

Resolving Ambiguous, Implicit and Non-Literal References by Jointly Reasoning over Linguistic and Non-Linguistic Knowledge

Nicholas L. Cassimatis

Department of Cognitive Science, Rensselaer Polytechnic Institute
Troy, NY 12180 USA
cassin@rpi.edu

Abstract

The problem of resolving ambiguous, implicit and non-literal references exemplifies many difficult issues in understanding language. We describe an approach for dealing with these by representing and jointly reasoning over linguistic and non-linguistic knowledge (including structures such as scripts and frames) within the same inference framework. This approach enables a treatment of several reference resolution phenomena that to our knowledge have not previously been the subject of a unified analysis. These results suggest that treating language understanding as an inference problem encompassing nonlinguistic knowledge can expand the ability of computational systems to use language.

1 Difficult references

Ambiguous, implicit and non-literal references embody several difficult problems in language understanding. We propose an approach for dealing with these problems that involves representing and jointly reasoning over syntactic, semantic and non-linguistic knowledge.

Formal and computational accounts of language use have difficulties with utterances whose meaning cannot be compactly and unambiguously captured as a function the meanings of the elements of an utterance. Many of these problems are evident in reference resolution and they involve the interaction of linguistic and nonlinguistic information. The following examples (adapted from (Hobbs, Stickel, Appelt, & Martin, 1990)) illustrate this:

- (1) Dave hid Paul's keys. He was drunk.
- (2) Dave hid Paul's keys. He often jokes with him.

Finding the most likely antecedent for "he" in each sentence depends in part on nonlinguistic factors such as the relationship of the people discussed, the necessity of keys in driving and the effects of drunken driving.

The following cases (copied or adapted from (McShane, in preparation)) illustrate references items not explicitly mentioned in discourse.

- (3) The couple went for a walk. He held her hand. (Referent of "he" is member of set).
- (4) The home goalie played hard but the visiting goalie played even harder. Both of them got special mention after the game. (Referent of "both" is set formed by previously mentioned objects).
- (5) The storm lasted for hours. The thunder scared my dog. (Referent for "thunder" implied by noun).
- (6) It thundered for hours. The thunder scares my dog. (Verbal antecedent).
- (7) George Bush signed the bill. (Referent of "George Bush" is in common knowledge).

Finally, the actual reference of an utterance can be entirely different from the literal reference:

- (8) The author began the book. (Pustejovsky, 1995) (The *writing* of the book was begun.)
- (9) The ham sandwich ordered some coffee. (Nunberg, 1979) (The person who is eating the ham sandwich ordered the coffee).

Several approaches have been used to resolve references computationally. One approach is first the "structured knowledge" approach. Relatively large and complex structures such as scripts (Schank & Abelson, 1977) and frames (Minsky, 1975) can be used to encode much of the knowledge needed to resolve references. For example, consider a typical frame-based account of (3). It

presumes a “couple frame” that has two slots, one for each member. One slot (M) is marked as male and the other (F) as female. During processing, when “couple” is encountered, an instance of the couple frame is instantiated. When “he” and “she” are processed, the task is to “match” them to slots in frames that have already been instantiated. In this case, “he” and “she” match the male and female member of the couple frame respectively.

Although capable of dealing with many otherwise difficult cases, the structured knowledge approach has several problems. First, structures often do not work in cases that vary slightly from those for which they were designed. Second, the matching process is not always smooth. For example, in “The couple went for a walk with their daughter. They held her hand”, “her” could match the female member of the couple and also the daughter. Matching algorithms that deal with such ambiguities are very complex and imperfect. They cannot easily incorporate “common sense” reasoning that would in this case infer that “her” refers to the daughter since otherwise the couple would be holding the hand of the female member of the couple, which is highly unusual.

The statistical, corpus driven approach to reference resolution (e.g., (Mitkov, 2000)) relies on the premise that there is enough information latent in actual instances of language use to successfully resolve inferences. By not involving complex structures or matching algorithms, they do not raise many of the difficulties of the structured knowledge approach. On some corpora, they can achieve upwards of 90% accuracy. However, many of these results rely on very specific assumptions, e.g., that the antecedent to a referent is explicitly mentioned in the text. However, as (3)-(7) illustrate, there are many cases where antecedents do not occur anywhere in the text. Current corpora cannot be used to deal with such cases because they only mark antecedents that explicitly occur. Additionally, there are many cases where even infants can find referents of novel words for which they have no statistical information. Finally, performance even in cases where corpora can be used has plateaued in the field, suggestion limits to the potential purely statistical approaches.

The “inferential” approach to reference resolution (e.g., (Hobbs et al., 1990)) eases the combination of reasoning over world and linguistic knowledge that seems to be required for reference reso-

lution. The inference approach views utterances as actions taken by people and the problem of language understanding as a kind of action understanding or abduction problem. By formulating both the linguistic and nonlinguistic constraints using the same inferential framework, the hope is that the right meaning for an utterance can be inferred using general-purpose and flexible inference engines rather than precarious structure-matching algorithms. In (1), for example, an abduction process would explain Dave’s hiding of Paul’s keys by Dave’s desire to prevent Paul from inuring himself while driving drunk. Once this has been inferred, then the subject of “was drunk” must refer to Paul and thus the coreference of “he” and “Paul” is inferred.

Although the inferential approach has achieved some success, it has suffered from the lack of powerful enough inference mechanisms has not so far yielded an analysis of non-literal uses of language.

2 An inferential approach

The work described in this paper is based on a new incarnation of the inference approach that is intended to address some of its past deficiencies and broaden the range of linguistic phenomena explainable within a single formal or computational framework. It is based on several precepts:

Action understanding. We adopt the view (Clark, 1996) that conceives of utterances as actions taken by a user and the problem of language understanding as one of finding the best explanation of these actions.

Non-linguistic constraints. We believe that non-linguistic knowledge and information is often key to inferring the meaning of an utterance. This can include knowledge about the world, people’s beliefs and desires and perceptual salience. One consequence is that nonlinguistic knowledge must be part of explaining many linguistic phenomena.

Single inferential substrate. In order to explain how linguistic and nonlinguistic knowledge constrain language understanding, we use the same “substrate” of relations and inference methods to encode and reason over this knowledge. Some constraints “span” linguistic and nonlinguistic information. For example, lexical entries often include phonological information about a word as well as what aspects of the world the word normally refers to. Thus, by combining linguistic, non-

linguistic and spanning constraints into one set of constraints, a reasoning engine that operates over this set will automatically and without any special provision use linguistic and non-linguistic information to constraint interpretation.

Structures as sets of constraints. Although we take the inferential approach, we presume that scripts, frames and other elements from structured knowledge approaches are required to explain language use. We thus encode this knowledge using the same constraint language use to encode other knowledge. Since structures tend to have exceptions (e.g., although a room script would include slots for windows, not all rooms have windows), it is important to use a framework that uses “soft” constraints that can be violated.

3 Inferential framework

Our approach relies on a language for expressing probabilistic constraints over relations among objects. Many aspects of the language has characteristics common to typical logical and probabilistic reasoning frameworks. Although work with such languages typically involves logical or probabilistic reasoning methods (such as MCMC, SAT solving or resolution), we remain agnostic in this paper as to which mechanisms are used. We do suspect however, that analogical, case-based and neural-network methods not normally associated with logical and probabilistic inference will also be required for the kind of inferences we describe here to be made in any kind of realistic scenario.

In this language, constraints are probabilistic conditionals whose antecedent and consequents are possibly negated first-order literals. Variables all start with “?”. For example, the constraint, $Wet(?x) + Iron(?x) \rightarrow .95 Rust(?x)$ states that if something is wet and iron, it has a 95% of rusting because of this. Facts can be stated as constraints with antecedents that are always true, e.g., $True() \rightarrow 1 Rises(sun)$. This can be abbreviated simple as $Rises(sun)$.

Constraints can also be followed by “posited variables” that license the positing of objects. Consider, for example:

$$Plane(?p) \wedge InRange(?p, ?r) \rightarrow .87 \\ Blip(?r), ?p$$

This constraint states that a plane in range of a radar station has an 87% chance of causing a blip. In the case where a particular radar station has a blip, one can infer the existence of a plane that caused that blip, even if the plane was not known about in advance.

Finally, we presuppose the ability to find the most likely world(s) given a set of constraints. Specifically, given a set of constraints C , there are several worlds consistent with it. A world is simply an assignment of truth values to the propositions in or licensed by C . For example, if $C = (True(Rain(today)) \rightarrow .8 \neg Rain(today))$, there are two worlds consistent with that: $w_1 = (Rain(today), true)$ and $w_2 = (Rain(today), false)$. Worlds have a probability of being actual. The probability of w_1 is .2 and of w_2 is .8.

Finally, identity is an important relation in what follows. $Same(x, y)$ states that x and y name the same object. We will presume the axioms of identity. E.g., if $P(x)$ and $Same(x, y)$, then $P(y)$.

Although apparently straightforward, inference approaches for languages with identity that license the positing of objects raise several difficult technical issues that have not begun to be dealt with until recently (e.g., (Milch et al., 2005)).

4 Fundamentals

Our overall goal is to represent and jointly reason over linguistic and nonlinguistic knowledge in order to provide a unified account of some difficult aspects of language use. This section presents some basic precepts of how to use the inferential framework described in the last section to accomplish this.

4.1 Linguistic knowledge

We will assume that the totality of a language understander’s linguistic and non-linguistic knowledge is encoded in a set of constraints, C . In what follows, we illustrate the kinds of constraints our approach uses.

Utterances can be represented using logical propositions. For example, we indicate that “Mary likes John” was uttered with the following propositions:

$$IsA(w1, Word), Phonology(w1, "mary"), \\ Occurs(w1, t1), IsA(w2, Word), \dots$$

In our analyses, we presume that the syntactic structure of utterances as given. However, work casting parsing as an inference problem (Murugesan & Cassimatis, 2006) makes us optimistic that syntactic parsing can also be dealt with as an inference problem.

The literal reference and semantic information of a word or phrase can be indicated thus:

$$LitRef(w1, litRef) \wedge Name(litRef, "Mary") \wedge IsA(litRef, Female)$$

The literal reference of a word is not always its actual reference. For example, in the case where the speaker means that Mary's dog likes John's dog, we can say: $Ref(w1, dog12)$.

How this reference is determined will be discussed in the next section.

We can represent that that the literal reference is often the actual references with:

$$LitRef(?w, ?litRef) \rightarrow p_{lit} Ref(?w, ?litRef).$$

How the actual value of p_{lit} , the probability that phrases are used literally, is arrived at is left for future research. All we assume in what follows that it is relatively close to 1.

Coreference in this framework is an identity relationship. For example, if in "John likes himself" the actual referent of John is $j-ref$ and "himself" is $h-ref$, then "John" and "himself" corefer if $Same(h-ref, j-ref)$.

Ambiguity of reference is uncertainty about identity. For example, in (10), the reference of "he" can be John or Fred.

(10) John and Fred are friends because he is rich.

This is represented thus: $Ref(w1, john)$, $Ref(w3, fred)$, $Ref(w7, h)$. "he" can refer to John: $Same(h, john)$ or to $Same(h, fred)$. If we assume the background knowledge that $\neg Same(fred, john)$, then h cannot equal both $fred$ and $john$.

We have thus far presumed several components of C, the constraints representing the listener's knowledge. To summarize, these include the set of utterances (U) heard, syntactic knowledge (SYN), semantic knowledge (SEM) (e.g, that literal refer-

ence of Mary is a female named Mary), pragmatic ($PRAG$, e.g., that the literal reference tends to be the actual referent) and non-linguistic knowledge ($WORLD$, e.g., that John is not the same person as Fred).

The debate as to whether or how to precisely distinguish between syntactic, semantic and pragmatic knowledge need not be settled to proceed with our analyses. All forms of knowledge, linguistic and non-linguistic, are treated identically within the inferential framework we are using. They are all simply constraints.

4.2 Non-linguistic knowledge

How to represent the full range of non-linguistic knowledge is of course a very broad and difficult problem. However, for our purposes, it is enough to describe how we use constraints to represent structures such as frames and scripts.

Frames and scripts can both be characterized in terms of "slot-filler" pairs together with properties of the fillers and relationships between the fillers. For example, imagine a couple frame that has two slots. The filler of one slot has the property of being a male and the filler of the other is a female.

The information in scripts and frames can be captured by constraints. For example, the following constraints represent the information encoded in the couple frame:

$$Couple(?c) \rightarrow 1$$

$PartOf(?m, ?c) \wedge PartOf(?m, ?c), ?m, ?f$
(Couple frames have two slots, each of whose filler is a part of the couple).

$$Couple(?c) \wedge PartOf(?x, ?c) \rightarrow (.5) Male(?x) \\ Couple(?c) \wedge PartOf(?m, ?c) \wedge Male(?m) \\ \wedge PartOf(?f, ?c) \rightarrow Female(?f).$$

(One slot filler of a couple is male and the other is female.)

As we have noted, one of the problems with such structures has been that they have exceptions. This can be straightforwardly dealt with by using probabilities near, but less than, 1 on the constraints characterizing a structure. This high probability biases inference according to the information in the structure while permitting exceptions.

A key aspect in using scripts and frames is the matching process. For example, in a typical ap-

proach to resolving the pronominal references in “the couple went for a walk, he held her hand”, the referent of “he” and the possessor indicated by “her” are matched to the male and female members of the couple frame. In our approach, these matches are represented by identity propositions. To say that the referent of “he” is the male of the couple is to say that they are identical, i.e., *Same(heRef, maleSlotFiller)*.

In this approach, therefore, the procedural problem of matching an object to a slot becomes a factual question about identity. As will be illustrated in the next section, this helps explain how the full range of a person’s knowledge and inference abilities can be used to find the best filler for a slot. This is a much harder task when matching is a procedural matter that is conducted by a separate algorithm or subsystem from inference about the world.

Finally, we need one more constraint to infer that matches are made at all. For example, the constraints mentioned thus far would not favor a world where a pronoun has an antecedent, e.g., where *heRef* is identical to some other object introduced into the discourse. We can favor such worlds with the following “minimal interpretation” constraint: $True() \rightarrow .51 \textit{Same}(?x, ?y)$.

All else being equal, this reduces the probability of worlds where two objects are not equal. Of course, other constraints, e.g., that “he” typically refers to a male, can override this bias.

This is, of course, a gross oversimplification. A much richer set of constraints are involved in favoring interpretations with reference, but this will be sufficient for our purposes.

4.3 Language understanding as a MAP inference problem

It is now possible to somewhat more precisely characterize the language understanding problem within the inferential approach. The listener’s knowledge is characterized by the set of constraints, $C = U \cup \textit{SYN} \cup \textit{SEM} \cup \textit{PRAG} \cup \textit{WORLD}$, i.e., the union of knowledge of the specific utterances made together with linguistic and nonlinguistic knowledge. The goal of listening is characterized as finding the most likely world given this knowledge. This world includes the identity relationships characterizing the references of phrases in U . For example, if one of the utterances encoded in U is “John likes Mary because

she is funny”, the most likely world will have the statement *Same(sheRef, mary)*.

More technically, this characterization treats language as a maximum a posteriori inference (MAP) problem. This does not however fully capture the listener’s situation. For example, the case where the most likely interpretation of a sentence has 99% probability is different from the situation where the most likely interpretation has 33% and the next most likely has 31%. In the work presented here, it will be sufficient to illustrate the benefits of the inferential approach by treating understanding as a MAP problem, although future work will need to address this issue.

5 Analyses

We now demonstrate how the substrate approach enables a unified analysis of the difficult kinds of utterances that motivated this investigation.

This is a new incarnation of the inferential approach and thus many aspects of it are oversimplified and provisional. In particular, many of the analyses below rely on simplified constraints that use probabilities that are at present guessed at. These were adequate and necessary for the goal of this work, namely to begin to develop an approach that provides a unified treatment of many difficult aspects of language. Once the outline of an approach exists, it will then become possible to more carefully elaborate aspects of the theory.

5.1 Nonlinguistic inference

We begin first by illustrating how cases where nonlinguistic inference help disambiguate a reference. Consider:

- (11) John paid Fred for the car he gave him.
- (12) John paid Fred for the car he wrecked.

In the most likely interpretation of (11), Fred (“he”) sold the car to John. In (12), John (“he”) wrecked Fred’s car and compensated him. Constraints such as the following capture the relevant knowledge:

$$\textit{Give}(?x, ?y) \rightarrow 1 \textit{Pay}(?y, ?x)$$

$$\textit{Damage}(X, Y) + \textit{Owned}(X, O) \rightarrow .8 \textit{Pay}(X, O)$$

In each sentence, there are two possible referents for “he” in the text: John (*Same(heRef,*

john) and Fred (*Same(heRef, Fred)*). In the first sentence, there is a world where Fred is the referent. In this case, the commercial transaction would explain the paying of the money. In the world John is the referent, then John giving Fred a car would not explain John paying Fred, the paying event would be unexplained and that world would have a lower probability. Thus, we infer that Fred is the referent. A similar pattern of reasoning yields the correct referent in the second sentence.

In general, each possible identity relation will imply a possible world. Since world and linguistic knowledge are represented using constraints, they both jointly determine a probability for that world. The best referent is the one that is true in the world with the highest probability.

5.2 Implicit co-referent

Although we deal with several kinds of implicit reference (e.g., references from common knowledge, members of sets and sets composed of past elements in discourse), it is possible to give them a unified treatment. In each of these cases, world knowledge licenses the inference of *implied objects* not explicitly referred to in the utterance. Then the minimal interpretation constraint favors possible identities between the explicit referent and the implied objects. Finally, world knowledge helps rank these identities according to their likelihood. The following cases illustrate this chain of inference.

Referent is member of set. Consider (3). As described in section 4, “couple” licenses the inference of two entities who are likely to be a male (*m*) and a female (*f*). “he” licenses a male (*hm*) and “she” a female (*sf*).

There are several worlds based on combinations of identity propositions. These are a few:

1. $m = fm; f = hm$. “he” refers to the female member of the couple and “she” to the male.
2. $m = hm; f = fm$. “he” and “she” refer to the male and female member of the couple respectively.
3. $m = hm; f \neq fm$. “he” refers to the male of the couple but *f* refers to someone not mentioned in the couple.

...

All worlds (e.g., world 1) where “he” and “she” refer to people of the wrong gender are given very

low probability because of the conditionals describing the semantics of the pronouns. Worlds where one of the pronouns do not refer to something explicit or implicit in the discourse (worlds 3 onward) have their probability decreased by the minimal interpretation constraint. World 2, where “he” and “she” refer to the male and female of the couple respectively is the only one that does not violate the semantic and minimal interpretation constraints and thus is the one with the highest probability.

A more complicated variation of (3) is (13):

- (13) The couple went for a walk with their daughter. They held her hand.

There are two antecedent females for “her”, the daughter and the female member of the couple, although the case where “her” refers to the “daughter” is clearly more likely. Such examples pose severe difficulties for algorithms used in structured knowledge approaches. Properly matching “her” to the female slot of the couple frame requires ruling out the match with daughter based on a chain of inference involving the fact that people tend not to hold their own hands and matching the referent of “her” to the female member implies that she is holding her own hand. Matching algorithms generally do not themselves make such inferences and are difficult to integrate with algorithms that do.

In the inference approach, one needs simply add a constraint indicating that people tend to hold other people’s hands:

$$\text{HoldHand}(?x, ?y) \rightarrow .98 \neg \text{Same}(?x, ?y).$$

Adding this constraint makes the interpretation where “her” refers to the female member of the couple less likely and thus the most likely world will have “she” refer to the daughter.

Composed referent. In (4), “both” refers to the set composed of the home and visiting goalies. The analysis here is similar to the previous case. “Both” implies the existence of a set with two objects, *o1* and *o2*. The minimal interpretation constraint favors worlds where *o1* and *o2* are identical to elements in the discourse. This leaves possible worlds with two different sets of assumptions:

$$\begin{aligned} o1 &= \text{visiting goalie}; o2 = \text{home goalie or} \\ o1 &= \text{home goalie}; o2 = \text{visiting goalie.} \end{aligned}$$

Each possibility is equally likely and thus there will be a tie for the most likely world. Ideally, there would be some way of inferring the equivalence of these worlds and thus in some sense collapsing them into one world or interpretation. However, even under the present circumstance, in each of the most likely worlds, the goalie NPs refer to members of the set “both” refers to.

Part of event. In (5), “the thunder” refers to the thunder that was a part of the storm mentioned in the first sentence. If the listener knows the following constraint, i.e., that thunder tends to be part of storms, then the finding the referent is simple.

$$\text{Storm}(?s) \rightarrow (.7)\text{Thunder}(?t) \wedge \text{PartOf}(?t, ?s), ?t$$

“The storm” licenses the inference of a storm, *s*, and the constraint licenses the possible existence of thunder, *t*, that is part of the storm. “the thunder” licenses the inference of thunder (*thunder*). The minimal interpretation constraint favors worlds where “*Same(t, thunder)*”, i.e., the interpretation that “the thunder” refers to the thunder that is part of “the storm”. As in the previous cases, the interpretation follows directly from the meaning of phrases and the minimal interpretation constraint.

Verbal antecedent. The analysis of (6) is nearly identical. The only difference is that the thunder is inferred from the mention of the thundering event.

Common Knowledge. In (7), George Bush is not introduced or implied previously in the utterance. However, for most people his existence is known and hence part of the *C* via *WORLD*. Thus, “George Bush” licenses the existence of a person whose name is “George Bush” and the minimal interpretation constraint favors worlds in which the referent of “George Bush” is identical to the George Bush of common knowledge.

5.3 Non-literal reference

Cases of non-literal reference can also be dealt with using a combination of identity matching and the minimal interpretation constraint. In this case, non-literal referents are identified on the basis of their relation to the literal referent.

To illustrate, in (9), while the ham sandwich did not order the coffee, the person who did was related to that ham sandwich (by virtue of having ordered it). This suggests that in cases where the

actual referent is not the literal referent, that there is often nevertheless a relation between them. If so, then many non-literal references can be understood by first determining the literal reference to be unlikely and then searching for an object related to it that can plausibly be the actual referent.

We represent the possibility of non-literal references being related to literal references with the “related referent constraint”:

$$\text{LitRef}(?w, ?litRef) \rightarrow p_{nonlit} \text{Ref}(?w, ?ref) \wedge ?R(?ref, ?litRef).$$

This constraint is similar in spirit to formula occurring in Pustejvsky’s (1995) treatment of coercion, though in its present manifestation, it is used to explain a wider range of other phenomena.

Treatments of coercion and other phenomena try to limit the set of relations that can be involved in these phenomena. Although we take no stance on the content of such sets, constraints can be formulated to restrict the set of relations that can relate literal and non-literal referents.

We can deal with several cases of non-literal reference. Each analysis only involves the few very general constraints about reference already discussed, the literal semantics of each utterance and some world knowledge. No special provision need be made for each phenomenon. Much variation is accounted for by non-linguistic “context”.

Coersion. In coercion, a phrase appears in a position that calls for a different type of referent than the actual referent of the phrase. For example, in (8) and (14), “begin” requires an event or action, while book is an object.

(14) The student began the book.

In (8) the action is the writing of the book while in (14) it is the reading of the book. In both cases, the literal referent is “coerced” from being of type object to type action. These also illustrate that coercion can be ambiguous and context-sensitive.

Coercion and the disambiguation of coercion are straightforwardly explained by the related referent constraint. We call the actual referent of “the book”, *bookRef*, and the literal reference, *bookLitRef*. The semantics of “began” constrains the category of book-ref to be an action (*ISA(bookRef, Action)*). Thus, worlds where the actual reference of “the book” is the literal refer-

ence (i.e., $Same(litRef, bookLitRef)$) have very low probability because the category of literal reference is an object ($isA(bookLitRef, object)$) and not an event. (This presumes background knowledge that books are not objects).

Only worlds where the actual referent is a non-literal referent (that is, related (by relation R) to the literal referent) remain. At least two actions are related to books: reading and writing.

$Author(?p) \rightarrow (.2) Write(?p, ?b) \wedge$

$Book(?b), ?b, ?w$

(Some authors write books.)

$Student(?p) \rightarrow (.9) Read(?p, ?b) \wedge Book(?b)$

(Most students read books.)

$Students(?p) \rightarrow (.99) \neg Author(?p).$

(Most students are not authors.)

Since students are more likely to read books than write them, the world where the student began reading ($Same(R, Read)$) the book is favored. Likewise, since authors write books, the world where the actual referent is a writing event ($Same(R, Write)$) is more likely in (8).

Metonymy. The account of metonymic references is almost identical. Worlds where the literal referent is the actual referent have low probability because they clash with world knowledge. Thus, a metonymic reference is more probable.

For example, with respect to (9), ham sandwiches do not order coffee. Thus, “ham sandwich” cannot refer to the ham sandwich and must instead refer to something related to it. There are many things related to the ham sandwich: e.g., the chef, the plate, the ham in it, and the waiter who served it. However, since these order coffee infrequently, if ever, worlds where they are the actual referent have low probability. Since customers do often order coffee, the world where the customer is the referent is most likely.

6 Conclusions

Language use involves difficult problems, many of which are manifest in resolving ambiguous, implied and non-literal references. These problems seem to involve the interaction of linguistic and nonlinguistic factors. Our approach attempts to deal with these problems by framing language understanding as an action understanding inference problem. This enables a unified treatment of phe-

nomena that to our knowledge have not yet been given a single explanation. It accounts for subtle variations in reference judgments based on nonlinguistic context. This work differs from past inferential approaches by extensively using identity constraints and thereby enabling an account of non-literal references, which had not been heretofore possible in inferential frameworks.

Fully realizing this approach will require a much more linguistic and non-linguistic knowledge. Acquiring it will involve learning from many instances of actual utterances. The preceding analysis provides a target for this effort and suggests significant benefits would result from it.

References

- Clark, H. (1996). *Using Language*. New York, NY: Cambridge University Press.
- Hobbs, J. R., Stickel, M. E., Appelt, D., & Martin, P. (1990). *Interpretation as Abduction* (No. 499). Menlo Park, California: AI Center, SRI International. (Document Number)
- McShane, M. (in preparation). *The Requirements of Robust Reference Resolvers*.
- Milch, B., Marthi, B., Sontag, D., Russell, S., Ong, D. L., & Kolobov, A. (2005). *BLOG: Probabilistic Models with Unknown Objects*. Paper presented at the IJCAI-05, Edinburgh, Scotland.
- Minsky, M. L. (1975). *A Framework for Representing Knowledge*. In P. H. Winston (Ed.), *The Psychology of Computer Vision*. New York, NY: McGraw-Hill.
- Mitkov, R. (2000). *A new fully automatic version of Mitkov’s knowledge-poor pronoun resolution method*. Paper presented at the CICLING-2000.
- Murugesan, A., & Cassimatis, N. L. (2006). *A Model of Syntactic Parsing Based on Domain-General Cognitive Mechanisms*. Paper presented at the 8th Annual Conference of the Cognitive Science Society, Vancouver, Canada.
- Nunberg, G. (1979). *The non-uniqueness of semantic solutions: polysemy*. *Linguistics and Philosophy*, 3(1).
- Pustejovsky, J. (1995). *The Generative Lexicon*. Cambridge, MA: MIT Press.
- Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Hillsdale, NJ: Lawrence Erlbaum.