes that, through such a process of providing a user with more specific lexical items in response to questions containing interrogative proforms, an enrichment of the user's implicature scale occurs. Ideally, the computer should know, for example, which games the kid actually likes (as stored knowledge in the KB). But in the case of not actually knowing, the system can present an answer that gets a little closer to information the user may be able to use. It is imagined that if the user is a student of English as a Second Language (ESL), then presenting stronger terms along whichever relational axis is available can act as an automatic vocabulary tutor by providing the more specialized service of offering new, potentially unknown and more specific items to the user's vocabulary. Such items, when encountered during the question answer exchange, promote the lexical contrast inherent to pragmatic reasoning. Subsequently, the user may find increased exposure to the implicature scales in context more rewarding and more insightful for understanding what, as well as the extent of which, native English speakers implicate during fluent discourse.

In this chapter, I have presented components of a dialogue agent's intentional structure, including system goals and the lexical knowledge that it contains. By connecting the lexical knowledge of WordNet to the attentional state model through the Ci-list, I illustrated how lexemes in memory can be actualized in a question answer exchange. Finally, I demonstrated how system recognition of interrogative proforms in a user's question are handled and may be used to traverse sets of lexical items which are useful for interpreting and producing conversational implicature.

Chapter 7

Conclusion

From the theories and practice investigated in this thesis, a dialogue agent that uses interrogative proforms to fix reference and mine implicature scales is feasible, but complicated. With an attentional state model (i.e. the S-List), it resolves interrogative anaphora while tracking discourse entities among turns in dialogue. This was illustrated in the example question-answer exchanges in chapter 5, showing that both interrogative pronominals and adnominals are handled with an augmentation to the S-List algorithm. Thus, a computational approach to handling the functional aspects of interrogative proforms is feasible, and in fact, implemented.

Representing and activating sets of lexical items that share semantic relations is an essential ingredient to designing and developing a system for scalar enrichment. I have suggested how lexical semantic scales can be represented by the agent in an intentional centers list (i.e. Ci-list). When the S-List is connected to the Ci-list, items in the S-List activate lexical sets contained in WordNet and bring them into the Ci-list based on basic level category constraints. In chapter 6, I have detailed these processes.

With information contained in the knowledge base, the S-List and Ci-list work well together in structuring a simple question answering system. The traversal of scalar sets is entirely possible given the proposed framework. However, several areas need to be designed, developed, and tested before it is satisfactorily autonomous as a dialogue agent that promotes scalar enrichment. These areas of limitation include: initiative management, user modeling, pragmatic reasoning, and complex question answering.

7.1 Future Directions

Initiative management would enable the system to operate on a more sophisticated arrangement of goals and plans for interacting with the user. Enabling the system to propose certain topics and propositions from a plan of interaction that is mutually shared with the user could make the system more engaging, for it could introduce new domains of discussion and avoid its beating a dead horse, so to speak. Moreover, the user may not presently have the English lexical knowledge to pursue other topics. Through initiative management, the agent could make topic transitions without relying entirely on the user. Grosz and Sidner (1990) describe a system for initiative management that embodies a "shared plan" infrastructure. Incorporating their work into the present system may help in developing this capability. Moreover, through machine learning techniques, such as Bayesian Networks, the agent could be made to determine which topics are more likely to be of interest to a user as well as to connect those potential topics to a current dialogue session.

User modeling offers the capability of tailoring dialogue to an individual user based on the user's knowledge of the language, personal interests, and needs as a language learner. As mentioned, machine learning techniques could be used to estimate user interests. They could also be used in determining the language knowledge and size of a lexical set with which the user is familiar. More importantly, user modeling requires that the system be able to represent the beliefs, knowledge, and activities of many users simultaneously. At present, the dialogue agent assumes there to be only one user. Designing the agent to represent multiple users requires a new knowledge base for each user as well as a way of tracking or indexing the lexical scales with which the user is familiar. This addresses the issue of scaling up the system to larger audiences.

Furthermore, enhancing the Prolog knowledge base with an abductive, or non-monotonic, system of inference making could lend a great deal of functionality to the agent. Presently, the agent makes deductive inferences to match answers with questions. However, to understand as well as to generate conversational implicature in statements, a procedure for cancellation of certain semantically stronger or alternate lexemes in a given lexical set needs to be implemented. For example, if the user says, “I rented a car,” then the agent should not assume that the user rented a minivan. Perhaps an agent response to this statement could be, “Why didn’t you rent a bus?” given the {car, bus} alternate set in which *bus* is cancelled by the use of *car*. Similarly, if the user says, “I rented a minivan,” then the agent should assume that the user rented a car, since *minivan* entails *car* though a hypernym relation. Thus, further development of non-monotonic inferential operations could automate pragmatic reasoning and enrich the dialogue even more.

Finally, the agent needs a way of understanding and generating complex questions, sometimes expressed through the question words *how* and *why*. Such questions require a trace of the inference procedure performed by the Prolog logic engine (i.e. the KB). Answers to “how” and “why” involve collecting propositions of a whole chain of logical entailment, both at the level of propositions and at the term level. For example, if the user asks the system “Why are you so happy?” then agent might be able to answer “because I am not sad,” due to the contrast it represents between *happy* and *sad*. For a more satisfactory answer, the agent might reason that chatting with the user makes the agent happy and if the agent is chatting with the user, then the agent is happy. Thus, the response “Because I am chatting with you” could be generated.

Through initiative management, user modeling, non-monotonic inference, and the handling of complex questions, the agent would be more equipped to promote lexical scalar enrichment while at the same time offering a robust system of interaction to its users. In future developments, I look forward to incorporating these areas into an automated system for scalar enrichment.

References

Asher, Nicholas (1998) Questions in Dialogue. *Linguistics and Philosophy* 21: 237-309.

Berlin, Brent and Paul Kay (1969) *Basic Color Terms: Their Universality and Evolution.* Berkeley: University of California Press.

Bhat, D.N.S. (2004) *Pronouns.* Oxford: Oxford University Press.

Brennan, Susan, Marilyn Friedman and Carl Pollard (1987) A Centering Approach to Pronouns. *Proceedings of the 25th ACL.* Stanford, CA. 155-162.

Fellbaum, Christiane (ed.) (1998) *WordNet: An Electronic Lexical Database and Some of its Applications.* Cambridge: MIT press.

Green, Rebecca (2006) Vocabulary Alignment via Basic Level Categories. OCLC/ALISE research grant report published electronically by OCLC Research. Available online at `http://www.oclc.org/research/grants/reports/green/rg2005.pdf`

Grice, H. P. (1975) Logic and conversation. *Speech Acts*, eds. P. Cole and J. L. Morgan. New York: Academic Press. 41-58

Grosz, Barbara and Candace Sidner (1998) Lost Intuitions and Forgotten Intentions. *Centering Theory in Discourse*, eds. Walker, Marilyn, Avarind Joshi and Ellen Prince. Oxford, UK: Clarendon Press. 39-51.

Grosz, Barbara and Candace Sidner (1990) Plans for Discourse. *Intentions in Communication*, eds. P. R. Cohen, J. Morgan and M. E. Pollack. Cambridge, MA: MIT Press. 417-444.

Grosz, Barbara, Avarind Joshi and Scott Weinstein (1995) Centering: A Framework for the Modeling of Local Coherence in Discourse. *Computational Linguistics* 21(2): 202-225.

Grosz, Barbara, Avarind Joshi and Scott Weinstein (1986) Towards a Computational Theory of Discourse Interpretation. Unpublished manuscript.

Grosz, Barbara and Candace Sidner (1986) Attention, intention, and the structure of discourse. *Computational Linguistics* 12(3): 175-204.

Grosz, Barbara (1977) The Representation and Use of Focus in Dialogue Understanding. Stanford, CA: Stanford University Ph.D. thesis

Gundel, Jeanette (2003) Information Structure and Referential Givenness/Newness: How Much Belongs in the Grammar? http://csli-publications.stanford.edu/HPSG/4/gundel.pdf Stanford, CA: CSLI Publications

Gundel, Jeanette (1998) Centering Theory and the Givenness Hierarchy. *Centering Theory in Discourse*, eds. Walker, Marilyn, Avarind Joshi and Ellen Prince. Oxford, UK: Clarendon Press. 183-198.

Gundel, Jeanette, Nancy Hedberg and Ron Zacharski (1993) Cognitive status and the Form of Referring Expressions in Discourse. *Language* 69 (2): 274-307. Linguistics Society of America

Hahn, Udo, Katja Markert and Micheal Strube (1996) A Conceptual Reasoning Approach to Textual Ellipsis. *Proceedings of the 12th Eurpoean Conference on Artificial Intelligence (ECAI '96)*. Budapest, Hungary, 12-16 August 1996, 572-576. Chichester: John Wiley.

Harabagiu, Sanda (2006) Questions and Intentions. *Advances in Open-Domain Question Answering*, eds. Strzalkowski, T. and Sanda Harabagiu. Dordrecht: Kluwer Academic Publishers.

Harabagiu, Sanda (2001) Just-In-Time Question Answering. *Proc. of the Sixth Natural Language Processing Pacific Rim Symposium (NLPRS-2001)*. Tokyo, Japan, pp. 27-34.

Hirschberg, Julia (1985) *A Theory of Scalar Implicature*. Ph.D. thesis, University of Pennsylvania.

Hobbs, Jerry (1978) Resolving Pronoun References. *Lingua* 44: 311-338.

Horn, L. R. (1984) Toward a New Taxonomy for Pragmatic Inference: Q- and R-based implicature. *Meaning, form, and use in context*, ed. D. Schriffin. Washington, DC: Georgetown University Press.

Hudson-D'mzura, Susan and Michael Tanenhaus (1998) Assigning antecedents to ambiguous Pronouns. *Centering Theory in Discourse*, eds. Walker, Marilyn, Avarind Joshi and Ellen Prince. Oxford, UK: Clarendon Press. 71-88.

Joshi, Aravind and Scott Weinstein (1981) Control of Inference: Role of some aspects of discourse structure - centering. *Proceedings of the IJCAI*. Vancouver, CA. 385-387.

Jurafsky, Daniel and James Martin (2001) *Speech and Language Processing: an Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. New Jersey: Prentice Hall.

Kamp, Han (1981) A Theory of Truth and Semantic Representation. *Formal Methods in the Study of Language, Part I*, eds. J. Groenendijk, TH. Janssen and M. Stokhof. Amsterdam: Mathematisch Centrum, 277-322

Kamp, Hans and Uwe Reyle (1993) *From Discourse to Logic Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Dordrecht: Kluwer Academic Publishers.

Kibble, Rodger and Richard Power (2004) Optimizing Referential Coherence in Text Generation. *Computational Linguistics* 30(4): 401-416.

Lakoff, George (1987) *Women, Fire, and Dangerous Things: What categories reveal about the mind*. Chicago: University of Chicago Press.

Lappin, Shalom and Herbert Leass (1994) An Algorithm for Pronominal Anaphora Resolution. *Computational Linguistics* 20(4): 535-562.

Levinson, Stephen C. (2000). *Presumptive Meanings: The theory of generalized conversational implicature*. Cambridge: MIT Press.

Levinson, Stephen C. (1983) *Pragmatics*. Cambridge University Press. New York, NY.

Liu, Hugo (2004). *MontyLingua: An end-to-end natural language processor with common sense*. Available at: web.media.mit.edu/ hugo/montylingua.

Manning, Christopher and Hinrich Scheutze (1999) *Foundations of Statistical Natural Language Processing*. Cambridge: MIT Press.

Poesio, Massimo, Rosemary Stevenson, Barbara Di Eugenio and Janet Hitzeman (2004) Centering: A Parametric Approach. *Computational Linguistics* 30 (3): 309-359.

Prince, Ellen (1992) The ZPG Letter: Subjects, definiteness, and information-status. *Discourse Description: Diverse Linguistic Analyses of a Fund-Raising Text*, eds. W. C. Mann and S. A. Thompson. Amsterdam: John Benjamins. 295-325.

Prince, Ellen (1981) Towards a taxonomy of given-new information. *Radical Pragmatics*, ed. P. Cole. New York, NY: Academic Press, 223-255.

Priss, Uta (1998) The Formalization of WordNet by Methods of Relational Concept Analysis. *WordNet: An Electronic Lexical Database and Some of its Applications*, ed. Christiane Fellbaum. Cambridge: MIT press. 179-196.

Resnik, Philip (1998) WordNet and Class-Based Probabilities. *WordNet: An Electronic Lexical Database and Some of its Applications*, ed. Christiane Fellbaum. Cambridge: MIT press. 239-263

Rosch, Eleanor (1973) Natural categories. *Cognitive Psychology* 4: 328-50.

Russell, Stuart and Peter Norvig (2003) *Artificial Intelligence A Modern Approach.* Upper Saddle River, New Jersey: Pearson Education, second edition.

Siddharthan, Advaith (2002) An Architecture for a Text Simplification System. *Proceedings of the Language Engineering Conference 2002 (LEC 2002).* Pages 64-71.

Sidner, Candace (1983) Focusing in the comprehension of definite anaphora. *Computational Models of Discourse*, eds. M. Brady and R. C. Berwick. Cambridge, MA: MIT Press. 267-330.

Sidner, Candace (1981). Focusing on the Interpretation of Pronouns. *Computational Linguistics* 7: 217-231.

Strube, Michael and Udo Hahn (1999) Functional Centering: Grounding referential coherence in information structure. *Computational Linguistics* 25 (3): 309-344.

Strube, Michael (1998) Never Look Back: An alternative to centering. *Proceedings of the 36th Meeting of the Association for Computational Linguistics and the 17th International Conference on Computational Linguistics*, 2: 1251-1257.

Toutanova, Kristina and Christopher Manning (2000) Enriching the Knowledge Sources Used in a Maximum Entropy Part-of-Speech Tagger. *Proceedings of the Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora (EMNLP/VLC-2000)*, Hong Kong.

Toutanova, K., Klein, D., Manning, C. and Singer, Y. (2003) Feature-Rich Part-of-Speech Tagging with a Cyclic Dependency Network. *Proceedings of HLT-NAACL*, 252-259.

Tetreault, Joel (2001) A Corpus-based Evaluation of Centering and Pronoun Resolution. *Computational Linguistics* 27(4): 507-520.

Walker, Marilyn, Arvarind Joshi, Ellen Prince (eds.) (1998) *Centering Theory in Discourse*. Oxford: Clarendon Press.