Inference through Alternative-Set Semantics

GANN BIERNER, *Division of Informatics, University of Edinburgh, E-mail: gbierner@cogsci.ed.ac.uk*

BONNIE WEBBER, *Division of Informatics, University of Edinburgh, E-mail: bonnie@dai.ed.ac.uk*

Abstract

We show how alternative set semantics can be used to analyze a class of lexical items frequently found in natural language requests to search engines and databases (including *other (than)*, *such (as)*, and *besides* as well as ***better*, *faster*, and other comparatives**). We show how the analyses allow us to infer properties of entities, as in the bolded phrase above where *other* licenses the inference that *better* and *faster* are comparatives.

Keywords: alternative set semantics, inference, presupposition

1 Introduction

An *alternative set* is a set of propositions which differ with respect to how one or more arguments are filled [11, 17, 18]. For example, the alternative set $\{like(mary, jen), like(mary, bob), ...\}$, summarized as $\lambda X.like(mary, X)$, represents the entities that Mary likes. Rooth uses the concept of alternative sets to describe the semantics of the focus particles *even* and *only*. *Only* involves the restriction of an alternative set to a single element. For example, in the sentence *Mary likes only Bob*, out of the alternative set $\lambda X.like(mary, X)$, only one element, $like(mary, bob)$, is true.

Other lexical items such as *such (as)*, *other (than)*, and *besides* can also be understood in terms of alternative sets ("alt-sets"). For example, *besides*, in the question *Who does Mary like besides Bob?*, appeals to the alt-set $\lambda X.like(mary, X)$ and considers all elements except for $like(mary, bob)$.

Alt-sets are of more than theoretical interest: [16] and [23] have already shown that alt-sets form the semantic basis for contrast in intonation, and this can be exploited in speech interpretation and speech generation [2]. We believe that alternative set semantics can also benefit query handling in natural language information retrieval (NLIR) systems, such as *The Electric Monk*[1] [5]. Among queries sent to the *Monk* are the following.

(1) a. *What is the drinking age in Afghanistan?*
What is the drinking age in **other** *countries?*

b. *Where can I find web browsers for download?*
Where can I find **other** *web browsers* **than** *netscape for download?*

[1] http://www.electricmonk.com

c. *Where can I find a list of all the shoe manufacturers in the world?*
Where can I find shoes made by Buffalino, **such as** *the Bushwackers?*

For the *Monk* to answer the second query in (1a), it must identify Afghanistan as a country and exclude it from the current search. Similar reasoning in (1b) should conclude that Netscape is a web browser to be excluded from the search. Incorrectly including it can so overwhelm the results that no other answers are returned to the user. In (1c), one must conclude that Bushwackers are shoes made by Buffalino which can be non-exclusively included in the search.

Hearst [8] demonstrates cases where pattern matching can extract knowledge from these constructions. In this paper, we present an alternative, formal approach to alt-set words, showing how relevant information can be systematically inferred through limited accommodation of their presuppositions. It is unlikely that *all* such information would already be available in an NLIR system's knowledge base (KB), so it is important that it be able to infer this knowledge when given the opportunity.

Using a lexicalized grammar, we show that dividing the semantics of lexical items into assertion and presupposition[2], as in [24] and [26], reveals a simple and elegant analysis which exploits regularities of alt-set words. This analysis provides a new approach to interpreting alt-set words, including those with discourse anaphora such as (1a), which [8] fails to handle. At the same time, the analysis provides a method for inferring new information without any further machinery.

2 Background

2.1 CCG

Our semantic analysis is tied to a syntactic analysis using Combinatory Categorial Grammar (CCG) [22]. CCG is a lexicalized grammar that encodes both the syntactic and semantic properties of a word in the lexicon. For instance, a transitive verb such as *find* might have the lexical entry in (2).

(2) $find = \begin{cases} syntax: & (S\backslash NP)/NP \\ semantics: & \lambda X \lambda Y.find(Y,X) \end{cases}$

The symbol / refers to a rightward looking category and \ to a leftward looking category. (2) states that the syntactic category of *find* is a functor that requires its argument, a noun phrase, on its right. The corresponding semantic argument is simultaneously collected and bound to the outer variable X. A new functor is returned whose syntactic argument, another noun phrase, must be on its left. The corresponding semantic argument is bound to Y. The result is a sentence whose semantics is $find(Y, X)$ with X and Y bound as described above.

Categories combine using rules such as forward and backward function application:

(3) a. $X/Y : f \quad Y : a \quad \Rightarrow \quad X : (f\ a)$
b. $Y : a \quad X\backslash Y : f \quad \Rightarrow \quad X : (f\ a)$

The derivation in (4) shows how these rules combine the lexical items in *I find shoes* into single syntactic and semantic categories. Here, $>$ indicates forward application and $<$, backward application.

[2] We take the pragmatic view of presuppositions explored by [13] and [21] which, stated loosely, sees them as propositions that must be true for an utterance to make sense. For an overview of presupposition, see [1].

(4)
$$\dfrac{\dfrac{I}{NP : I} \quad \dfrac{\dfrac{find}{(S\backslash NP)/NP : \lambda X \lambda Y.find(Y,X)} \quad \dfrac{shoes}{NP : shoes}}{S\backslash NP : \lambda Y.find(Y, shoes)}>}{S : find(I, shoes)}<$$

2.2 Lexical Semantics as Assertion and Presupposition

We follow [24] and [26] in separating lexical semantics into *assertion* and *presupposition*. The assertion is computed during the derivation (Section 2.1), and presupposition is evaluated through the mechanism described in Section 5.1.

We will write lexical entries in the following form, where the semantic parameters scope both the assertion and presupposition:

$$\text{lexical item} = \begin{cases} syntax: & \text{syntactic category} \\ semantics: & \lambda ... \begin{cases} assertion: & \text{proposition} \\ presupposition: & \text{proposition set} \end{cases} \end{cases}$$

2.3 Set Semantics of Nouns

Following [12] and [14], a common noun is interpreted as the set of entities having the property of that noun. We take the property and the set as together forming an *alt-set*[3]. Neither speaker nor hearer need be aware of its complete extension; in fact, we demonstrate here how extensions can grow through presupposition accommodation.

To simplify matters, we will write alt-sets as $_{\mathsf{description}}\{...\}$, a set with a property attribute, called a *description*, whose elements, the *ground*, are entities with that property. For example, the alt-set represented by the noun *country* is $\{country(argentina), country(brazil), ...\}$. This alt-set has the description $\lambda X.country(X)$ and the ground $\{argentina, brazil, ...\}$. One agent's current semantics for *country* might be (5)[4]. The selector functions *ground* and *description* return the components of an alt-set.

While this way of writing the alt-set is less expressive because it does not allow abstraction over predicates, it is sufficient to demonstrate the claims in this paper.

(5) $country = {}_{\mathsf{country}}\{argentina, brazil, china, denmark\}$

3 Alt-Set Words

The *alt-set* words we have analyzed fall into two classes: those that assemble a set from elements and those that excise a set from a larger set. In either case, one particular set of elements is of interest, the *figure*. With assembly words, the figure is either admitted into the *ground* of the alternative set, or combined with a *complement* to form the ground. With excision, the *figure* is explicitly excluded from the *ground*. We show here that the *figure*, *ground*, and *complement* may derive from structurally-related constituents, or one or more of them may be presupposed.

For these analyses, we define the relation $alts(X, Y)$ with the following semantics, where $X, Y \subset entities$:

---

[3] Alt-sets are also formed by abstracting over propositions, discussed in Section 1.

[4] For readability, we abbreviate $\lambda X.country(X)$ to $\mathsf{country}$.

(6) $alts(X,Y) \iff \exists A \in \textit{alt-sets} \ s.t. \ (X \cup Y) \subset ground(A)$

Intuitively, this relation specifies that the two arguments can be found together in the ground of at least one alt-set in the KB. Note that the description component of the alt-set need not be known.

3.1 Besides

Besides is an excision word whose figure and ground are given structurally, as in Example (7), which was taken from the same corpus of questions as (1). We discover that the figure, *BidFind*, is a member of the ground, *auction search engines*, using the lexical information in (8)[5] in the derivation given in (10).

(7) *Are there auction search engines besides BidFind?*

$$(8)\quad besides = \begin{cases} syntax: & (NP\backslash NP)/NP \\ semantics: & \lambda F \lambda S \begin{cases} assertion: & S - F \\ presupposition: & F \subset description(S) \\ & alts(F, S-F) \end{cases} \end{cases}$$

In this analysis, the assertion specifies that the figure should be removed from the set of entities subject to clause-level predication, which may be true of the figure (9b) or may not be (9a).

(9) a. Fido is vicious and I hate him, but I like dogs besides Fido.
b. Fido may be my favorite, but I like dogs besides Fido.

In either case the speaker is committed to the fact that the figure (Fido) is a member of the set under consideration (dogs). This is expressed as the first presupposition in the analysis. (Since figures can be realized pronominally, the presupposition refers to the referent of the figure rather than its form. Section 5.1 discusses the consequences of this for implementation.)

The second presupposition states that the figure is an alternative to the other entities under consideration. Although this may seem redundant, we will see in Section 3.2 that it plays an important role.

$$(10)\quad \dfrac{\dfrac{\textit{auction search engines}}{NP : \mathsf{ASE}\{a,b,c\}} \quad \dfrac{\dfrac{\textit{besides}}{(NP\backslash NP)/NP : \lambda F \lambda S.(S-F)} \quad \dfrac{\textit{BidFind}}{NP : \{b\}}}{NP\backslash NP : \lambda S.(S - \{b\})}>}{NP : \mathsf{ASE}\{a,b,c\} - \{b\}}< \qquad \text{presupposition set: } \begin{cases} \{b\} \subset ASE \\ alts(\{b\}, \mathsf{ASE}\{a,c\}) \end{cases}$$

Here, the presuppositions of *besides* identify BidFind as an auction search engine and an alternative to the other auction search engines. If this knowledge is not already available to the system, it may be accommodated.

The semantics of the NP, *auction search engines besides BidFind*, uses the assertion of *besides* to yield $\mathsf{ASE}\{a,c\}$– i.e. the subset of auction search engines consisting of a and c. Search engines, including *The Electric Monk*, ultimately express requests as (possibly enriched) Boolean combinations of key phrases. So the interpretation

[5] In other constructions involving *besides*, the ground can derive from either the subject, object, and indirect object, requiring an analysis similar to *unlike* in Section 4.1. *Besides* also has an assembly use, as in the sentence *Besides John, Mary walks.*

of *auction search engines besides BidFind* would map to a rather coarse term like: `''auction search engine'' AND NOT BidFind`, but the `AND NOT` can make a significant difference in what is returned.

3.2 Such

Such, as in (1c), differs from *besides* in two ways. First, it is an assembly word: Bushwackers, the figure, is required to be included in the ground, the set of shoes made by Buffalino. Second, the range of constructions involving *such* and *as* suggest that it is more systematic to treat them separately. The related sentence (1c′) shows that they are separable within a sentence, and example (11) shows that *such* is a discourse anaphor, with a similar presupposition when it occurs in a later clause.

(1c′) Where can I find such shoes as the Bushwackers?

(11) Bushwackers are very comfortable.
Where can I find such shoes?

We treat *such* in (12) as an NP modifier with two presuppositions: (1) the figure, F, has the properties associated with the ground, S, and (2) the figure and members of the ground are alternatives[6]. This allows the simple syntactic/semantic derivation of *such shoes* given in (13).

$$(12)\quad such = \begin{cases} syntax: & NP/NP \\ semantics: & \lambda S \begin{cases} assertion: & S \cup F \\ presupposition: & F \subset description(S) \\ & alts(F, S \cup F) \end{cases} \end{cases}$$

$$(13)\quad \dfrac{\dfrac{such}{NP/NP : \lambda S.S \cup F} \quad \dfrac{shoes}{NP : \mathsf{shoe}\{a,b,c\}}}{NP : \mathsf{shoe}\{a,b,c\} \cup F}> \qquad \text{presupposition set: } \begin{cases} F \subset shoe \\ alts(F, \mathsf{shoe}\{a,b,c\} \cup F) \end{cases}$$

At this point, we simply know that some shoe, the figure F, is anaphorically presupposed, but in some cases we can identify it. Our analysis of NP-taking *as* is given in (14), from which we can perform the derivation in (15).

$$(14)\quad as = \begin{cases} syntax: & (NP\backslash NP)/NP \\ semantics: & \lambda X \lambda Y \begin{cases} assertion: & Y \\ presupposition: & alts(X, Y) \end{cases} \end{cases}$$

$$(15)\quad \dfrac{\dfrac{\dfrac{such}{NP/NP : \lambda S.S \cup F} \quad \dfrac{shoes}{NP : \mathsf{shoe}\{a,b,c\}}}{NP : \mathsf{shoe}\{a,b,c\} \cup F}< \quad \dfrac{\dfrac{as}{(NP\backslash NP)/NP : \lambda X \lambda Y.Y} \quad \dfrac{the\ Bushwackers}{NP : \{b\}}}{NP\backslash NP : \lambda Y.Y}>}{NP : \mathsf{shoe}\{a,b,c\} \cup F}< \qquad \text{presupposition set: } \begin{cases} a.\ F \subset shoe \\ b.\ alts(F, \mathsf{shoe}\{a,b,c\} \cup F) \\ c.\ alts(\{b\}, \mathsf{shoe}\{a,b,c\} \cup F) \end{cases}$$

The presupposition set is the union of the presuppositions of *such* and *as*, as bound during the derivation. The remaining variable, F, we can determine solely from the presupposition set of (15) using the old AI planning heuristic "use existing objects"

[6] The assertional semantics provided here is a very coarse approximation. The set returned by the assertion should actually be $S \cup F$ restricted by salient properties of F, such as *comfortable* in (11). If no properties are available, an effective default might be to return the most specific subsumer of F with respect to S.

[19] to avoid inventing new objects when others are already available. In particular, we can unify (15b) and (15c), discovering that F, the figure, is the unary set *Bushwackers*. This then instantiates (15a), yielding $\{b\} \subset shoe$. Unifying logical forms to instantiate variables in this way follows the "interpretation as abduction" paradigm [9, 10].

That *Bushwackers* (b) are a type of shoe may or may not already be in the discourse. If it is not, we can accommodate the fact. But this is a limited form of accommodation in that we assume that these presuppositions can be resolved with entities already available from the discourse or a very limited set of the common ground (Section 4.3). We do not postulate new entities or keep around partial, or underspecified, representations for later instantiation.

3.3 Other

The semantic analysis in (16) defines *other*, like *besides*, as an excision word that excludes the figure from the ground[7]. However, it differs in that the figure derives anaphorically rather than being contained structurally.

(16) $$other = \begin{cases} syntax: & NP/NP \\ semantics: & \lambda S \begin{cases} assertion: & S - F \\ presupposition: & F \subset description(S) \\ & alts(F, S - F) \end{cases} \end{cases}$$

For example, in (17), *other countries* is interpreted as the set of countries not including the figure, which must be available from elsewhere in the sentence, from the discourse, or from the common ground (Section 4.3).

(17) $$\dfrac{\dfrac{other}{NP/NP : \lambda S.(S - F)} \quad \dfrac{countries}{NP : \mathsf{country}\{a, b, c\}}}{NP : \mathsf{country}\{a, b, c\} - F}> \qquad \text{presupposition set: } \begin{cases} F \subset country \\ alts(F, \mathsf{country}\{a, b, c\} - F) \end{cases}$$

While the problem of identifying a presupposed figure is not yet solved (cf. Section 4), we can find the figure for *other* in constructions containing *than*. The word *than* has the same analysis as *as* except for features on the syntax not shown here, marking it as comparative rather than equative. Given this, we can perform the derivation in (18) for the relevant portion of (1b).

(18) $$\dfrac{\dfrac{\dfrac{other}{NP/NP : \lambda S.(S - F)} \quad \dfrac{web\ browsers}{NP : \mathsf{browser}\{e, n\}}}{NP : \mathsf{browser}\{e, n\} - F}> \quad \dfrac{\dfrac{than}{(NP\backslash NP)/NP : \lambda X \lambda Y.Y} \quad \dfrac{Netscape}{NP : \{n\}}}{NP\backslash NP : \lambda Y.Y}>}{NP : \mathsf{browser}\{e, n\} - F}< \qquad \text{presupposition set: } \begin{cases} F \subset browser \\ alts(F, \mathsf{browser}\{e, n\} - F) \\ alts(\{n\}, \mathsf{browser}\{e, n\} - F) \end{cases}$$

As with *such as* (Section 3.2), we unify the last two presuppositions and determine that the figure is Netscape. Thus our first presupposition is instantiated to indicate that Netscape is a browser, which we can accommodate if not already known. In addition, the assertional semantics results in $\{e, n\} - \{n\} = \{e\}$. In the context of an NLIR system, the resulting Boolean request would be something like: `''web browser'' AND NOT netscape`.

A remaining problem involves *other* NPs with relative clauses: the clause can convey either old material (for identifying what is to be excluded) or new material (to be

[7] *Other* can appear in other syntactic constructions which we lack space to consider here.

predicated of the new entity). In speech, intonation can disambiguate: for instance, in *penguins and other birds that CAN fly*, *bird* but not *flying* is predicated of *penguin*, while in *robins and other birds that can fly*, both are predicated of *robin*. For NLIR, however, disambiguation remains a problem.

4 Finding the Figure

Identifying a presupposed figure when interpreting words like *such* and *other* (i.e. when the figure is not given structurally) is comparable to determining the referent of a definite pronoun. However, there are several common constructions, besides those created with *as* and *than*, that make the presupposed figure easily identifiable.

4.1 Exploiting the Presence of Other Alt-Set Words

Like and *unlike* are assembly words, but unlike other alt-set words we have discussed, they do not take a ground. Rather, they take a figure and complement and presuppose that they are alternatives. For example, in *Unlike Mary, John likes spam*, we presuppose that Mary and John are alternatives. In this case, no explicit evidence is given for their belonging to an alt-set. However, one source of evidence is the appearance of *other* in one of the arguments– e.g. in the first sentence of this paragraph. The analysis of *unlike* in (19)[8], provides the partial derivation in (20).

(19) $$unlike = \begin{cases} syntax: & ((S/VP)/NP)/NP \\ semantics: & \lambda X \lambda Y \lambda P \begin{cases} assertion: & \neg(P\ X) \wedge (P\ Y) \\ presupposition: & alts(Y, X) \end{cases} \end{cases}$$

As in Sections 3.2 and 3.3, we can "use existing objects" and unify the last two presuppositions, determining that the figure is the set of *like* and *unlike* which, through the first presupposition, are inferred to belong to the set of alt-set words.

(20)
$$\dfrac{\dfrac{\dfrac{unlike}{((S/VP)/NP)/NP : \lambda X \lambda Y \lambda P.\neg(P\ X) \wedge (P\ Y)} \quad \dfrac{\dfrac{other}{NP/NP : \lambda S.(S - F)} \quad \dfrac{alt\text{-}set\ words,}{NP : \mathsf{as\_words}\{...\}}}{NP : \mathsf{as\_words}\{...\} - F}>}{(S/VP)/NP : \lambda Y.\lambda P.\neg(P(\mathsf{as\_words}\{...\} - F) \wedge (P\ Y)}> \quad \dfrac{like\ and\ unlike}{NP : \{like, unlike\}}}{S/VP : \lambda P.\neg(P(\mathsf{as\_words}\{...\} - F) \wedge (P\ \{like, unlike\})}>$$

$$\text{presupposition set: } \begin{cases} F \subset alt\text{-}set\ words \\ alts(F, \mathsf{as\_words}\{...\} - F) \\ alts(\{like, unlike\}, \mathsf{as\_words}\{...\} - F) \end{cases}$$

4.2 List Contexts

List contexts such as (21) below also provide a situation in which we can identify the figure. We include a presupposition with *and* (and list-forming commas) that states the coordinated items are alternatives – the same presupposition given for *as* and *than* (Section 3.2). Given that the current semantics for *country* is $\mathsf{country}\{a, b, c\}$, the presupposition of *and* is instantiated in (21) to $alts(\{a\}, \mathsf{country}\{a, b, c\} - F)$. The

[8] The syntactic type shown here is unusual, but it is designed to work with appositive constructions, which are not discussed in detail due to space considerations.

presuppositions of *other* are instantiated as in (17). This is now equivalent to (18) where we can identify the figure, Afghanistan, though unification.

(21) *What is the drinking age in Afghanistan and other countries?*

In the case of *...Afghanistan, other countries, and Dallas*, however, we must ensure that *Dallas* is not considered a country. Assuming that the grammar collects list items from left to right (easily implemented in CCG or through an incremental parser), *Afghanistan* and *other countries* will be combined first. Because we unify presuppositions as soon as possible (as in [9, 10]), by the time *Dallas* is combined into the list, the figure has already been resolved to *Afghanistan* in the manner described above. Therefore, *Dallas* cannot be the figure and is not identified as a country.

4.3 Intersentential Reference

Finding the figure outside of the sentence is handled as standard discourse anaphora. (Our implementation (Section 5) does this through an analysis based on c-command and salience.) A presupposition concerning the figure, then, becomes a further restriction placed on an antecedent. If no antecedent is found that is known to meet this restriction, an antecedent consistent with it is chosen and the presupposition, accommodated. Thus, in (1a), when evaluating the second sentence, we look for an antecedent that is a country. If Afghanistan is known to be a country, we choose it. Otherwise, if it is consistent with being a country, we choose it and accommodate that fact. Otherwise we try the common ground before failing.

That is, presuppositions (including presupposed figures) can also be licensed by elements of the common ground that come from the speaker's and hearer's shared physical or cultural situation. While in general, completely specifying user's and system's common ground is impossible, in the constrained domain of NLIR queries, we can do quite well. In particular, the user, the location of the user, and the user's web browser can all be considered part of the common ground. Thus, in the absence of alternative evidence, *other countries* probably excludes the user's current location, while *other browsers* almost certainly excludes the one being used.

5 Practical Implementation

We have implemented the ideas proposed in this paper in *Grok*, a modular NLP system written in the Java programming language. *Grok* provides a CKY-style chart parser and maintains a discourse model including alternative sets, a salience list based on *centering* [7], and a dynamic ISA hierarchy.[9]

5.1 Presuppositions

The paper has shown how the rules of CCG direct the evaluation of syntax and assertional semantics in a derivation, but not how presuppositions are evaluated.

We have indicated that presuppositions are stored lexically and are scoped by the same parameters as the assertional semantics. During parsing, as contiguous strings

[9] *Grok* and more information about *Grok* can be obtained at http://www.iccs.informatics.ed.ac.uk/~gbierner/grok/.

are combined via CCG rules, the corresponding presuppositions are combined as well. This is a compositional, monotonic approach that does not address the projection problem [11], which has not yet been a pressing issue for us.

Packaging presuppositions within the derivation allows *Grok* to entertain multiple interpretations simultaneously. At this time, we do not do incremental consistency checking, but rather wait until a particular parse is chosen. This maintains the monotonicity of the KB but does not exploit the disambiguating effect of presuppositions and, in turn, the potential reduction of parsing time. We are currently evaluating this trade-off. Incremental interpretation also requires incremental anaphora resolution since presupposed figures can be realized by pronouns. To this end, we plan to incorporate the incremental centering model given in [25].

5.2 Alt-Set Inclusion and Hierarchical Relations

The presuppositions in this paper refer to the inclusion of an NP in an alternative set. These sets are represented in *Grok* as nodes in a single-rooted inheritance hierarchy with two types of relations, member and subtype.

We have taken names, e.g. *Netscape* and *Bushwackers*, to denote individuals. Inclusion in an alt-set therefore corresponds to being a member of its corresponding node in the hierarchy: e.g. from *browsers other than Netscape*, *Netscape* becomes a member of *browsers*. Plural NPs, e.g. *browsers*, denote sets. Inclusion in an alt-set here corresponds to being a subtype. For example, in *applications other than browsers*, *browser* must be inserted as a subtype of the node corresponding to *application*.

As discussed in [3], there is a general ambiguity in English as to whether singular NPs (definite or indefinite) should be interpreted specifically (22) or generically (23).

(22) a. The lion frightened my sister.
b. A lion walks into a bar.

(23) a. The lion lives in Africa.
b. A lion is a powerful beast.

This same ambiguity leads to a problem for determining whether a singular NP should be represented as a member or a subtype. We have not tried to solve the problem in full generality. Rather, we use the following heuristic:

Given $X \subset A$ *s.t.* $A \in$ *alt-sets*, we always interpret A as a *kind*. For X, indefinite singulars are interpreted as kinds and definite singulars as individuals if an antecedent can be found in the discourse or common ground, and kinds otherwise. As above, if X denotes a kind, then X is made a subtype of A. If X is an individual, X becomes a member of A.

For example, in *an application other than a browser*, *a browser* is interpreted as a kind and therefore made a subtype of *application*. In *an animal other than the lion*, if a particular lion is in the discourse or common ground, that entity becomes a member of *animal*. Otherwise, *the lion* is interpreted as a kind and becomes a subtype of *animal*.

Another issue regarding the ISA hierarchy is that a complete KB is far too large to load into memory when running *Grok*. The solution is to have an independent server to supply this information on demand. When *Grok* requires information about browsers, for instance, it sends the request to the server and receives the entry for *browser* as well as all its paths to the root. This caches relevant information in the local KB but does not make it discourse relevant since the discourse model is a separate data structure. We assume these hierarchies will be highly branching but relatively shallow. This is supported by the hypernym structure of WordNet [15]– whose maximum depth is sixteen nodes.

5.3 Resolution

Given the anaphoric nature of some alt-set words, resolution is an important issue. The two important aspects of resolution are restrictions and salience.

One type of restriction is the dominance restriction. Given a logical form, we compute three things for each portion, x, of the logical form: external dominators of x, local dominators of x, and things dominated by x. Local dominators are within the same predicate-argument structure and external dominators are outside. So, the sentence *John thinks he likes other dogs* produces the logical form in (24a) and the dominance relations in (24b). The dominance relation described here is C-Command as it is described in [22, p.19]. It is basically the traditional view of c-command except that we compute on semantic, not syntactic, structures.

(24) a. $think(john, like(X, dog - F))$

b.

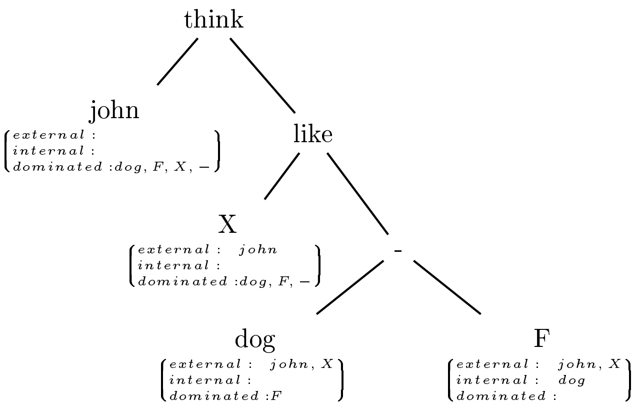

Steedman [22] summarizes how these relations are used to restrict anaphora resolution in a theory similar to that of [4]. Reflexives, for instance, must be bound to a local dominator but cannot be bound to anything it dominates, Condition A. Other anaphors are restricted by Condition B which states that they cannot bind to local dominators and Condition C, that they cannot be bound by anything they dominate. Finally, referring expressions such as definite NPs cannot resolve to any dominator or anything it dominates.

The consequence of these restrictions in the above example, is that the pronoun, *he*, cannot resolve to *dog*, *other dogs*, or F. Thus, it must resolve to *John* or something else earlier in the discourse. F, the figure of *other dogs*, cannot resolve to *dog* but can resolve to anything else, including *John* and the referent of X.

As noted in Section 4.3, figures have further restrictions placed on them. A presupposition of *John thinks other dogs like Mary* is that the figure of *other dogs* is a dog. This restriction is attached to F so that it can be used in the resolution process. Thus, when *Grok* attempts to bind F to *John*, it checks to see if John is a dog. If so, John is chosen. If not, John is rejected. If John is consistent with being a dog and there are no salient dogs available in the discourse, John is chosen and the fact that he is a dog is accommodated.

5.4 Information Retrieval

As noted earlier, this research has been motivated by the practical problem of natural language information retrieval. The main aim of applying this research to NLIR is to increase the amount of information in a query recognized as relevant and use it correctly. `"web browsers" AND netscape` is not an appropriate search request for *What are some web browsers besides Netscape.* With the approach presented in this paper a more reasonable query can be formed.

The same theory can be used to allow a search engine to inform the user of ways to improve the results of a query. For instance, if the results of the query *What are some web browsers?* are overwhelmed by pages about Netscape, the message, *Shall I search for web browsers other than Netscape?*, could be generated as text or speech. *Grok* provides this capability with a semantic head-driven generator [20].

The acquisition of knowledge in an ISA hierarchy is also pertinent to NLIR. As described in [5], *The Electric Monk* uses an ISA hierarchy to increase the scope of queries when few pertinent pages are found. By inferring knowledge from user's questions, as described in this paper, topical information is automatically made available to the *Monk* to improve the results for future queries.

6 Preliminary Evaluation

We have hand-tested the heuristics given in Section 4 for finding the figure of *other* on three corpora: a subset of the British National Corpus (BNC), a corpus of home maintenance instructions (RD) [6], and one month of queries from the *Monk*. The BNC contains a great deal of literature and literary criticism, while the RD is a more constrained "how to" text. The dialogues in the *Monk's* user interactions are the shortest and most constrained of all.

	Without Default		With Default	
	precision	recall	precision	recall
BNC	100%	10%	47%	47%
RD	100%	43.4%	57.8%	57.8%
Monk	100%	69.6%	78.3%	78.3%

TABLE 1. Accuracy of figure finding heuristics

Table 1 gives scores for precision and recall where precision is the number of *figures* correctly identified out of those attempted and recall is the number correct out of all instances of the word *other*[10]. We show two sets of scores[11]. The first uses only the techniques described in Section 4. The second includes a default heuristic which chooses the most recent sentential subject that is not the *other* phrase itself. With

[10] We excluded idioms like *on the other hand.*

[11] Note that the gold standard is not 100% because there are cases in the data sets where not enough of the discourse was available to identify the figure.

the default, precision and recall are the same because the procedure identified a figure for all instances of *other*.

These scores simply suggest that the current approach is practical since the heuristics give significantly greater accuracy with more constrained texts such as are found in the queries of NLIR systems like the *Monk*.

Comparing results with and without the default, in the latter case, 30.4% of the time, no element is chosen as the figure, causing NPs of the form "other y" to be translated to just simply `y` in the Boolean query. This will cause *false positives*, pages incorrectly containing the figure, to be included in the results of the query.

With the default, the set of instances where false positives are returned is reduced to 21.7%. In these cases, the NP "other y" is translated into `y AND NOT z` where z is an incorrectly identified figure. Here there are two possibilities: z could be of interest to the user, causing important documents to not be returned (*false negatives*). Alternatively, z could be irrelevant, in which case no false negatives occur, and the query is essentially just `y`, causing false positives. Further investigation into the frequency of these situations and the effect they have on users' ability to effectively complete their tasks will help pinpoint the appropriate tradeoffs to make in a practical system. Further evaluation will be possible when *Grok* is integrated with the *Monk*.

7 Conclusion

We have shown how a coordinated syntactic/semantic analysis of common words involving alternative sets, together with a "use existing objects" heuristic also used in abductive approaches to discourse interpretation, allows properties of entities to be automatically inferred. Although the work applies in general, it has been motivated with respect to helping to direct the search of NLIR systems like *The Electric Monk*.

We believe that significant knowledge can be accumulated and the precision of searches greatly improved in this way. Moreover, with no more machinery, we believe that analyses similar to those in this paper can be produced for a large set of commonly occurring words such as *like*, *different (than)*, *similar to*, *except (for)*, *rather than*, *apart from*, *including*, *for example*, *also*, *too*, *instead (of)*, and *another*, as well as comparatives such as *taller* and *better*. This is work we are currently carrying out.

A parallel effort is required to understand the attentional characteristics of their presuppositions, so that more accurate resolution procedures can be developed for those constructions whose arguments are not all given structurally.

8 Acknowledgments

Thanks to Mark Steedman for his encouragement and comments on earlier drafts, and to Jason Baldridge, Constance Bierner, and Julia Hockenmaier for countless useful discussions on this material, to the management of *The Electric Monk* for sharing data and information, and to the anonymous reviewers for their insightful comments.

References

[1] David Beaver. Presupposition. In Johan van Benthem and Alice ter Meulen, editors, *Handbook of Logic and Language*, pages 939–1008. North Holland, Amsterdam, 1997.

[2] Gann Bierner. TraumaTalk: Content-to-speech generation for decision support at point of care. In *Proceedings of AMIA*. Hanley and Belfus, November 1998.

[3] Gregory N. Carlson and Francis Jeffry Pelletier, editors. *The Generic Book*. The University of Chicago Press, 1995.

[4] Noam Chomsky. *Lectures on Government and Binding*. Foris, Dordrecht, 1981.

[5] Edwin Cooper. The On Point system: A natural language search engine for the world wide web. Master's thesis, University of Chicago, 1997.

[6] Reader's Digest. *Reader's Digest New Complete Do-It-Yourself Manual*. Reader's Digest Publishing Co., 1991.

[7] Barbara J. Grosz, Aravind K. Joshi, and Scott Weinstein. Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, 21(2):203–225, 1995.

[8] Marti Hearst. Automatic acquisition of hyponyms from large text corpora. In *Proceedings of the Fourteenth International Conference on Computational Linguistics*, Nantes, France, July 1992.

[9] Jerry R. Hobbs, Mark Stickel, Paul Martin, and Douglas Edwards. Interpretation as abduction. In *Proc. of the 28^{th} Annual Meeting of the Association for Computational Linguistics. Buffalo, NY, 7–10 June 1988*, pages 95–103, 1988.

[10] Jerry R. Hobbs, Mark E. Stickel, Douglas E. Appelt, and P. Martin. Interpretation as abduction. *Artificial Intelligence*, 63(1–2):69–142, October 1993.

[11] Lauri Karttunen and Stanley Peters. Conventional implicature. In Choon-Kyu Oh and David Dinneen, editors, *Syntax and Semantics 11: Presupposition*. New York, Academic Press, 1979.

[12] Fred Landman. Groups, I. *Linguistics and Philosophy*, 12:559–605, 1989.

[13] David Lewis. Scorekeeping in a language game. *Journal of Philosophical Logic*, 8:339–359, 1979.

[14] Godehard Link. The logical analysis of plurals and massterms. In R. Bäuerle, C. Schwarze, and A. von Stechow, editors, *Meaning, Use, and Interpretation in Language*, pages 302–323. de Gruyter, Berlin, 1983.

[15] George Miller. Wordnet: an on-line lexical database. *International Journal of Lexicography*, 3, 1990.

[16] Scott Prevost and Mark Steedman. Specifying intonation from context for speech synthesis. *Speech Communication*, 15:139–153, 1994.

[17] Mats Rooth. *Association with Focus*. PhD thesis, University of Massachusetts, Amherst, 1985.

[18] Mats Rooth. A theory of focus interpretation. *Natural Language Semantics*, 1:75–116, 1992.

[19] Earl Sacerdoti. *A Structure for Plans and Behavior*. Elsevier/North-Holland, Amsterdam, 1977.

[20] Stuart Shieber, Gertjan van Noord, Robert Moore, and Fernando Pereira. A semantic head-driven generation algorithm for unification-based formalisms. In *Proceedings of the 27th Conference of the Association for Computational Linguistics*, 1989.

[21] Robert Stalnaker. Pragmatic presuppositions. *Semantics and Philosophy*, pages 129–214, 1974.

[22] Mark Steedman. *Surface Structure and Interpretation*. MIT Press, Cambridge Mass., 1996. Linguistic Inquiry Monograph, 30.

[23] Mark Steedman. Information structure and the syntax-phonology interface. *Linguistic Inquiry*, 1999. To appear.

[24] Matthew Stone and Bonnie Webber. Textual economy through close coupling of syntax and semantics. In *International Workshop on Natural Language Generation*, pages 178–187, Niagara-on-the-Lake, Canada, August 1998.

[25] Michael Strube. Never look back: An alternative to centering. In *Proceedings of COLING-ACL '98, Montreal*, pages 1251–1257, 1998.

[26] Bonnie Webber, Alistair Knott, Matthew Stone, and Aravind Joshi. Discourse relations: A structural and presuppositional account using lexicalised tag. In *Proceedings of the 37th Conference of the Association for Computational Linguistics*, pages 41–48, College Park, MD, 1999.

Received January 13, 2000