

Using Language and Context in the Analysis of Text*

Yigal Arens

Computer Science Division
Department of EECS
University of California, Berkeley
Berkeley, California 94720

Abstract

We describe a theory of natural language understanding within which we identify two separate components, a language centered one and a context centered one. The former component uses a knowledge base consisting of pairings of phrases with the concepts associated with them to determine the meaning of utterances. The latter component clarifies the meaning found by the first one and makes it more specific by attempting to reconcile it with the context of the utterance.

We have constructed a program called PHRAN (PHRasal ANalyzer) which performs the task of the language centered component.

1.0 INTRODUCTION

A recent tendency among some researchers in natural language processing has been towards a uniform, single stage process of understanding text. In such an understanding process, all the knowledge the system has is available for use at any time, and any existing routine may be run at any point. One version of this view has been called, by Schank et al (1980) and Lebowitz (1980), "the integrated parser".

At the root of this approach is the realization that previous attempts to process natural language in separate stages -- first finding a syntactic parse and then using semantic information to determine the meaning of an utterance -- were in error. Much research has been done in recent years demonstrating that semantic information is and should be used continually when processing text, and that it is wrong to assume that a separate syntactic parsing stage exists. But in arguing against the syntax-semantics distinction in processing, these natural language researchers have denied the existence of any discernible levels, or stages, in the process of understanding language and have argued that it is totally uniform.

We claim that in fact there are two distinct levels of processing in the language understanding process:

1. A language centered component, involving the application of knowledge about phrasal patterns and their meaning.
2. A context centered component, involving the application of knowledge about the context of the utterance.

These two components together perform the task of understanding the language of the utterance. After these processes have done their job the full range of knowledge we have available to us may be applied to the result. This final application of knowledge will enable us to make inferences, etc., in order to comprehend the input.

We will describe our model briefly here, and will then go over the separate components in more detail. A system called PHRAN (PHRasal ANalyzer) based on our view of the language centered component has been implemented, and will be extensively described in most of the paper. An outline of some of the ideas behind context centered processing will be given in section 5.

1.1 Brief Outline

According to our model, in understanding a text the language centered component reads the text, recognizes the words and phrasal patterns in it, and produces conceptual fragments representing the meaning of the patterns plus other contextual information it may determine from its knowledge of the language. These fragments are accumulated in a buffer which is considered the context of the utterance. The context centered component looks at the conceptual fragments as they are introduced to this buffer and uses some world knowledge to introduce to it yet other related conceptualizations. It matches the appropriate fragments to slots that were left empty by the language centered process, and it clarifies the meaning of some fragments in the manner described in section 5., thereby forming more complete conceptualizations in the buffer. These conceptual fragments are then used to represent the complete meaning of the text the system is reading. Context is also used by the language centered process when needed to resolve ambiguities.

In our model of language understanding, the language centered component has knowledge about the meaning of words, but in addition, much of the knowledge is about larger forms of utterances. This knowledge is stored in the form of pattern-concept pairs, where the pattern is a phrasal construct of varying specificity, and the concept is a notation that represents the meaning of the phrase. Together, this pair associates various forms of utterances with their meaning, and possibly with other aspects of the context in which they may appear.

The understanding process matches incoming utterances against known patterns and uses the concepts associated with the matched patterns to represent the meaning of the text.

PHRAN (PHRasal ANalyzer) is a system we have developed that models the language centered component of natural language understanding. It reads English text and produces structures that represent its meaning, so that a context centered component may act upon them. As it reads an utterance, PHRAN searches its knowledge base of pattern-concept pairs for patterns that best interpret the text. The concept portion of these pairs is then used to produce the meaning representation for the utterance.

PHRAN is able to handle phrasal language units -- which are found with great frequency in ordinary speech and common natural language text -- in the same manner that it handles the rest of the text. Since control and representation are kept separate, it is quite simple to add new information to the system. All one needs to do to extend PHRAN is add new pattern-concept pairs to the data-base. Mike Morgan, a graduate student at Berkeley, has

* Research sponsored by the Office of Naval Research under contract N00014-80-C-0732.

constructed data-bases of pattern-concept pairs that enable PHRAN to analyze Spanish and Chinese sentences, with essentially no knowledge of PHRAN's control. The knowledge base used by PHRAN is declarative, and is shareable by a system for language production. Such a system, named PHRED (Phrasal English Diction), has been constructed at Berkeley by Steve Upstill.*

2.0 PHRASAL LANGUAGE CONSTRUCTS

Many natural language processing systems assume that the meaning of utterances can be computed as a function of their constituents, with the constituents assumed to be words. Thus all the knowledge about the semantics of the language is stored at the word level (Birnbaum and Selfridge, 1979) (Riesbeck and Schank, 1975) (Wilks, 1973) (Woods, 1970). However, many natural language utterances have interpretations that cannot be found by examining their components. Idioms, canned phrases, lexical collocations, and structural formulas are instances of large classes of language utterances whose interpretation requires knowledge about the entire phrase independent of its individual words (Becker, 1975) (Mitchell, 1971).

For this reason the language centered process needs to know more about the language than the meaning of words.

By the term "phrasal language constructs" we refer to those language units of which the language user has specific knowledge. Our notion of a phrasal language construct is similar to a structural formula (Fillmore, 1979). Our constructs vary greatly in flexibility, from fixed expressions such as "a perfect stranger" to most general phrases expressing the usage of a word sense. For example, to express one usage of the verb *kick*, the phrase "<person> <kick-form> <object>" is used. (This denotes a person followed by some verb form involving *kick* (e. g., *kick*, *kicked*, "would have kicked") followed by some utterance denoting an object.)

3.0 PHRAN**

PHRAN is a system which models the processing done by the language centered component while understanding English text. PHRAN integrates both generative and non-productive language abilities. It has knowledge about individual words combined with knowledge about longer utterances of the English language.

Here are some examples of sentences for which PHRAN is able to produce conceptual fragments representing their meaning. These fragments should then be passed on to the buffer upon which a system modeling the context centered process should act to reconcile them with the context and produce more complete conceptualizations. Since no such system exists as yet, PHRAN has been extended to do some of its processing. The understanding of the following is done, therefore, entirely by PHRAN.

- * Oilmen are encouraged by the amount of natural gas discovered in the Baltimore Canyon, an undersea trough about 100 miles off the New Jersey coast. (Newsweek, Feb 1980)
- * Tenneco, one of 39 companies engaged in drilling in the area, thinks its leased tract contains a marketable supply of gas.
- * The young man has gotten into an another argument with his boss.
- * Willa will drive Bill to The Big Apple if she is given twenty five dollars.

* For more information on these projects see R. Wilensky's paper in this volume.

** For a more extensive discussion of PHRAN and phrasal patterns see Wilensky and Arens (1980).

- * Bill rewarded Mary for marrying him by making a payment on her expensive car.

(These sentences are analyzed by an uncompiled version of PHRAN on the DEC KL10 system at UC Berkeley in from 1/2 to 6 seconds of CPU time).

PHRAN is centered around a knowledge base of phrasal patterns. These range from literal strings such as "by hook or by crook" and "in a jiffy", through patterns such as "vice <title>" and "<parent> in law", to most general phrases such as "<number> <plural noun>" and "<person> <go> to <location>".

Associated with each phrasal pattern is a conceptual template. A conceptual template is a piece of meaning representation with possible references to pieces of the associated phrasal pattern. For example, the conceptual template associated with the pattern "<parent> in law" denotes the <parent> of the spouse of the person involved; and with the phrasal pattern "<person> <go> to <location>" is associated a template denoting the movement of <person> from the place s/he was to <location>.

4.0 HOW PHRAN WORKS

4.1 Overall Algorithm

PHRAN is made up of three parts - a database of pattern-concept pairs, a set of comprehension routines, and a routine which suggests appropriate pattern-concept pairs. PHRAN takes as input English text, and as it reads it from left to right, it compares the input against patterns from the database. Whenever a matching pattern is found, PHRAN interprets that part of the sentence that matched the pattern as describing the concept associated with the pattern in the pattern-concept pair, and passes the resulting conceptualizations on to the buffer where further processing will be performed by the context centered level.

4.1.1 Overview Of Processing -

When PHRAN analyzes a sentence, it reads the words one at a time, from left to right. The pattern suggesting routine determines if any new patterns should be tried, and PHRAN checks all the new patterns to see if they agree with that part of the sentence already analyzed, discarding those that don't. A word's meaning is determined simply by its matching a pattern consisting of that literal word. At this point a term is formed with the properties specified in the conceptual template associated with the word, and this term is added to the buffer PHRAN uses. PHRAN checks if the term it just added to the list completes or extends patterns that had already been partially matched by the previous terms. If a pattern is completely matched, the terms matching that pattern are removed from PHRAN's buffer and a new term, specified by the concept part of the pattern-concept pair, is formed and replaces the terms the pattern matched in its buffer. In addition, an appropriate conceptual fragment is passed on to the buffer on which the context centered process operates, for further processing.

When PHRAN finishes processing one word it reads the next, iterating the procedure just described, periodically passing conceptual fragments to the context centered level of processing.

4.1.2 Overview Of PHRAN Patterns -

A pattern-concept pair consists of a specification of the phrasal unit, an associated concept, and some additional information about how the two are related. When PHRAN instantiates a concept in its buffer, it creates an item called a term that includes the concept as well as some additional information.

A pattern is a sequence of conditions that must hold true for a sequence of terms. A pattern may specify optional terms too, the place where these may appear, and what effect (if any) their appearance will have on the properties of the term

formed if the pattern is matched. For example, consider the following informal description of one of the patterns suggested by the mention of the verb 'to take' in certain contexts.

```
{ pattern to recognize -
  <first term: represents a person>
  <second term: is an active form of TAKE>
  <third term: represents a physical object>
  <OPTIONAL part:
    <fourth term: the word FROM>
    <fifth term: represents a person>> }
term to form -
  (PTRANS (ACTOR <first term>)
    (OBJECT <third term>)
    (FROM <fifth term if present,
      otherwise consult context>)
    (TO <first term>))) }
```

Notice that the fourth and fifth terms are marked as optional. If they are not present in the text, PHRAN will not fill the FROM slot, and the context centered process will use the context to find the right slot filler.

4.1.3 Simple Example -

The following is a highly simplified example of how PHRAN processes the sentence "John dropped out of high school":

First the word "John" is read. "John" matches the pattern consisting of the literal John, and the concept associated with this pattern causes a term to be formed that represents a noun phrase and a particular male person named John. No other patterns are suggested. This term is added to *PHRAN-BUF*, the list of terms PHRAN keeps and in which it does its processing. So at this point *PHRAN-BUF* is

```
< [JOHN1 - person, NP] >
```

"Dropped" is read next. It matches the literal "dropped", and an appropriate term is formed. The pattern suggesting routine instructs PHRAN to consider the 'basic' pattern associated with the verb 'to drop', which is:

```
{ <person> <DROP> <object> } [ ... ] }
```

Its initial condition is found to be satisfied by the first term in *PHRAN-BUF* -- this fact is stored under that term so that succeeding ones will be checked to see if this partial match continues. The term that was formed after reading "dropped" is now added to the list. *PHRAN-BUF* is now

```
< [JOHN1 - person, NP] , [DROP - verb] >
```

PHRAN now checks to see if the pattern stored under the first term matches the term just added to *PHRAN-BUF* too, and indeed it does. This new fact is now stored under the last term.

Next the word "out" is read. The pattern suggestion mechanism is alerted by the occurrence of the verb 'drop' followed by the word 'out', and it instructs PHRAN to consider the pattern

```
{ <person> <DROP> "out" of <school> } [ ... ] }
```

The list in *PHRAN-BUF* is checked against this pattern to see if it matches its first two terms. It does, and this fact is stored under the second term. A term associated with 'out' is now added to *PHRAN-BUF*:

```
< [JOHN1 - person, NP] , [DROP - verb] , [OUT] >
```

The two patterns that have matched up to DROP are checked to see if the new term extends them. This is true only for the second pattern, and this fact is stored under the next term. The pattern

```
[ <person> <DROP> <object> ]
```

is discarded.

Now the word "of" is read. A term is formed and added to *PHRAN-BUF*. The pattern that matched up to OUT is extended by OF so this fact is stored under the fourth term.

The word "high" is read and a term is formed and added to *PHRAN-BUF*. Now the pattern that matched up to OF is compared against HIGH. It doesn't satisfy the next condition. PHRAN reads

"school", and the pattern suggestion routine presents PHRAN with two patterns:

1. { ["high" "school"] [representation denoting a school for 10th through 12th graders] }
2. { [<adjective> <noun>] [representation denoting noun modified by adjective] }

Both patterns are satisfied by the previous term and this fact is stored under it. The new term is added to *PHRAN-BUF*, now:

```
< [JOHN1 - person, NP] , [DROP - verb] , [OUT] ,
  [OF] , [HIGH - adj] , [SCHOOL - school, noun] >
```

The two patterns are compared against the last term, and both are matched. The last two terms are removed from *PHRAN-BUF*, and the patterns under OF are checked to determine which of the two possible meanings we have should be chosen. Patterns are suggested such that the more specific ones appear first, so that the more specific interpretation will be chosen if all patterns match equally well and if there is no possibility of further patterns being suggested that will clarify matters. That is the case here since we have reached the end of the sentence.

A term is formed and added to *PHRAN-BUF*, which now contains

```
< [JOHN1 - person, NP] , [DROP - verb] , [OUT] ,
  [OF] , [HIGH-SCHOOL1 - school, NP] >
```

The pattern under OF is checked against the last term in *PHRAN-BUF*. PHRAN finds a complete match, so all the matched terms are removed and replaced by the concept associated with this pattern.

PHRAN-BUF now contains this concept as the final result:

```
< [ ($SCHOOLING (STUDENT JOHN1)
  (SCHOOL HIGH-SCHOOL1)
  (TERMINATION PREMATURE)) ] >
```

4.2 Pattern-concept Pairs In Some More Detail

4.2.1 The Pattern -

The pattern portion of a pattern-concept pair consists of a sequence of predicates. These may take one of several forms:

1. A word; which will match only a term representing this exact word.
2. A class name (in parentheses); will match any term representing a member of this class (e.g. "(FOOD)" or "(PHYSICAL-OBJECT)").
3. A pair, the first element of which is a property name and the second is a value; will match any term having the required value of the property (e.g. "(Part-Of-Speech VERB)").

In addition, we may negate a condition or specify that a conjunction or disjunction of several must hold.

The following is one of the patterns which may be suggested by the occurrence of the verb 'give' in an utterance:

```
[ (PERSON) (ROOT GIVE) (PERSON) (PHYSOB) ]
```

4.2.1.1 Optional Parts -

To indicate the presence of optional terms, a list of pattern concept-pairs is inserted into the pattern at the appropriate place. These pairs have as their first element a sub-pattern that will match the optional terms. The second part describes how the new term to be formed if the main pattern is found should be modified to reflect the existence of

the optional sub-pattern.

The concept corresponding to the optional part of a pattern is treated in a form slightly different from the way we treat regular concept parts of pattern-concept pairs. As usual, it consists of pairs of expressions. The first of each pair will be placed as is at the end of the properties of the term to be formed, and the second will be evaluated first and then placed on that list.

For example, another pattern suggested when 'give' is seen is the following:

```
[(PERSON) (ROOT GIVE) (PHYSOB)
  ([ (TO (PERSON))
    (TO (OPT-VAL 2 CD-FORM)) ])]
```

The terms of this pattern describe a person, the verb give, and then some physical object. The last term describes the optional terms, consisting of the word to followed by a person description. Associated with this pattern is a concept part that specifies what to do with the optional part if it is there. Here it specifies that the second term in the optional pattern should fill in the TO slot in the conceptualization associated with the whole pattern.

4.2.2 The Concept -

When a pattern is matched, PHRAN removes the terms that match it from *PHRAN-BUF* and replaces them with a new term, as defined by the second part of the pattern-concept pair. For example, here is a pattern-concept pair that may be suggested when the verb 'eat' is encountered:

```
(((PERSON) (ROOT EAT) ([((FOOD))
  (FOOD (OPT-VAL 1 CD-FORM)) ]]))
[P-O-S 'SENTENCE
  CD-FORM '(INGEST (ACTOR ?ACTOR) (OBJECT ?FOOD))
  ACTOR (VALUE 1 CD-FORM)
  FOOD '(*FIND* (FOOD))]
```

The concept portion of this pair describes a term covering an entire sentence, and whose meaning is the action of INGESTING some food (Schank, 1975). The next two descriptors specify how to fill in variable parts of this action. The expression (VALUE n prop) specifies the 'prop' property of the nth term in the matched sequence of the pattern (not counting optional terms). OPT-VAL does the same thing with regards to a matched optional sub-pattern. Thus the concept description above specifies that the actor of the action is to be the term matching the first condition. The object eaten will be the term corresponding to the optional pattern if it is found, or else it will be determined by the context centered process from context.

4.3 Pattern Manipulation In More Detail

4.3.1 Reading A Word -

When a word is read PHRAN compares the patterns offered by the pattern suggesting routine with the list *PHRAN-BUF* in the manner described in the example in section 4.1.3. It discards those patterns that conflict with the information already in it. Then PHRAN tries to determine which meaning of the word to choose, using the active patterns (those that have matched up to the point where PHRAN has read). It checks if there is a particular meaning that will match the next slot in some pattern or, if no such definition exists, if there is a meaning that might be the beginning of a sequence of terms whose meaning, as determined via a pattern-concept pair, will satisfy the next slot in one of the active patterns. If this is the case, that meaning of the word is chosen. If not enough information is available to facilitate a decision, all meanings are recorded with the hope that a pattern suggested at a later point will make a choice possible then.

In principle, of course, the context centered level should be able to help in such situations, using the context.

A new term is formed and is both passed the buffer of the context centered level and recorded in *PHRAN-BUF*. If it satisfies the next condition in one of the active patterns, the appropriate pattern is moved to the pattern-list of the new term. If the next condition in the pattern indicates that the term specified is optional, then PHRAN checks for these optional terms, and if it is convinced that they are not present, it checks to see if the new term satisfies the condition following the optional ones in the pattern.

4.3.2 A Pattern Is Matched -

When a pattern has been matched completely, PHRAN continues checking all the other patterns on the pattern-list. When it has finished, PHRAN will take the longest pattern that was matched and will consider the concept of its pattern-concept pair to be the meaning of the sequence. If there are several patterns of the same length that were matched PHRAN will group all their meanings together.

New patterns are suggested and a disambiguation process follows, exactly as in the case of a new word being read, with the resulting meaning(s) being passed to the context centered process.

For example, the words "the big apple", when recognized, will have two possible meanings; one being a large fruit, the other being New York City. PHRAN will check the patterns active at that time to determine if one of these two meanings satisfies the next condition in one of the patterns, and if so, that meaning will be chosen. Otherwise, both meanings will be entertained until future information confirms one of them.

4.3.2.1 Possible Extension Of A Pattern -

It sometimes happens that we have two patterns where the smaller one ends before the longer one. For example, consider

```
[ (PERSON) (ROOT SIT) (EVENT) OUT ]
```

and

```
[ (PERSON) (ROOT SIT) ],
```

when PHRAN reads the sentence "Jane sat the game out". In this case, upon analyzing the word "sit", PHRAN realizes that the second pattern has matched completely, but it is also aware of the fact that there exists a pattern which has matched up to the same point and which also may be present in the text. Since PHRAN is intended to match the longest pattern possible, it will not treat the shorter pattern as matched yet. PHRAN will continue reading and analyzing the text until either the longer pattern is found to be present, in which case it will be used to find the meaning of the sentence, or until the longer pattern fails and PHRAN concludes that it is not present. In the latter case PHRAN will back up and consider the shorter pattern matched.

4.4 Indexing And Pattern Suggestion

Retrieving the phrasal pattern matching a particular utterance from PHRAN's knowledge base is an important problem that we have not yet solved to our complete satisfaction. We find some consolation in the fact that the problem of indexing a large data base is a necessary and familiar problem for all knowledge based systems.

In the current version we index the pattern-concept pairs of the database in a tree. As words are read, the pattern suggesting mechanism travels down this tree, choosing branches according to the meanings of the words. It suggests to PHRAN the patterns found at the nodes at which it has arrived. This list of nodes is remembered, and when the next word is read the routine continues to branch from them, in addition to starting from the root. In practice, the number of nodes in the list is small.

For example, whenever a noun-phrase is followed by an active form of some verb, the node the suggesting routine arrives at in the tree instructs PHRAN to consider the simple declarative forms of the verb. The phrasal pattern that will recognize the expression "by and large" is found at the node reached only after seeing those three words consecutively. In this manner this pattern will be suggested only when necessary.

The main problem with this scheme is that it does not lend itself well to allowing contextual cues to influence the choice of patterns PHRAN should try.

5.0 CONTEXT CENTERED PROCESSING AND RECONCILIATION

5.1 The Context, And Reconciling Conceptual Fragments With It

Consider the well known example of disambiguation:

(1) The old man's glasses were filled with sherry.

When most people read this sentence, the word "glasses" is interpreted as meaning "eye-glasses", as opposed to "drinking glasses". This decision is based on the fact that the mention of "the old man" activates some knowledge concerning older people and their normal physical appearance.

The difficulty people have with this sentence supports the claim that in understanding text people form a context, which includes the concepts appearing up to the point where they've read and related knowledge. Upon processing future utterances and producing additional conceptual fragments, the understander attempts to reconcile these fragments -- whether they describe objects or higher level concepts -- with the context. We call this process reconciliation.

The existence of a context will also help the understanding system in the processing of the following sentence:

(2) Mary cut the salami.

Here, reconciling the general notion of cutting with this particular context will enable the system to realize that, among other things, the cutting was done with a knife, it was done in a direction perpendicular to the length of the salami, or at a slight angle to it, and that the salami was probably peeled too.

On the other hand, when we hear mention of cutting in the context of a meal, the reconciliation process enables the system to realize that the instrument used was a knife, and that with the help of a fork the person doing the cutting separated bite sized pieces from the food and proceeded to ingest them.

In comparing the limited number of inferences that can be made from a generalized concept of cutting to what can be understood from the fully specified one the system arrives at after reconciling it with the context, we see the utility of this process to the natural language understander.

5.2 The Context Centered Process

In view of the previous discussion we have decided to include in our model of natural language understanding another component, the context centered component. This component attempts to reconcile conceptual fragments which are introduced to the buffer with what is already known about the context of the utterance.

No program has yet been written implementing this component of our model. However, we can at this point state in general terms what it is expected to do:

1. Add recognized concepts to the context.

2. Recall related concepts from memory and add them to context (making use of other routines).
3. Whenever adding a concept to context, check if context requires a particular interpretation of it. (Thus resolving some ambiguities that the language centered level is unable to resolve by itself.)
4. Determine references (based on the known context).

In performing its task the context centered process makes use of the context of the utterance as previously recognized. In determining references, for instance, it will not be able to resolve a reference if the thing referred to has not been previously introduced to the context.

Unlike the knowledge base used by the language centered process, much of the knowledge used at this stage is not relevant only to understanding ideas communicated to us. We apply here, among other things, the same knowledge we might use in order to understand someone's observed actions, or to plan our own. For example, the knowledge about cutting salami used to understand the sentence dealt with earlier may be helpful when we decide to serve some to friends, or when we are trying to figure out what someone is doing when we see them bending over a salami with a sharp knife in their hand. Thus, the context centered component may call on other routines to supply it with further related conceptualizations -- which will contribute to the context of the utterances.

One should keep in mind that the two processes described in this paper are constantly influencing each other. Obviously, the language centered component supplies the context centered component with fragments to work on; but there is interaction in the other direction too. As we have seen, fragments recognized by the language centered process may be used as constituents in further processing by it. The context centered level, however, may add information to the fragment found by the previous level, and in this manner modify the meaning of any fragments of which this one is a part. In addition, the information concerning the context, which is collected during context centered processing, is used, when necessary, to determine the appropriate meaning of a word or phrase the language centered component has found.

5.3 Example

In order clarify some of the ideas described in the previous section, let us simulate the processing of the following story:

It was duck hunting season.
Mary went hunting.
She aimed at a bird and pulled the trigger.
She got the bird.

When the first sentence is processed by the language centered level, an appropriate conceptual fragment is produced, namely one representing the fact that the time of year described was that during which one could hunt ducks. When the context centered process sees that duck hunting is mentioned, it consults memory about ducks and about hunting them. It adds to the context, in addition to the fragment representing the first statement, conceptualizations representing:

Prototypical knowledge about ducks, knowledge about hunting, and in particular, guns and how they are used (e. g. representations for 'aiming' and 'shooting'), and perhaps others (the precise conceptualizations introduced to context will be a function of how memory is organized; cf. MOPS in Schank, 1980). Among the conceptualizations added there will be one of the form (\$HUNT (HUNTER ?Y) (OBJECT DUCK)), with some indication that it has not yet been seen in the text. This fragment will be of use later.

When reading the second sentence, the language centered level produces a fragment representing the fact described in it; i. e. something like (\$HUNT (HUNTER MARY) (OBJECT ?X)). Then the context centered process, in attempting to reconcile the fragment with the context, identifies the hunting in the context (i. e. hunting duck, plus the additional information there) with the hunting Mary was doing. This will happen because in the context we have, as we recall, (\$HUNT (HUNTER ?Y) (OBJECT DUCK)), which the previous fragment will match.

When the word 'she' is read the person is identified as Mary, since the context contains knowledge of only one person. The concept of aiming is identified with the aiming of a gun already in the context, in a fashion similar to that described for the previous sentence. The bird is identified with a duck, since the system is told that Mary aimed at a bird, context includes the concept of aiming at a duck, and specifying the bird as a duck is both possible -- in this case because a duck is a bird -- and enhances the match between the text and the context. Since the aiming in the context is marked as part of the shooting which is part of hunting, we will get that the aiming mentioned in the text is also understood as part of the process of hunting duck. The pulling of the trigger is, again, understood in terms of using the gun and hunting.

When reading the last sentence, the language centered component comes up with several possible meanings, among them that 'Mary received the bird from an unknown party', and that 'Mary succeeded in doing what she wanted to do to the bird'. The information about hunting already in the context enables the context centered component to decide in favor of the second meaning, because we know nothing about anyone giving Mary anything, but we have represented the fact that Mary wanted to do something to the bird (i. e. kill it). This conceptualization will match part of the second possible meaning of the phrase, so the last sentence will be understood as 'Mary succeeded in killing the bird'.

References

- Becker, Joseph D. (1975). The phrasal lexicon. In Theoretical Issues in Natural Language Processing. R. Schank and B.L. Nash-Webber (eds.). Cambridge, Mass.
- Birnbaum, L., and Selfridge, M. (1979). Problems in conceptual analysis of natural language. Yale University Department of Computer Science Research Report #168.
- Fillmore, C.J. (1979). Innocence: A Second Idealization for Linguistics. In Proceedings of the Fifth Berkeley Language Symposium, Berkeley, California.
- Lebowitz, Michael (1980). Generalization and Memory in an Integrated Understanding System. Yale University Dept. of Computer Science Research Report #186.
- Mitchell, T. F. (1971). Linguistic "Goings On"; Collocations and Other Matters Arising on the Syntactic Record. Archivum Linguisticum 2 (new series) 35-69.
- Riesbeck, C. K., and Schank, R. C. (1975). Comprehension by computer: expectation-based analysis of sentences in context. Yale University Research Report #78.
- Schank, R. C. (1975). Conceptual Information Processing. American Elsevier Publishing Company, Inc., New York.
- Schank, R. C. (1980). Language and Memory. Cognitive Science, vol. 4, no. 3.
- Schank, R. C., Lebowitz, M., and Birnbaum, L. (1980). An Integrated Understander. Am. J. of Computational Linguistics vol.6, no.1.
- Wilensky, R., and Arens, Y. (1980). PHRAN - a Knowledge Based Approach to Natural Language Analysis. University of California at Berkeley, Electronic Research Laboratory Memorandum No. UCB/ERL M80/34.
- Wilks, Yorick (1973). An AI Approach to Machine Translation. In Computer Models of Thought and Language, R.C. Schank and K.M. Colby (eds.), W.H. Freeman and Co., San Francisco, 1973.
- Woods, W. A. (1970). Transition Network Grammars for Natural Language Analysis. CACM 13, 591-606.