

THE USE OF DYNAMICALLY EXTRACTED CONTEXT
FOR ANAPHORIC REFERENCE RESOLUTION

David Klappholz and Abe Lockman
Department of Electrical Engineering
and Computer Science
Columbia University
New York, New York 10027

We define anaphora more broadly than is usually done, i.e. to include not only pronoun, proadverb (e.g. here, there, then), pro-verb (e.g. do as in John goes to bed early; I do too.), definite, etc. reference, but also indefinite reference (as in 1 below), as well as reference by an entire sentence or subsentence to a previous sentence or subsentence (as in 2 below). In 1 an extremely interesting tattooed lady is a reference to the circus, with the intended relation "was performing in" or something similar; In 2 My car wouldn't start is viewed as reference to I was late to work yesterday with an intended relation which, to sweep the dirt under the rug, might be termed "cause", but which in fact is far more complicated. An anaphoric reference in this extended sense is thus an implied relation from an item (word, phrase, subsentence, sentence) in a text to one or more previously occurring items (word, phrase, subsentence, sentence, group of sentences) in that text. Resolution of anaphora involves the identification of the latter item(s) as well as of the relation intended by the writer/speaker. Thus the resolution of all instances of anaphora in a text amounts to the determination of that part of the "meaning" of the text which is more than the collection of the "meanings" of its individual sentences taken in isolation.

Since, in general, there is no test as to whether a particular item in a text is or is not an anaphoric reference (even a definite noun phrase need not be--e.g. the first president of the United States in most contexts refers to a particular person, but not necessarily to a previously occurring item in the text), an anaphora resolver must be a "guess and test" type algorithm. For each item in a new sentence, a resolver must hypothesize relations between the item and some, hopefully limited, set of possible referents, and then use directed inferencing to judge which, if any, relation, to which, if any, possible referent was intended by the writer. This directed inferencing must use not only semantics and world knowledge, but also the structure of the previous text in order to properly resolve the reference.

We are concerned with that aspect of an algorithmic theory of anaphora resolution which is dependent upon the "structure" of the text. Most previous approaches have either assumed that the structure of the text is known in advance or have ignored it entirely. In particular we reject the notion that frames/scripts somehow "solve" the problem of anaphoric reference resolution.*

Most previous approaches make no real use of information that the text itself provides concerning the flow of context. A successful algorithm

must make use of the ways in which writers choose a particular form of presentation of their ideas in order to guide the flow of contexts and allow the user to recover and use it in processing the text. These choices provide clues which must be used to dynamically extract local context at every point in the processing of a text. This local context is then a guide to reference resolution.

The most important factor in determining the local context is the way in which each sentence as a whole develops the material presented in the previous sentences of the text. Each sentence in a text is or contains either an expansion on (a further detail, adverbial, cause, effect, "how", etc.) or a temporal continuation of one or more previous sentences in the text (expansion can be either of a whole sentence or of one or more items within a sentence). In fact, the default assumption made by readers is that a sentence expands or temporally continues the immediately previous sentence of its text; semantics and world knowledge based inferencing must demonstrate that this link hypothesis is relatively weak in order to allow a sentence's interpretation as an expansion or temporal continuation of a (several) sentence(s) further back in the text.

The set of expansion/temporal continuation relations among the sentences of a text constitutes a graph which defines the context structure of the text. Our approach is to fit a new sentence into its proper places in the context graph representing the previous text, a fit that is accomplished by searching up the graph and hypothesizing and testing sentence relations (expansion/temporal continuation). Item to item reference hypotheses are entailed by the sentence relation hypotheses, and validated if and only if the relations entailing them have been validated.

Finally we have demonstrated that, for a number of syntactic constructions, the choice that a writer makes as to how to group his semantic predicates into a surface sentence influences the way in which a reader will resolve later references to items in the sentence.

An integration of the above ideas into a sketch of an algorithm for resolving reference can be found in Columbia University AI Project Technical Report 77-2-1.

1. I went to the circus yesterday. I saw an extremely interesting tattooed lady.
2. I was late to work yesterday. My car wouldn't start.

* Our ideas can, however, be taken as the rudiments of a theory of how frames/scripts can be dynamically shifted in and out of focus under control of a text's structure.