

U N D E R S T A N D I N G  
B Y  
C O N C E P T U A L I N F E R E N C E

CHUCK RIEGER

Computer Science Department  
University of Maryland  
College Park

*Copyright 1975*

*Association for Computational Linguistics*

## ABSTRACT

Any theory of language must also be a theory of inference and memory. It does not appear to be possible to "understand" even the simplest of utterances in a contextually meaningful way in a system in which language fails to interact with a language-free memory and belief system, or in a system which lacks a spontaneous inference reflex.

People apply a tremendous amount of cognitive effort to understanding the meaning content of language in context. Most of this effort is of the form of spontaneous conceptual inferences which occur in a language-independent meaning environment. I have developed a theory of how humans process the meaning content of utterances in context. The theory is called Conceptual Memory, and has been implemented by a computer program which is designed to accept as input analyzed Conceptual Dependency (Schank et al.) meaning graphs, to generate many conceptual inferences as automatic responses, then to identify points of contact among those inferences in "inference space". Points of contact establish new pathways through existing memory structures, and hence "knit" each utterance in with its surrounding context.

Sixteen classes of conceptual inference have been identified and implemented, at least at the prototype level. These classes appear to be essential to all higher-level language comprehension processes. Among them are causative/resultative (those which predict cause and effect relations), motivational (those which predict and describe actors' intentions), enablement (those which predict the surrounding context of actions), state-duration (those which predict the fuzzy duration of various states in the world) normative (those which assess the "normality" of a piece of information - how unusual it is), and specification (those which predict and fill in missing conceptual information in a language-communicated meaning graph).

Interactions of conceptual inference with the language processes of (1) word sense promotion in context, and (2) ident-

ification of referents to memory tokens are discussed. A theoretically important inference-reference "relaxation cycle" is identified, and its solution discussed.

The theory provides the basis of a computationally effective model of language comprehension at a deep conceptual level, and should therefore be of interest to computational linguists, psychologists and computer scientists alike.

## TABLE OF CONTENTS

1. The Need for a Theory of Conceptual Memory and Inference .....	5
2. A Simple Example .....	7
3. Background .....	10
4. A Brief Overview of the Conceptual Memory's Inference Control Structure .....	15
5. The Sixteen Theoretical Classes of Conceptual Inference .....	20
5.1. CLASS 1: Specification Inferences .....	21
5.2. CLASSES 2 and 3: Resultative and Causative Inferences .....	25
5.3. CLASS 4: Motivational Inferences .....	26
5.4. CLASS 5: Enabling Inferences .....	27
5.5. CLASS 6: Action Prediction Inferences .	28
5.6. CLASS 7: Enablement Prediction Inferences .....	30
5.7. CLASS 8: Function Inferences .....	31
5.8. CLASSES 9 and 10: Missing Enablement and Intervention Inferences .....	33
5.9. CLASS 11: Knowledge Propagation Inferences .....	34
5.10. CLASS 12: Normative Inferences .....	35
5.11. CLASS 13: State Duration Inferences ..	38
5.12. CLASSES 14 and 15: Feature and Situation Inferences .....	41
5.13. CLASS 16: Utterance Intent Inferences .....	42
6. Summary of the Inference Component .....	43
7. The Inference-Reference Relaxation Cycle in Conceptual Memory .....	44
8. Word Sense Promotion, and Implicit Concept Activation in the Conceptual Memory .....	48
9. Conclusion .....	51
APPENDIX A. Causal Chain Expansion Computer Example .....	52
APPENDIX B. Inference-Reference Relaxation Cycle Computer Example .....	57
REFERENCES .....	62

## 1. The Need for a Theory of Conceptual Memory and Inference

Research in natural language over the past twenty years has been focussed primarily on processes relating to the analysis of individual sentences (parsing). Most of the early work was devoted to syntax. Recently, however, there has been a considerable thrust in the areas of semantic, and importantly, conceptual analysis (see (R2), (M1), (S1) and (C1) for example). Whereas a syntactic analysis elucidates a sentence's surface syntactic structure, typically by producing some type of phrase-structure parse tree, conceptual analysis elucidates a sentence's meaning (the "picture" it produces), typically via production of an interconnected network of concepts which specifies the interrelationships among the concepts referenced by the words of the sentence. On the one hand, syntactic sentence analysis can more often than not be performed "locally" that is, on single sentences, disregarding any sort of global context; and it is reasonably clear that syntax has generally very little to do with the meaning of the thoughts it expresses. Hence, although syntax is an important link in the understanding chain, it is little more than an abstract system of encoding which does not for the most part relate in any meaningful way to the information it encodes. On the other hand, conceptual sentence analysis, by its very definition, is forced into the realm of general world knowledge; a conceptual analyzer's "syntax" is the set of rules which can produce the range of all "reasonable" events that might occur in the real world. Hence, in order to parse conceptually, the conceptual analyzer must interact with a repository of world knowledge and world knowledge handlers (inferential processes). This need for such an analyzer-accessible world knowledge repository has provided part of the motivation for the development of the following theory of conceptual inference and memory

However, the production of a conceptual network from an isolated sentence is only the first step in the understanding

process. After this first step, the real question is: what happens to this conceptual network after it has been produced by the analyzer? That is, if we regard the conceptual analyzer as a specialized component of a larger memory, then the allocation of memory resources in reaction to each sentence follows the pattern: (phase 1) get the sentence into a form which is understandable, then (phase 2) understand it! It is a desire to characterize phase 2 which has served as the primary motivation for developing this theory of memory and inference. In this sense, the theory is intended to be a charting-out of the kinds of processes which must surely occur each time a sentence's conceptual network enters the system. Although it is not intended to be an adequate or verifiable model of how these processes might actually occur in humans, the theory described in this paper has nevertheless been implemented as a computer model under PDP-10 Stanford 1.6 LISP. While the implementation follows as best it can an intuitively correct approach to the various processes described, the main intent of the underlying theory is to propose a set of memory processes which, taken together, could behave in a manner similar to the way a human behaves when he "understands language".

## 2. A Simple Example

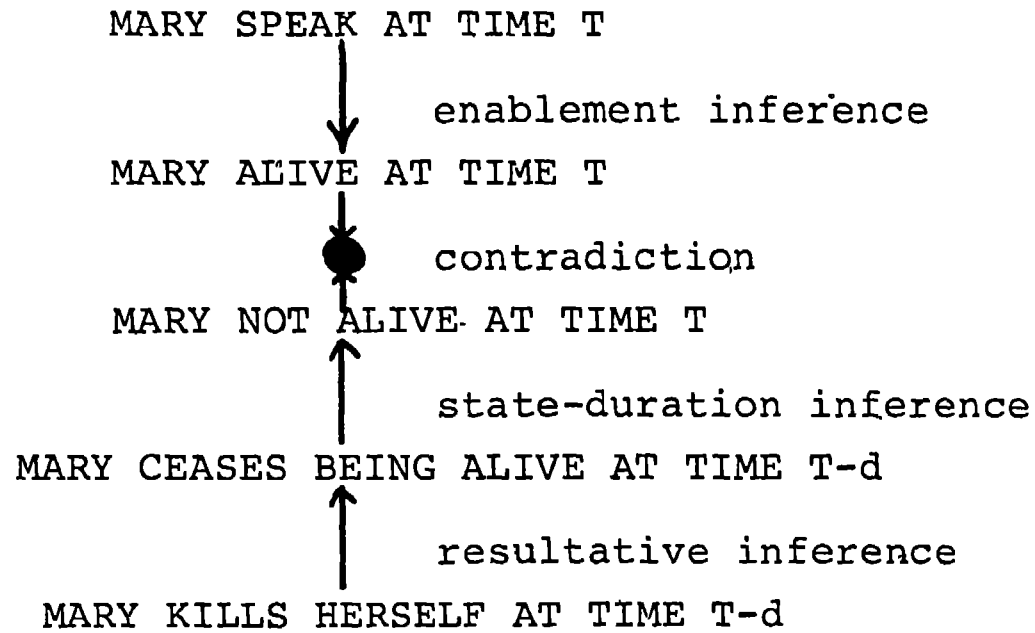
The attentive human mind is a volatile processor. My conjecture is that information simply cannot be put into it in a passive way; there are very primitive inference reflexes in its logical architecture which each input meaning stimulus triggers. I will call these primitive inference reflexes "conceptual inferences", and regard them as one class of subconscious memory process. I say "subconscious" because the concern is with a relatively low-level stratum of "higher-level cognition", particularly insofar as a human applies it to the understanding of language-communicated information. The eventual goal is to synthesize in an artificial system the rough flow of information which occurs in any normal adult response to a meaningfully-connected sequence of natural language utterances. This of course is a rather ambitious project. In this paper I will discuss some important classes of conceptual inference and their relation to a specific formalism I have developed (R1).

Let me first attempt, by a fairly ludicrous example, to convince you (1) that your mind is more than a simple receptacle for data, and (2) that you often have little control over the thoughts that pop up in response to something you perceive. Read the following sentence, pretending you were in the midst of an absorbing novel:

EARLIER THAT EVENING, MARY SAID SHE HAD KILLED HERSELF.

One of two things probably occurred: either you chose as referent of "herself" some person other than Mary (in which case everything works out fine), or (as many people seem to do) you first identified "herself" as a reference to Mary. In this case, something undoubtedly seemed awry: you realized either that your choice of referent was erroneous, that the sentence was part of some unspecified "weird" context, or that there was simply an out-and-out contradiction. Of course, all three interpretations are unusual in some sense because of a "patently obvious"

contradiction in the picture this utterance elicits. The sentence is syntactically and semantically impeccable; only when we "think about it" does the big fog horn upstairs alert us to the implicit contradiction:



Here is the argument: before reading the sentence, you probably had no suspicion that what you were about to read contained an implicit contradiction. Yet you probably discovered that contradiction effortlessly! Could there have been any a priori "goal direction" to the three simple inferences above? My conclusion is that there could not have been. If we view the mind as a multi-dimensional "inference space", then each incoming thought produces a spherical burst of activity about the point where it lands in this space (the place where the conceptual network representing it is stored). The horizon of this sphere consists of an advancing wavefront of inferences - spontaneous probes which are sent out from the point. Most will lose momentum and eventually atrophy; but a few will conjoin with inferences on the horizons of other points' spheres. The sum of these "points of contact" represents the integration of the thought into the existing fabric of the memory in that each point of contact establishes a new pathway between the new thought and existing knowledge (or perhaps among several existing pieces of knowledge). This to me is a pleasing memory paradigm, and there is a tempting analogy to be drawn with neurons and actual physical



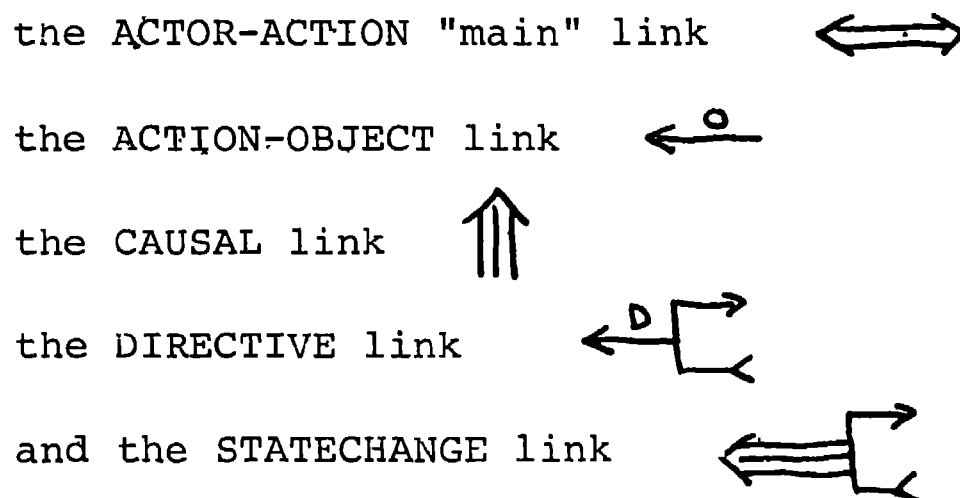
wavefronts as proposed years ago by researchers such as John Eccles (E1). The drawing of this analogy is, however, left for the pleasure of you, the reader.

This killing example was of course more pedagogical than serious, since it is a loaded utterance involving rather black and white, almost trivial inferences. But it suggests a powerful low-level mechanics for general language comprehension. Later, I will refer you to an example which shows how the implemented model, called MEMORY and described in (R1), reacts to the more interesting example MARY KISSED JOHN BECAUSE HE HIT BILL, which is perceived in a particular context. It does so in a way that integrates the thought into the framework of that context and which results in a "causal chain expansion" involving six probabilistic inferences.

### 3. Background

Central to this theory are sixteen classes of spontaneous conceptual inferences. These classes are abstract enough to be divorced from any particular meaning representation formalism. However, since they were developed concurrently with a larger model of conceptual memory (R1) which is functionally a part of a language comprehension system involving a conceptual analyzer and generator (MARGIE (S3)), it will help make the following presentation more concrete if we first have a brief look at the operation and goals of the conceptual memory in the context of the complete language comprehension system.

The memory adopts Schank et al.'s theory (S1,S2) of Conceptual Dependency (CD) as its basis for representation. CD is a theory of meaning representation which posits the existence of a small number of primitive actions (eleven are used by the conceptual memory), a number of primitive states, and a small set of connectives (links) which can join the actions and states together into conceptual graphs (networks). Typical of the links are:



Each primitive action has a case framework which defines conceptual slots which must be filled whenever the act appears in a conceptual graph. There are in addition TIME, LOCATION and INSTRUMENTAL links, and these, as are all conceptual cases, are obligatory, even if they must be inferentially filled in by the conceptual memory (CM). Figure 1 illustrates the CD representation of the

sentence MARY KISSED JOHN BECAUSE HE (JOHN) HIT BILL. That conceptual graph is read as follows: John propelled some unspecified object X from himself toward Bill, causing X to come into physical contact with Bill, and this entire event cause Mary to do something which resulted in her lips being in physical contact with John! Furthermore, the entire event occurred sometime in the past. Chapter 2 of (R1) contains a fairly complete overview of the CD representation.

Assuming the conceptual analyzer (see (R2)) has constructed, in consultation with the CM, a conceptual graph of the sort typified by Figure 1, the first step for the CM is to begin "integrating" it into some internal memory structure which is more amenable to the kinds of active inference manipulations the CM wants to perform. This initial integration occurs in three stages. First is an initial attempt to replace the symbols (JOHN, MARY, BILL, X, etc.) by pointers to actual memory concepts and tokens of concepts. Each concept and token in the CM is represented by a unique LISP atom (such as C0347) which itself bears no intrinsic meaning. Instead, the essence of the concept or token is captured in a set of features associated with the symbol. Thus, for instance, an internal memory token with no features is simply "something" if it must be expressed by language, whereas the token illustrated in Figure 2 would represent part of our knowledge about Bill's friend Mary Smith, a female human who owns a red Edsel, lives at 222 Avenue St., is 26 years old, and so forth. This set of features is called C0948's occurrence set, and is in the implementation merely a set of pointers to all other memory structures in which C0948 occurs. The process of referent identification will attempt to isolate one token, or second best, a set of candidate tokens for each concept symbol in the incoming graph by means of a feature-intersecting algorithm described in (R1).

Reference identification is the first stage of the initial integration of the graph into internal memory structures. The




FIGURE 1

Conceptual dependency representation of the sentence "Mary kissed John because he hit Bill."

C0948: (ISA # #PERSON)  
 (SEX # #FEMALE)  
 (NAME # MARY)  
 (SURNAME # SMITH)  
 (OWNS # C0713)  
 (RESIDENCE # C0846)  
 (TSTART # C0654)

...

(C0718 is the token representing the Edsel which Mary owns, C0846 is the token for Mary's place of residence, C0654 is a time token with a numeric value on the CM's time scale representing Mary's time of birth in 1948)

FIGURE 2

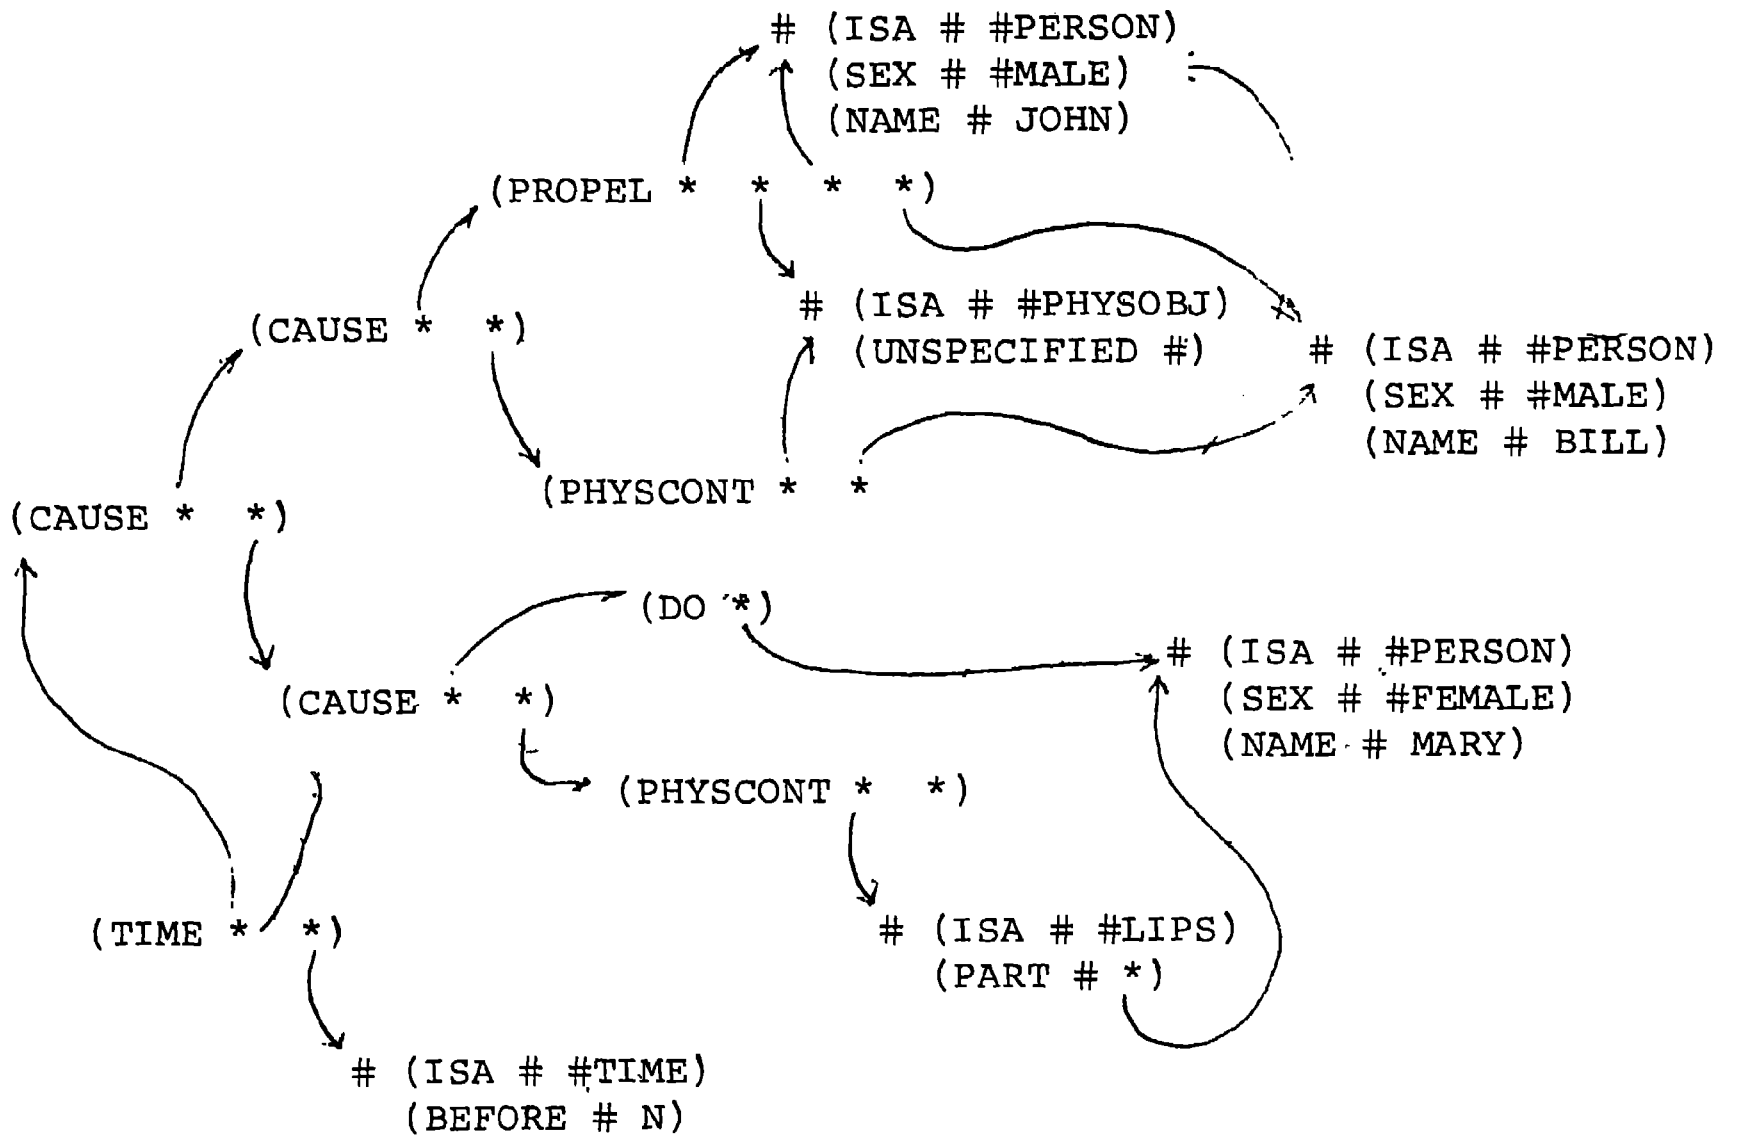
The memory token and its occurrence set which represent Mary Smith.

second and third stages are (2) the isolation of subgraphs which will form the beginning inference queue (input to the spontaneous inference component), and (3) the storage of the graph dependency links themselves as pointers in the memory. Just as for simple concepts and tokens, composite structures (actions and states) are stored under a unique internal symbol, and this symbol may also have an occurrence set. In addition, there are several other properties associated with each composite structure S: the recency of S's activation by explicit reference (RECENCY), the recency of S's activation by implicit (inferential) reference (TOUCHED), the degree to which S is held to be true (STRENGTH), a list of other composite structures from which S arose in the memory - its inferential antecedents-(REASONS), and a list of other composite structures in whose generation S played a role as antecedent (OFFSPRING). RECENCY and TOUCHED are also properties of concepts and tokens, and are used in the referent identification process.

Figure 3 shows the memory structures which result from the conceptual graph of Figure 1 after the initial integration. The net result of the initial integration is a set of starting memory structures, (actually, a list of pointers to their symbols, such as (C2496 C2301 C2207)). Each of these structures references memory concepts tokens and other composite structures.

Regarding the referent identification process, for those concepts and tokens which could not be uniquely identified, new temporary tokens will have been created, each having as its initial occurrence set a list of what is known about the entity so far.

After the initial integration, the inference component is applied simultaneously to each memory structure ("point in inference space") on the starting inference queue.



(N is the numeric "now"  
on the CM's time scale)

The internal memory structures resulting from the  
CD graph of Fig. 1.

Figure 3

#### 4. A Brief Overview of the Conceptual Memory's Inference Control Structure.

The control structure which implements the CM inference reflex is a breadth-first monitor whose queue at any moment is a list of pointers to dependency structures which have arisen by inference from the beginning structures isolated during the initial integration. It is the inference monitor's task to examine each dependency structure on the queue in turn, isolate its predicate, prepare its arguments in a standard format, collect several time aspects from the structure's occurrence set, then call the inference molecule associated with the predicate, passing along the arguments and time information.

All inferential knowledge in the CM is contained in inference molecules, which lie in one-one correspondence with conceptual predicates. An inference molecule is a structured LISP program which can perform arbitrary discrimination tests on a relevant dependency structure's features and features of all involved concepts and tokens, and which can call on specialist programs to carry out standard test information retrieval functions. ("CAUSER" is an example of such a specialist. It will scan back causal sequences from structure S until it locates a volitional action, then it returns the actor of that action as the primary causing agent of S.) Inference molecules are hence multiple-response discrimination networks whose responses are conceptual inferences (of the various theoretical types to be described) which can be made from the dependency structure. Each potential inference within the inference molecule is called an inference atom.

The contribution of an inference atom which has been found applicable to the dependency structure reports eight pieces of information to a component of the monitor called the structure generator, whose job it is to embody each new inference in a memory structure. These eight pieces of information are the following:

1. a unique mnemonic which indicates to which of the 16 theoretical classes the new inference belongs (this mnemonic is associated with the new structure only temporarily on the inference queue for subsequent control purposes)
2. the "reference name" of the generating inference atom

- (each atom has a unique name which is associated with the new memory structure for control purposes)
3. the dependency structure (a predicate which binds together several pointers to concepts, tokens and other structures), which is the substance of the new inference
  4. a default "significance factor" which is a rough, ad hoc measure of the inference's probable relative significance (this is used only if a more sophisticated process, to be described, fails)
  5. a REASONS list, which is a list of all other structures in the CM which were tested by the discrimination net leading up to this inference atom. Every dependency structure has a REASONS list recording how the structure arose, and the REASONS list plays a vital role in the generation of certain types of inference
  6. a "propagation strength factor" which, when multiplied by the STRENGTHs (degree of belief) of all structures on the REASONS list, produces the STRENGTH of the new inference. (There is a need for better heuristics here incidentally -- see (Z1) for instance.)
  7. a list of modifying structures (typically time aspects) which become the new inferred structure's initial occurrence set
  8. propagation and strength factors for each modifying structure Figure 4 illustrates the small implemented NEGCHANGE (something undergoes a negative change on some scale) inference molecule. It is included to communicate the gestalt rather than correct specifics at this early stage of development.

The two other main components of the inference monitor are the evaluator and the structure merger. It is the function of the evaluator to detect exact and fuzzy contradictions and confirmations (points of contact) between each new inference as it arises and existing memory dependency structures. Because "fuzziness" in the matching process implies access to a vast number of heuristics (to illustrate: would it be more like our friend the



---

```

(I PROG NEGCHANGE (UN PE SC) (X1 X2) (
(COND ( (EVENT UN)
      (COND ( (F1 (@ISA PE @#PERSON))
            (IR @NEGCHANGE1
              (@WANT PE (GU (@POSCHANGE PE SC))) ~PEOPLE OFTEN WANT TO BETTER
              (0.95 1.0 (CAR UN)) ~THEMSELVES AFTER SOME NEGCHANGE
              (@TS @* (TI UN))
              (1.0 (CAR UN)))
            (COND ( (AND (SETQ X1 (F1 (@*MFEEL* @_ @#NEGEMOTION PE)))
                      (SETQ X2 !GLOBALFIND))
                  (IR @NEGCHANGE2
                    (@POSCHANGE X1 @#JOY)
                    (0.9 1.0 (CAR UN) X2) ~PERSON GETS HAPPY WHEN ENEMY
                    (@TIME @* (TI UN)) ~SUFFERS NEGCHANGE
                    (1.0 (CAR UN)))
                )
            (COND ( (AND (SETQ X1 (CAUSER (CAR UN)))
                        (NOT (EQ (CAR X1) (C2 (CDR X1)))))
                  (IR @NEGCHANGE3
                    (@*MFEEL* PE @#NEGEMOTION (CAR X1)) ~PEOPLE DON'T LIKE
                    (0.95 1.0 (CAR UN) (CDR X1)) ~OTHERS WHO HURT THEM
                    (@TS @* (TI UN))
                    (1.0 (CAR UN)))
                )
            )
        )
      ( (HASPRP PE (@ISA PE @#PHYSOBJ))
        (COND ( (AND (SETQ X1 (F1 (@*OWN* PE _)))
                    (SETQ X2 (CAUSER (CAR UN)))
                    (NOT (EQ X1 (CAR C2))))
              (IR @NEGCHANGE4
                (@*MFEEL* X1 @#NEGEMOTION (CAR X2)) ~IF X DAMAGES Y'S PROPERT
                (0.85 1.0 (CAR UN) X1 (CDR X2)) ~THEN X MIGHT FEEL ANGER
                (TS @* (TI UN)) ~TOWARD Y
                (1.0 (CAR UN)))
              )
            )
          )
        )
      )
    )
  )
)

```

An inference molecule used by the current program.

---

FIGURE 4

The NEGCHANGE inference molecule.

lawyer or our friend the carpenter to own a radial arm saw?), the evaluator delegates most of the matching responsibility to programs - again organized by conceptual predicates - called normality molecules ("N-molecules"). N-molecules, which will be discussed more later, can apply detailed heuristics to ferret out fuzzy confirmations and contradictions. As I will describe, N-molecules also implement one class of conceptual inference

Confirmations and contradictions discovered by the evaluator are noted on special lists which serve as sources for possible subsequent responses by the CM. In addition, confirmations lead to invocation of the structure merger, which physically replaces the two matching structures by one new aggregate structure, and thereby knits together two lines of inference. As events go, this is one of the most exciting in the CM.

Inference cutoff occurs when the product of an inference's STRENGTH (likelihood) and its significance factor falls below a threshold (0.25). This ultimately restricts the radius of each sphere in the inference space, and in the current model, the threshold is set low to allow considerable expansion.

Figure 5 depicts the overall strategy of the inference monitor. (R1) contains a fuller account of the inference control structure, whose description will be terminated at this point.

Enter... sixteen theoretical classes of conceptual inference which fuel this inference reflex.

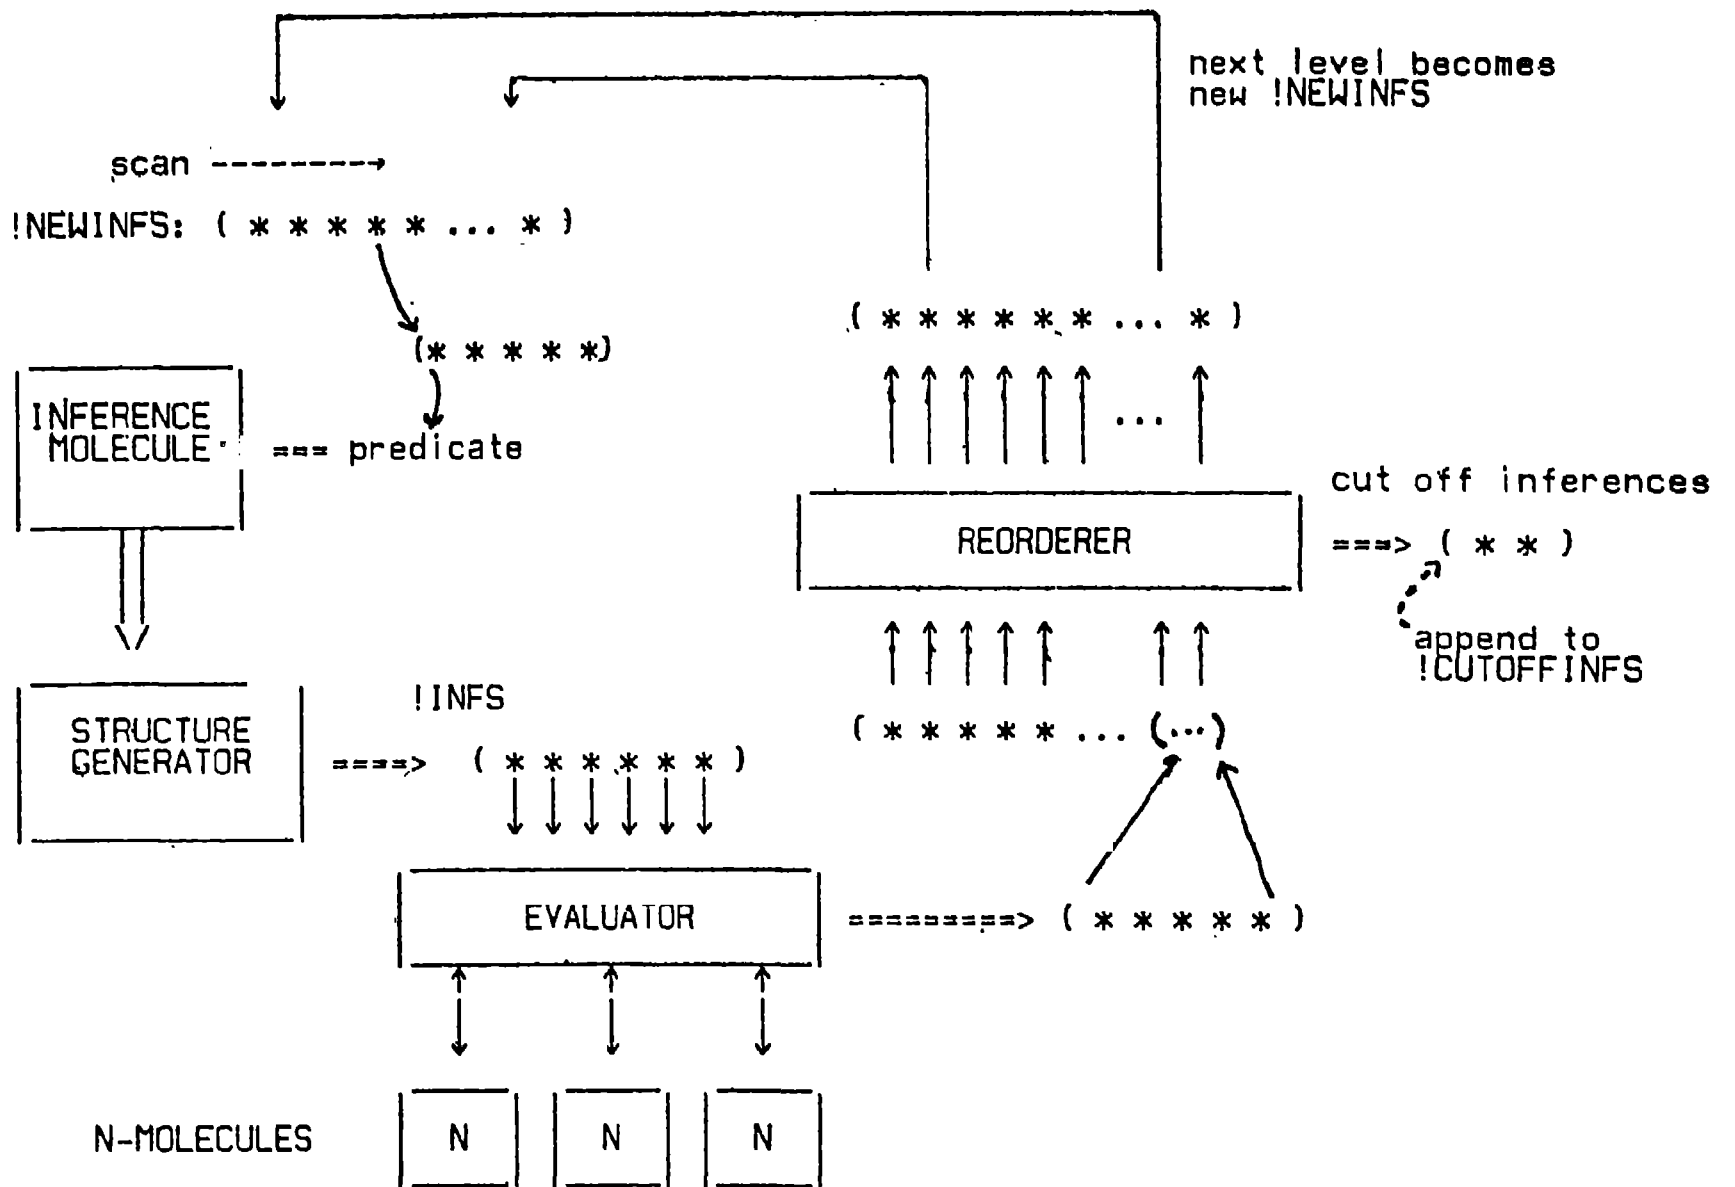


FIGURE 5

The inference monitor.

## 5. The Sixteen Theoretical Classes of Conceptual Inference.

It is phenomenological that most of the human language experience focuses on actions, their intended and/or resulting states, and the causality and enabling states which surround them. There seems to be an inescapable core of notions related to actions, causation and enablement which almost anyone who introspects long enough will independently discover. In his "Cold Warrior" model; Abelson (A1) was perhaps the first to attempt a computationally formal systematization of this fundamental core of meaning relations. It is of the utmost primacy in his system, which models the political ideologies and behavior patterns of a rabid right-winger, to discover and relate the underlying purposes, enablement and causality surrounding events in some hypothetical international scenerio or crisis. Again, in Schank et al's CD theory, the same emphasis arose more or less independently in a system of meaning representation for everyday utterances: causality, actions, state-changes and enablement were recurring themes. Not surprisingly, the same notions have emerged as central in my analysis of the inference reflex: over half of the 16 classes relate to this "action-intention-causality-enablement-knowledge" complex.

In the following descriptions of these 16 classes, keep in mind that all types of inference are applicable to every subcomponent of every utterance, and that the CM is essentially a parallel simulation. Also bear in mind that the inference evaluator is constantly performing matching operations on each new inference in order to detect interesting interactions between inference spheres. It should also be emphasized that conceptual inferences are probabilistic and predictive in nature, and that by making them in apparently wasteful quantities, the CM is not seeking one result or truth. Rather, inferential expansion is an endeavor which broadens each piece of information into its surrounding spectrum to fill out the information-rich situation to which the information-lean utterance might refer. The CM's gropings will resemble more closely the solution of a jigsaw puzzle than the more goal-directed solution of a crossword puzzle.

The following discussions can only sketch the main ideas behind each inference class. See (R1) for a more comprehensive treatment.

## 5.1 CLASS 1: SPECIFICATION INFERENCE

PRINCIPLE: The CM must be able to identify and attempt to fill in each missing slot of an incoming conceptual graph.

EXAMPLES: \*\*John was driving home from work. He hit Bill's cat.  
(inference) It was a car which John propelled into the cat.

\*\*John bought a chalk line.  
(inference) It was probably from a hardware store that John bought the chalk line.

### DISCUSSION:

Our use of language presupposes a tremendous underlying knowledge about the world. Because of this, even in, say, the most explicit technical writing, certain assumptions are made by the writer (speaker) about the comprehender's knowledge -- that he can fill in the plethora of detail surrounding each thought. In the CM, this corresponds to filling in all the missing conceptual slots in a graph.

The utility of such a process is twofold. First, CM failures to specify a missing concept can serve as a source of requests for more information (or goals to seek out that information by CM actions if CM is controlling a robot). Second, by predictively completing the graph by application of general pattern knowledge of the modeled world, novel relations among specific concepts and tokens will arise, and these can lead to potentially significant discoveries by other inferences.

To illustrate, a very common missing slot is the instrumental case. We generally leave it to the imaginative powers of the hearer to surmise the probable instrumental action by which some action occurred:

(husband to wife) I went to SEARS today.

(wife to husband), How? I had the car all day!

Here, wife fills in the instrumental slot as: "Husband drove a car to SEARS" (clearly relying on some specific heuristics, such as the distance from their

home to SEARS, etc.), and this led to her discovery of a contradiction. That she may have been premature in the specification (and had later to undo it) is of secondary importance to the phenomenon that she did so spontaneously.

In the CM specification inferences, as all inferences, are implemented in the form of structured programs which realize discrimination nets whose terminal nodes are concepts and tokens rather than inferences, as in general inference molecules. These specification procedures are called specifier molecules ("S-molecules"), and are quite similar to inference molecules. Fig. 6 shows a small prototype of the PROPEL specifier molecule which can predictively fill in the missing object of a PROPEL action, as in "John hit Pete." That particular "specifier atom" is sensitive to context along one simple dimension if the actor is known to be grasping an object (this prototype doesn't care whether it's a wet noodle or a bludgeon), at the time of the action, the molecule will infer that it was the grasped object which was propelled, as in "John picked up the flower pot. He hit Pete." Otherwise, the molecule will assume "hand of the actor". This is ridiculously oversimplified, but it represents a certain philosophy I will digress a moment to reveal.

I, as many other people (see W1, H1, C1, for instance), have come to believe that passive data structures are fundamentally awkward for representing knowledge in any detail, particularly for the purposes typified by

<pre> (SPROG *PROPEL* (UN V AC OB OF DT) (X1 X2 X3) ( (COND ( (NULL (CADR V)) (COND ( (AND (SETQ X1 (C (@ISA @_ @#HAND) @PART @_ AC))) (SETQ X2 (F1 (@*LOC* @_ X1))) (SETQ X3 !GLOBALFIND)) (SP V 2 X2 (LIST X3)) ) (T (SP V 2 X1 NIL) ) ) ) : (other specifier atoms go here) ) (RETURN V) )) </pre>	<pre>   This is a simplified specifier   molecule containing just an object   specifier atom. (NULL (CADR V)) is   test for lack of object specificatio   If unspecified, the atom locates   the hand of the actor, assigning   it to X1. It then checks to   see if anything is located in X1.   If something is found, it is bound   to X2, and the LOC structure which   expresses this information is   bound to X3. If nothing is located   in the actor's hand, his hand   itself (X1) is inferred. The   (LIST X3) in the first SP call   is the list of REASONS (just one   here) justifying the specification   of the object the actor was holding   as the object of the PROPEL. </pre>
---	--

FIGURE 6

The PROPEL specifier molecule.

this simple PROPEL example. The needs for "special case heuristics" in even such a modest operation as this quickly overtake one's prowess at devising "declarative" memory structures. Programs, on the other hand, are quick and to the point, quite flexible, and have as much "aesthetic potential" as even the most elegant declarative structures.

A life-size procedure for this very narrow process of specifying the missing object of a PROPEL action would obviously require many more tests for related contexts ("John was racing down the hill on his bike. He hit Bill.") But independent of the fidelity with which any given S-molecule executes its task, there is a very important claim buried both here and in the other inferential procedures in the CM. It is that there are certain central tasks in which the decision process must seek out the context, rather than context seeking out the appropriate decision process. In other words, much of the inference capability requires specialists who know a priori exactly what dimensions of context could possibly affect the generation of every potential inference, and these specialists carry out active probes to search for those dimensions before any inference is generated. I can imagine no "uniform context mechanism" which accounts for the human's diverse ability to attend to the relevant and ignore the superfluous. My conjecture is that the mechanism for contextual guidance of inference is highly distributed throughout the memory rather than centralized as a component of the memory's control structure.



## 5.2 CLASSES 2 and 3: RESULTATIVE and CAUSATIVE INFERENCES

PRINCIPLE: If an action is perceived, its probable resulting states should be inferred (RESULTATIVE). If a state is perceived, the general nature of its probable causing action (or a specific action, if possible) should be inferred (CAUSATIVE).

EXAMPLES: \*\*Mary hit Pete with a rock.

(inference) Pete probably became hurt. (RESULTATIVE)

\*\*Bill was angry at Mary.

(inference) Mary may have done something to Bill. (CAUSATIVE)

### DISCUSSION:

These two classes of inference embody the CM's ability to relate actions and states in causal sequences relative to the CM's models of causality. In addition to serving as the basis for MOTIVATIONAL inferences and contributing to the general expansion process, CAUSATIVE and RESULTATIVE inferences often achieve the rather exotic form of understanding I have termed "causal chain expansion." It is this process which makes explicit the oft-abbreviated statements of causality: language communicated predications of causality must always (if only subconsciously) be explained in terms of the comprehender's models of causality, and failures to do so signal a lack of understanding and form another source of CM queries for more information. Causal expansion successes on the other hand result in important intervening actions and states which draw out ("touch") surrounding context and serve as the basis for inferences in other categories. Appendix A contains the computer printout from MEMORY, tracing a causal expansion for "Mary kissed John because he hit Bill" in a particular context which makes the explanation plausible.

### 5.3 CLASS 4: MOTIVATIONAL INFERENCES

PRINCIPLE The desires (intentions) of an actor can frequently be inferred by analyzing the states (RESULTATIVE inferences) which result from an action he executes. These WANT-STATE patterns are essential to understanding and should be made in abundance.

EXAMPLES: \*\*John pointed out to Mary that she hadn't done her chores.  
(inference) Mary may have felt guilty. (RESULTATIVE)  
(inference) John may have wanted Mary to feel guilty.  
(MOTIVATIONAL)

\*\*Andy blew on the hot meat.  
(inference) Andy may have wanted the meat to decrease  
in temperature.

#### DISCUSSION:

Language is a dual system of communication in that it usually communicates both the actual, and, either explicitly or by inference, the intentional: Where the intentions of actors (the set of states they desire) are not explicitly communicated, they must be inferred as the immediate causality of the action. In the CM, candidates for MOTIVATIONAL inferences are the RESULTATIVE inferences  $\{R_i\}$  the CM can produce from an action A: for each RESULTATIVE inference  $R_i$  which the CM could make from A, it conjectures that perhaps the actor of A desired  $R_i$ .

Since the generation of MOTIVATIONAL inference is dependent upon the results of another class of inference (in general, the actor could have desired things causally removed by several inferences from the immediate results of his action), the MOTIVATIONAL inference process is implemented by a special procedure POSTSCAN which is invoked between "passes" of the main breadth-first monitor. These passes will be discussed more later.

Once generated, each MOTIVATIONAL inference will generally lead backward, via CAUSATIVE inferences, into an entire causal chain which lead up to the action. This chain will frequently connect in interesting ways with chains working forward from other actions.

## 5.4 CLASS 5: ENABLING INFERENCE

PRINCIPLE: Every action has a set of enabling conditions -- conditions which must be met for the action to begin or proceed. The CM needs a rich knowledge of these conditions (states), and should infer suitable ones to surround each perceived action.

EXAMPLES: \*\*John saw Mary yesterday,  
(inference) John and Mary were in the same general location sometime yesterday.

\*\*Mary told Pete that John was at the store  
(inference) Mary knew that John was at the store.

### DISCUSSION:

The example at the beginning of the paper contained a contradiction which could be discovered only by making a very simple enabling inference about the action of speaking (any action for that matter), namely that the actor was alive at the time! Enabling inferences can fruitfully lead from the known action through the enabling states to predications about other actions the actor might have performed in order to set up the enabling states for the primary action. This idea is closely related to the next class of inference.

## 5.5 CLASS 6: ACTION PREDICTION INFERENCES

PRINCIPLE: Whenever some WANT STATE of a potential actor is known, predictions about possible actions the actor might perform to achieve the state should be attempted. These predictions will provide potent potential points of contact for subsequently perceived actions.


EXAMPLES: \*\*John wants some nails.  
(inference) John might attempt to acquire some nails.  
\*\*Mary is furious at Rita.  
(inference) Mary might do something to hurt Rita.

### DISCUSSION:

Action prediction inferences serve the inverse role of MOTIVATIONAL inferences, in that they work forward from a known WANT STATE pattern into predictions about future actions which could produce the desired state. Just as a MOTIVATIONAL inference relies upon RESULTATIVE inferences, an ACTION PREDICTION inference relies upon CAUSATIVE inferences which can be generated from the state the potential actor desires. Because it is often impossible to anticipate the specific causing action, ACTION PREDICTION inferences typically will be more general expectancies for a class of possible actions. In the nails example above, the general expectancy is simply that John may do something which normally causes a PTRANS -(in CD terminology, a change of location of some object) of some nails from somewhere to himself. Often the nature of the desired state is such that some specific action can be predicted ("John is hungry... John will ingest food.") By making specific action predictions, a new crop of enabling inferences can be predicted ("John must be near food.", etc.), and those conditions which cannot be assumed to be already satisfied can serve as new WANT-STATES of the actor. Thus it is through MOTIVATIONAL, ACTION PREDICTION and ENABLING inferences that the CM can model (predict) the problem-solving behavior of each actor.

Predicted actions which match up with subsequently perceived conceptual input serve as a very real measure of the CM's success at piecing together connected discourse and stories. I suspect in addition that ACTION PREDICTION inferences will play a key role in the eventual solutions of the "contextual guidance of inference" problem. Levy (L1) has some interesting beginning thoughts on this topic.

Fig. 7 illustrates the ACTION PREDICTION inference cycle.



What the action prediction inference process tries to do:

$A_i E_j$  is enabling state  $j$  for action  $i$ .

FIGURE 7

The action prediction inference process.

## 5.6 CLASS 7: ENABLEMENT PREDICTION INFERENCES.

PRINCIPLE: If a potential actor desires a state which is a common enabling condition for some specific action, then it can be inferred that the actor might wish to execute that action.

EXAMPLES: \*\* Mary asked John to turn on the light.  
(inference) Mary probably wants to see something.  
\*\* Andy wants the meat to be cool.  
(inference) Andy might want to eat the meat.

### DISCUSSION:

Inferences in this class are, in a sense, the inverse of ENABLING inferences, because they attempt to predict an action from an enabling state known to be desired by a would-be actor. Whereas an ACTION PREDICTION inference predicts a possible future action to fulfill the desired state, enablement prediction draws out the motivation of the desire for the state by identifying a probable action the state would enable. Although (as with ACTION PREDICTION inference) it will frequently happen that no specific action can be anticipated (since most states could enable infinitely many specific actions), it is nevertheless possible to form general predictions about the nature of (restrictions on) the enabled action. If, for example, John walks over to Mary, then a RESULTATIVE inference is that he is near Mary, and a MOTIVATIONAL inference is that he wants to be near MARY. At this point an ENABLEMENT PREDICTION inference can be made to represent the general class of interactions John might have in mind. This will be of particular significance if, for instance, the CM knows already that John had something to tell her, since then the inferred action pattern would match quite well the action of verbal communication in which the state of spatial proximity plays a key enabling role.

## 5.7 CLASS 8: FUNCTION INFERENCE

PRINCIPLE: Control over some physical object P is usually desired by a potential actor because he is engaged in an algorithm in which P plays a role. The CM should attempt to infer a probable action from its knowledge of P's normal function.

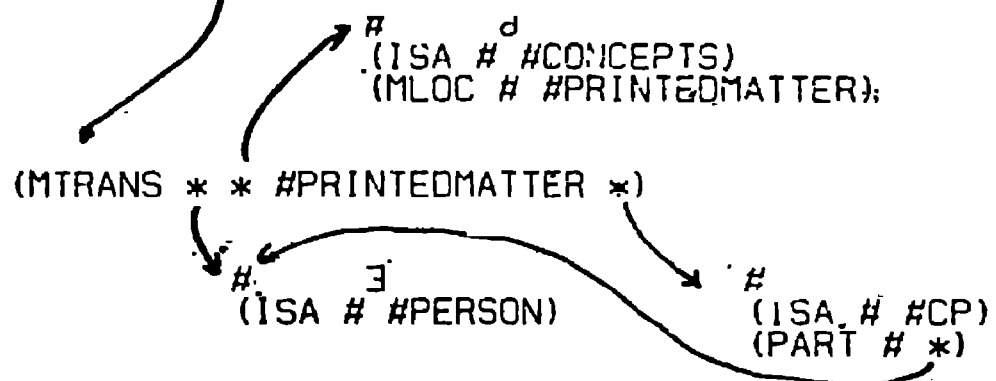
EXAMPLES: \*\*Mary wants the book.  
(inference) Mary probably wants to read the book.  
\*\*John wants a knife.  
(inference) John probably wants to cut something with the knife.  
\*\*Bill likes to pour sundaes down girls' dresses.  
Bill asked Pete to hand him the sundae...

### DISCUSSION:

Function inferences form a very diverse, rather colorful subclass of ENABLEMENT PREDICTION inference. The underlying principle is that desire of immediate control over an object is usually tantamount to a desire to use that object in the normal function of objects of that type, or in some function which is peculiar to the object and/or actor (third example above). In the CM, normal functions of objects are stored as (NFCT X Y) patterns, as in Fig. 8 for things that are printed matter. Before applying NFCT patterns, the CM first checks for unusual relations involving the specific actor and specific object (by excluding paths which include the normal ISA relations between, say sundae and food). Thus, that Bill is known to require sundaes for slightly different algorithms from most people will be discovered and used in the prediction. The result of a FUNCTION inference is always some predicted action, assumed to be part of some algorithm in which the actor is engaged,

---

(NFCT #PRINTEDMATTER \*)



The memory structure which stores  
the normal function of printed matter.

---

FIGURE 8

A "normal-function-of"  
memory structure.



## 5.8 CLASSES 9 and 10: MISSING ENABLEMENT and INTERVENTION INFERENCES

PRINCIPLE: If a would-be actor is known to have been unsuccessful in achieving some action, it is often possible to infer the absence of one of the action's enabling states (MISSING ENABLEMENT). If a potential actor is known to desire that some action cease, it can be predicted that he will attempt to remove one or more enabling states of the action (INTERVENTION).

EXAMPLES: \*\*Mary couldn't see the horses finish.  
(inference) Something blocked Mary's view. (MISSING ENABLEMENT)  
She cursed the man in front of her...

\*\*Mary saw that Baby Billy was running out into the street.  
(inference) Mary will pick Billy off the ground (INTERVENTION)  
She ran after him...

### DISCUSSION:

Closely related to the other enabling inferences, these forms attempt to apply knowledge about enablement relations to infer the cause of an action's failure (in the case of MISSING ENABLEMENT), or to predict a WANT NOT-STATE which can lead by action prediction inference to possible actions of intervention on the part of the WAnTer. In the second example above, Mary (and the CM) first must realize (via RESULTATIVE inferences) the potentially undesirable consequences of Billy's running action (i.e., possible NEGCHANGE for Billy). From this, the CM can retrace, locate the running action which could lead to such a NEGCHANGE, collect its enabling states, then conjecture that Mary might desire to annul one or more of them. Among them for instance would be that Billy's feet be in intermittent PHYS-CONT with the ground. From the (WANT (NOT (PHYSCONT FEET GROUND))) structure, a subsequent ACTION PREDICTION inference can arise, predicting that Mary might put an end to (PHYSCONT FEET GROUND). This will in turn require her to be located near Billy, and that prediction will match the RESULTATIVE inference made from her directed running (the next utterance input), knitting the two thoughts together.

## 5.9 CLASS 11: KNOWLEDGE PROPAGATION INFERENCES

PRINCIPLE: Based on what the CM knows an actor to know, it can often infer other knowledge which must also be available to the actor. Since most conceptual inferences involve the intentions of actors, this modeling of knowledge is crucial.

EXAMPLES: \*\*John saw Mary beating Pete with a baseball bat.  
(inference) John probably knew that Pete was getting hurt.  
\*\*Betty asked Bill for the aspirin.  
(inference) Bill probably surmised that Betty wasn't feeling well.

### DISCUSSION:

Modeling the knowledge of potential actors is fundamentally difficult. Yet it is essential, since most all intention/prediction-related inferences must be based in part on guesses about what knowledge each actor has available to him at various times. The CM currently models others' knowledge by "introspecting" on its own: assuming another person P has access to the same kinds of information as the CM, P might be expected to make some of the same inferences the CM does. Since the CM preserves a logical connectivity among all its inferred structures (by the REASONS and OFFSPRING properties of each structure), after inferences of other types have arisen from some unit of information U, the CM can return, determine who knew the original fact U, locate U's OFFSPRING (those other memory structures which arose by inference from U), then infer that P may also be aware of each of the offspring. As with MOTIVATIONAL inferences (which rely on the RESULTATIVE inferences from a structure), KNOWLEDGE PROPAGATION inferences are implemented in the procedure POSTSCAN which runs after the initial breadth-first inference expansion by the monitor.

Modeling others' knowledge demands a rich knowledge of what is normal in the world. ("does John Smith know that kissing is a sign of affection?"). In fact, all inferences must rely upon default assumptions about normality, since most of the CM's knowledge (and presumably a human's) exists in the form of general patterns, rather than specific relations among specific concepts and tokens. The next class of inference implements my belief that patterns, just as inferences, should be realized in the CM by active programs rather than by passive declarative data structures.

## 5.10 CLASS 12: NORMATIVE INFERENCE

PRINCIPLE: The CM must make heavy reliance upon programs which encode commonsense pattern information about the modeled world. When the retrieval of a sought-after unit of information fails, the relevant normality program should be executed on (pattern applied to) that information to assess its likelihood in the absence of explicit information.

EXAMPLES: \*\*Does John Smith own a book?  
(inference) Probably so; middle-class business executives normally own books.  
\*\*Was John likely to have been asleep at 3 pm yesterday?  
(inference) Most likely not, since he has a normal day-time job, and yesterday was a workday.

### DISCUSSION:

There are several low-level information retrieval procedures in the CM which search for explicit information units as directed by specific inference molecules. Such searches are on the basis of form alone, and successes result in precise matches, while failures are total. If there were no recourse for such failures, the CM would quickly grind to a halt, being unable to make intelligent assumptions. There must be some more positive and flexible mechanism to ameliorate "syntactic" lookup failures.

In the CM, this ability to make intelligent assumptions is implemented by having the low-level lookup procedures defer control to the appropriate normality molecule (N-molecule) which will perform systematic tests organized in single-response discrimination nets, to the unlocatable information. The goal is to arrive at a terminal node in the net where a real number between 0 and 1 is located. If some sequence of tests leads to such a number, the N-molecule returns it as the assessed likelihood ("compatibility" in fuzzy logic terminology (Z1)) of X being true.

Although the test in the N-molecules are themselves discrete, they result in the fuzzy compatibility. The point of course is that the tests can encode quite diverse and very specific heuristics peculiar to each small domain of patterns: For instance, based on known (or N-molecule inferrable -- one N-molecule can call upon others in its testing process!) features of either John or the hammer, we would suspect the compatibility of each of the following four conjectures to form a decreasing sequence:

1. John Smith owns something. (very likely, but dependent on his age, society in which he lives, etc.)
2. John Smith owns a hammer. (probably, but potentially related to features of John, such as his profession)
3. John Smith owns a claw hammer with a wooden handle. (maybe, but again dependent on features of John and models of hammers in general -- i.e., how likely is any given hammer to have a claw and wooden handle?)
4. John Smith owns a 16 oz. Stanley claw hammer with a steel-reinforced wooden handle and a tack puller on the claw. (likelihood is quite low unless the N-molecule can locate some specific hints, such as that John usually buys good equipment, etc.)

A successful N-molecule assessment results in the creation of the assessed information as a permanent, explicit memory structure whose STRENGTH is the assessed compatibility. This structure is the normative inference.

One is quickly awed by his own ability to rate (usually quite accurately) commonsense conjecture such as these, and the process seems usually to be quite sensitive to features of the entities involved in the conjecture. It is my feeling that important insights can be gained via a more thorough investigation of the "normative inference" process in humans.

Another role of N-molecules is mentioned in (R1) with respect to the inference-reference cycle I will describe shortly. Fig. 9 shows the substance of a prototype N-molecule for assessing dependency structures of the form (OWN P X) (person P owns object X).



---

is P a member of a pure communal society, or is it an infant?  
   if so, very unlikely that P owns X  
   otherwise, does X have any distinctive conceptual features?  
     if so, assess each one, form the product of likelihoods, and call it  
     M. M will be used at the end to mitigate the likelihood which would  
     normally be assigned.

is X living?  
   if so, is X a person?  
     if so, is P a slave owner; and does X possess characteristics  
     of a slave? if so, likelihood is low but non-zero  
     otherwise likelihood is zero  
   otherwise, is X a non-human animal or a plant?  
     if so, is X domestic in P's culture?  
       if so, does P have a fear of X's or is  
       P allergic to X's of this type?  
       if so, likelihood is low  
       otherwise, likelihood is moderate  
     otherwise, is X related to actions P does in any special  
     way?  
       if so, likelihood is low, but non-zero  
       otherwise, likelihood is near-zero

otherwise, does X have a normal function?  
   if so, does P do actions like this normal function? (Note here  
   that we would want to look at P's profession, and actions commonly  
   associated with that profession.)  
     if so, likelihood is moderately high  
   otherwise, is X a common personal item?  
     if so, is it's value within P's means  
       if so, likelihood is high  
       if not, likelihood is low, but non zero  
   otherwise, is X a common household item?  
     if so, is P a homeowner?  
       if so, is X within P's means?  
       if so, likelihood is high  
       otherwise, likelihood is moderate  
     otherwise, likelihood is low, but non-zero

and so on ...

How we might go about  iding  
 whether person P owns an 

---

FIGURE 9

The normality-molecule  
 discrimination network for the  
 pattern (OWNS P X).

## 5.11 CLASS 13: STATE DURATION INFERENCES

PRINCIPLE: Most interesting states in the world are transient. The CM must have the ability to make specific predictions about the expected (fuzzy) duration of an arbitrary state so that information in the CM can be kept up to date.


EXAMPLES: \*\*John handed Mary the orange peel.  
(tomorrow) Is Mary still holding the orange peel?  
(inference) Almost certainly not.  
\*\*Rita ate lunch a half hour ago.  
Is she hungry yet?  
(inference) Unlikely.

### DISCUSSION:

Time features of states relate in critical ways to the likelihood those states will be true at some given time. The thought of a scenario wherein the CM is informed that Mary is holding an orange peel, then 50 years later, uses that information in the generation of some other inference is a bit unsettling! The CM must simply possess a low-level function whose job it is to predict normal durations of states based on the particulars of the states, and to use that information in marking as "terminated" those states whose likelihood has diminished below some threshold.

My conjecture is that a human notices and updates the temporal truth of a state only when he is about to use it in some cognitive activity -- that most of the transient knowledge in our heads is out of date until we again attempt to use it in, say, some inference. Accordingly, before using any state information, the CM first filters it through the STATE DURATION inference process to arrive at an updated estimate of the state's likelihood as a function of its known starting time (its TS feature, in CD notation).

The implementation of this process in the CM is as follows: an (NDUR S ?) structure is constructed for the state S whose duration is to be predicted, and this is passed to the NDUR specifier molecule. The NDUR S-molecule applies discrimination tests on features of the objects involved in S. Terminal nodes in the net are duration concepts (typically fuzzy ones), such as #ORDERHOUR, #ORDERYEAR. If a terminal node can be successfully reached, thus locating such a concept D, the property CHARF (Characteristic time function) is retrieved from D's property list. CHARF is a step function of STRENGTH vs. the amount of time some state has been in existence (Fig. 10). From this function a STRENGTH is computed for S and becomes S's predicted likelihood. If the STRENGTH turns out to be sufficiently low, a (TF S now) structure is predictively generated to make S's low likelihood explicit. The STATE\_DURATION inference thus acts as a cleansing filter on state information which is fed to various other inference processes.



A typical STRENGTH function for fuzzy duration #ORDERHOUR.

---

( (W1MAX S1) (W2MAX S2) (W3MAX S3) ... (WnMAX Sn) )

$0 \leq T'-T < W1MAX$  has strength S1  
 $W1MAX \leq T'-T < W2MAX$  has strength S2  
 $T'-T \geq WnMAX$  has strength 0

The format of a fuzzy duration concept's step function.

---

FIGURE 10

A typical characteristic STRENGTH function for the state-duration inference process.

---



## 5.12 CLASSES 14 and 15: FEATURE and SITUATION INFERENCES

PRINCIPLE: Many inferences can be based solely on commonly observed or learned associations, rather than upon "logical" relations such as causation, motivation, and so forth. In a rough way, we can compare these inferences to the phenomenon of visual imagery which constructs a "picture" of a thought's surrounding environment. Such inferences should be made in abundance.

EXAMPLES: \*\* Andy's diaper is wet.  
(inference) Andy is a youngster. (FEATURE)  
\*\* John was on his way to a masquerade.  
(inference) John was probably wearing a costume. (SITUATION)

### DISCUSSION:

Many "associative" inferences can be made to produce new features of an object (or aspects of a situation) from known features. If something wags its tail, it is probably an animal of some sort, if it bites the mailman's leg, it is probably a dog, if it has a gray beard and speaks, it is probably an old man, if it honks in a distinctive way, it is probably some sort of vehicle, etc. These classes are inherently unstructured, so I will say no more about them here, except that they frequently contribute features which help clear up reference ambiguities and initial reference failures.

---

## 5.13 CLASS 16: UTTERANCE INTENT INFERENCE

PRINCIPLE: Based on the way a thought is communicated (especially the often telling presence or absence of information), inferences can be made about the speaker's reasons for speaking.

EXAMPLES:

- \*\*Don't eat green gronks.  
(inference) Other kinds of gronks are probably edible
- \*\*Mary threw out the rotten part of the fig.  
(inference) She threw it out because it was rotten.
- \*\*John was unable to get an aspirin.  
(inference) John wanted to get an aspirin.
- \*\*Rita like the chair, but it was green.  
(inference) The chair's color is a negative feature to Rita (or the speaker).

### DISCUSSION:

I have included this class only to represent the largely unexplored domain of inferences drawn from the way a thought is phrased. The CM will eventually need an explicit model of conversation, and this model will incorporate inferences from this class. Typical of such inferences are those which translate the inclusion of referentially superfluous features of an object into an implied causality relation (the fig example), those which infer desire from failure (the aspirin example) those which infer features of an ordinary X from features of special kinds of X. (the gronk example), and so forth. These issues will lead to a more goal directed model than I am currently exploring.

---

## 6. Summary of the Inference Component

I have now sketched 16 inference classes which, I conjecture, lie at the core of the human inference reflex. The central hypothesis is that a human language comprehender performs more subconscious computation on meaning structures than any other theory of language comprehension has yet acknowledged. When the current CM is turned loose, it will often generate upwards of 100 inferences from a fairly banal stimulus such as "John gave Mary the book." While most are irrefutable, they are for the most part quite mundane and "uninteresting" to a critical human observer, and are, after the fact, "wasteful." But change the context and the banal becomes salient -- even crucial -- while the crucial can become irrelevant! I can see no other mechanism for explaining contextual interaction of information than this spontaneous, subconscious groping.

I should perhaps briefly address the adequacy and applicability of the inference classes in the current model. There is undoubtedly a number (16, 32, 64?) of equally interesting inference classes I have ignored or overlooked. But I feel the number is not large, and that other classes will submit to the same sorts of systematization as described here. While the piecemeal examples I have used to illustrate the various inferences were not drawn from any coherent domain such as a "blocks world" (W1) -- and this is a weakness -- I believe the net result (these inference classes and their control structure) will prove central to any restricted domain which involves volitional actors. It is a current challenge to find such a restricted, yet interesting, domain to which these ideas can be transplanted and applied in slightly more goal-directed environments.

I want to describe now how two other important language components reference and implicit concept activation -- aid and abet the inference reflex.

## 7. The Inference-Reference Relaxation Cycle in Conceptual Memory.

A competent listener is frequently incapable of instantaneously identifying the referent (concept or token in memory) of a language construction (noun group, pronoun, etc.). Yet an attentive listener seldom fails eventually to identify the intended referent, and he will seldom lose information because of the reference delay. Furthermore, incorrect reference decisions are empirically few and far between. I believe that these phenomena are intimately related to the inference reflex.

In the CM, initial reference attempts are made for concepts and tokens from descriptive sets -- collections of conceptual features gleaned from an utterance by Riesbeck's conceptual analyzer (R2). Fig. 11 illustrates the descriptive set for the "the big red dog who ate the bird." Potential memory concepts and token referents are identified by an intersection search procedure which locates memory objects whose features satisfy all the features of the descriptive set. Such a search will result in either (a) a unique identification of some memory entity, (b) a failure to locate any satisfactory entities, or (c) a set of candidates, one of which is the probable referent. Case (a) requires no decision, but (b) and (c) do. In either case, a new, possibly temporary token T is created and, for case (b), T receives as its initial occurrence set the descriptive set identically. In case (c), where a set of candidates can be located, T receives the set of features lying in the intersection of all candidates' occurrence sets (this will be at least the descriptive set). In either case, the CM then has an internal token to work with, allowing the conceptual graph in which references to it occur to be tentatively integrated into memory structures.

The inference reflex I have described then generates all the various inferences, and eventually returns to its quiescent state. One byproduct of the inferencing is that the occurrence set of each memory object involved in the original structures will emerge with a possible enhanced occurrence set which may contain inferred information sufficient either (1) to identify the temporary token of category (b) above, or (2) to narrow the set of candidates associated with the temporary token of category (c) (hope-

---

```

X: { (ISA X #DOG)
      (COLOR X #RED)
      (RELSIZE X #LARGE)
      (Y: (INGEST X
           { (ISA # #BIRD) (REF # *THE*) }
           { (ISA # #MOUTH) }
           { (ISA # #STOMACH) }
           )
          (TIME Y { (ISA # #TIME) (BEFORE # #NOW) } )
      )
      (REF X *THE*)
    }

```

Descriptive set for "The big red dog who ate the bird"

---

FIGURE 11

An example of a descriptive set.

fully to exactly one). Thus, when the inference reflex has ceased, the CM re-applies the reference intersection algorithms to each unidentified token to seek out any inference-clarified references. Successful identifications at this point result in the merging (by the same structure merger mentioned earlier) of the temporary token's occurrence set with the identified token's occurrence set, thus preserving all information collected to that point about the temporary token. (Implicit in the merge operation is the substitution of all references to the temporary token by references to the identified one.) If, on the other hand, the results of inferencing serve only to narrow the candidate set of case (c), the occurrence sets of the remaining candidates are re-intersected, and (if this increases the size of the set) the set is re-attached to the temporary token. In either case progress has been made.

Now comes a key point. If any referents were in fact identified on this second attempt (making their entire occurrence sets accessible), or if any candidate set decreases caused new features to be associated with the temporary tokens, then there is the possibility that more inferences (which can make use of the newly-accessible features) can be made. The CM thus re-applies the inference reflex to all memory structures which were produced on the first pass. (The monitor is conditioned not to duplicate work already done on the first pass.) But a potential byproduct of the second pass is further feature generation which can again restrict candidate sets or produce positive identifications. This inference-reference interaction can proceed until no new narrowings or identifications occur; hence the term "relaxation cycle." Fig. 12 illustrates two examples of this phenomenon which are handled by the current CM, and Appendix B contains the computer trace of the second example.

EXAMPLE.1 Andy Rieger is a youngster.  
 Andy Moorer is an adult.  
 Andy's diaper is wet.  
 INFERENCE-REFERENCE: Andy Rieger's diaper is wet.

EXAMPLE 2 John was in Palo Alto yesterday.  
 Jenny Jones was in Palo Alto yesterday.  
 Jenny Smith was in France yesterday.  
 Bill loves Jenny Jones.  
 Bill saw John kiss Jenny yesterday.  
 INFERENCE-REFERENCE, FIRST PASS: It was Jenny Jones  
 that John kissed.  
 INFERENCE-REFERENCE, SECOND PASS: Bill felt anger  
 toward John

FIGURE 12a

Two examples of inference-reference interaction.

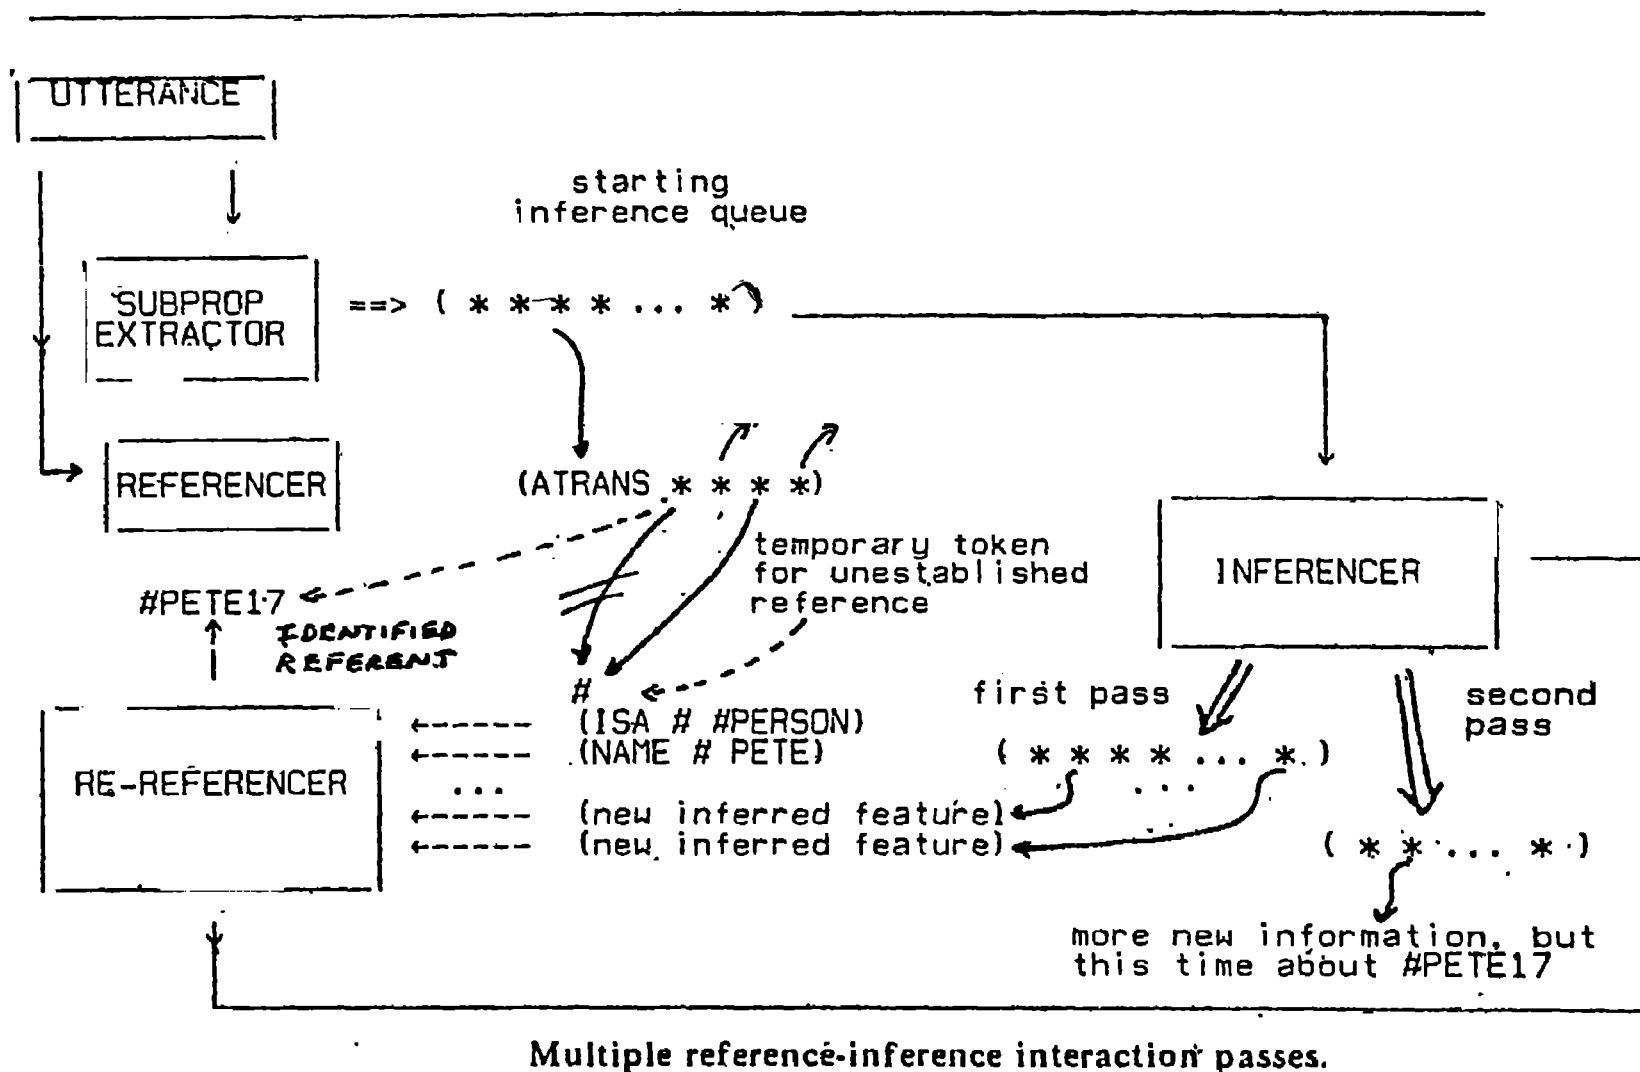


FIGURE 12b

The inference-reference relaxation cycle.

## 8. Word Sense Promotion and Implicit Concept Activation in the Conceptual Memory

Another byproduct of the generation of an abundance of probabilistic conceptual patterns from each input is that many related concepts and tokens implicitly involved in the situation are activated, or "touched." This can be put to use in two ways.

First, implicitly touched concepts can clarify what might otherwise be an utterly opaque subsequent reference. If, for instance, someone says (outside of a particular context): "The nurses were nice", you will probably inquire "What nurses?" If, on the other hand, someone says: "John was run over by a milk truck. When he woke up the nurses were nice" you will experience neither doubt about the referents of "the nurses", nor surprise at their mention. I presume that a subconscious filling-out of the situation "John was run over by a milk truck" implicitly activates an entire set of conceptually relevant concepts, "precharging" ideas of hospitals and their relation to patients.

Other theories founded more on concept associationism than conceptual inference have suggested that such activation occurs through word-word or concept-concept free associations (see (A2) and (Q1) for instance). While these more direct associations play an undoubted role in many language functions, it is my belief that these straight associative phenomena are not fundamentally powerful enough to explain the kind of language behavior underlying the nurse example. It is more often than not the "gestalt" meaning context of an utterance which restricts the kinds of meaningful associations a human makes. In contrast to the nurse example above, most people would agree that the reference to "the nurses" in the following situation is a bit peculiar:

In the dark of the night, John had wallowed through the knee-deep mud to the north wall of the deserted animal hospital. The nurses were nice.



A simple hospital-nurses association model cannot account for this. On the other hand, those concepts touched by the more restrictive conceptual inference patterns would presumably be quite distant from the medical staff of a hospital in this example, thus explaining the incongruity.

Related to this idea of concept activation through conceptual inference structures is another mechanism which, I presume, underlies comprehenders' ability to select (almost unerringly) the proper senses of words in context during the linguistic analysis of each utterance. This mechanism is frequently called word sense promotion, and its exact nature is one of the major conundrums of language analysis. It underlies our ability to avoid -- almost totally -- backing up to reinterpret words. It is as though at each moment during our comprehension we possess a dynamically shifting predisposition toward a unique sense of just about any word we are likely to hear next. Fig. 13 contains some illustrations of this phenomenon.

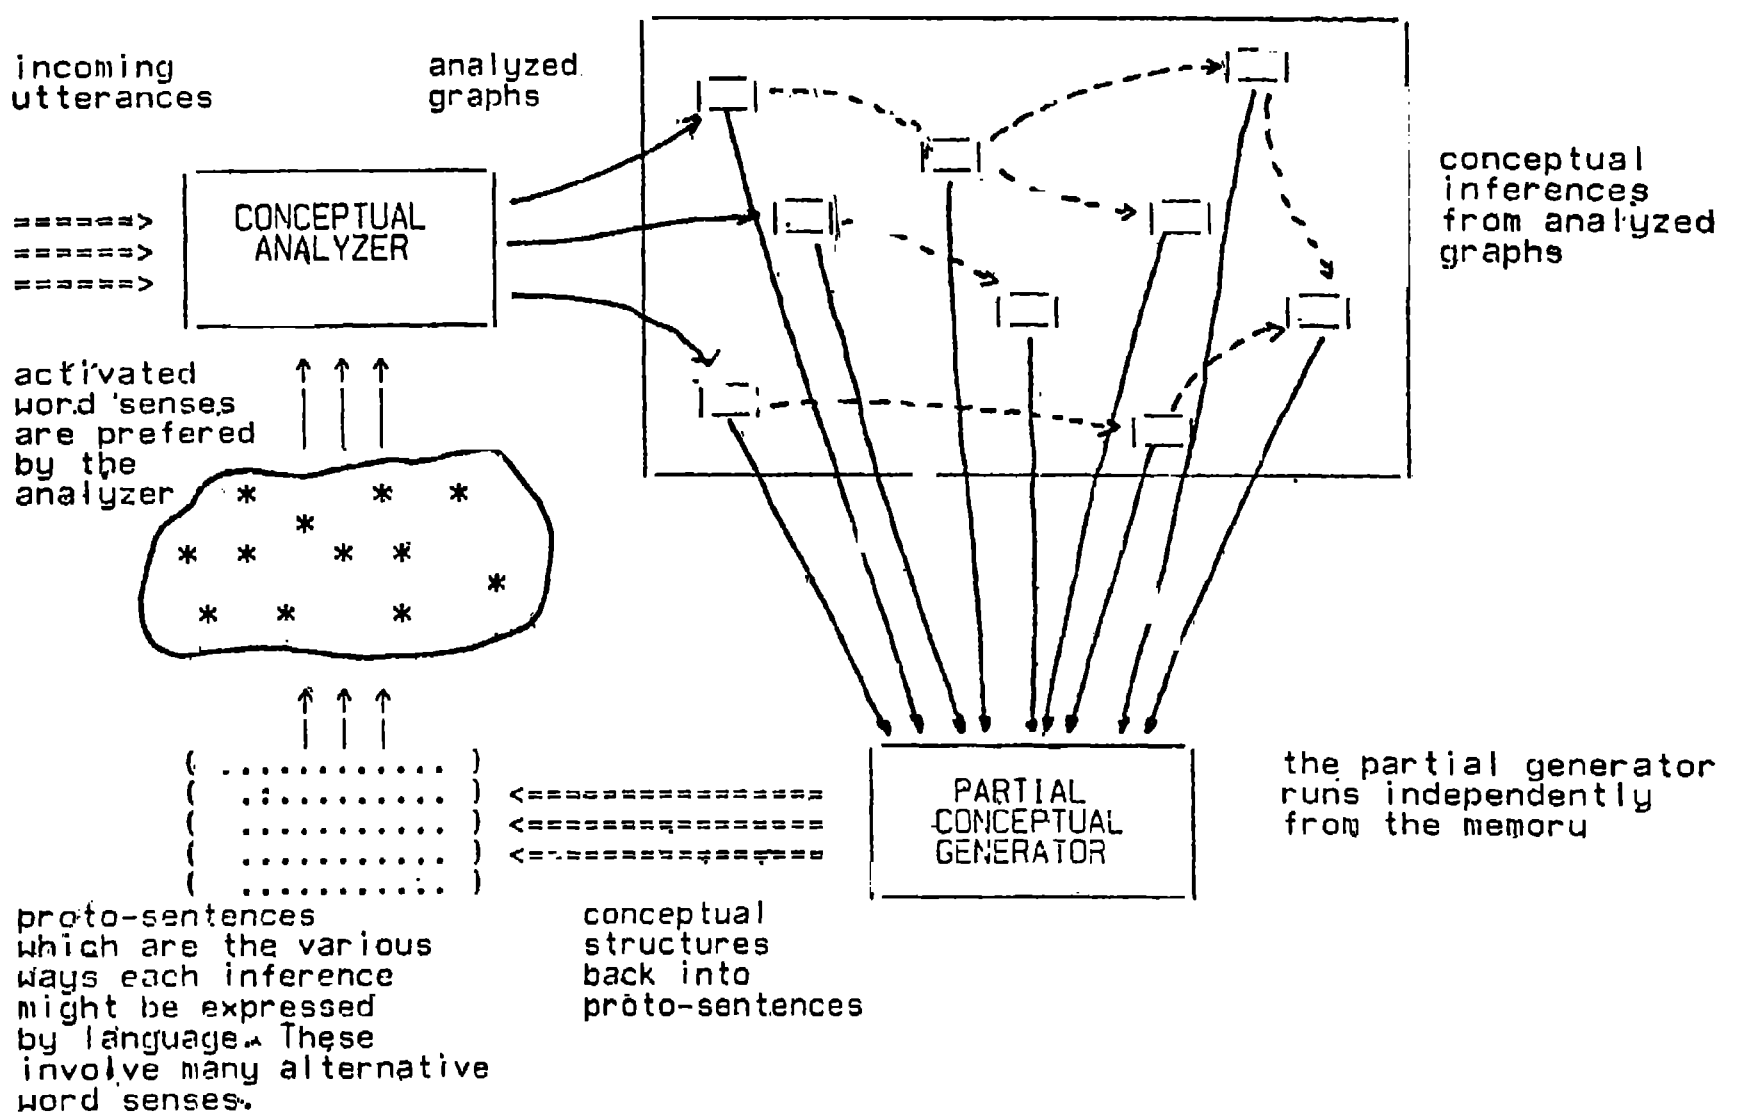
I have only a thought (which I plan to develop) on this issue. At each instant in the CM, there is a powerful inference momentum which is the product of conceptual inferences. Obviously, these concepts which the inference patterns touch will correspond to senses of words. These senses can be "promoted" in the same way implicit activation promotes certain referents. This is a partial explanation of word sense promotion. Suppose, however, that in addition the CM had an independent parallel process which took each inference as it arose and mapped it back into a near-language "proto-sentence", a linear sequence of concepts which is almost a sentence of the language, except that the actual word realizations of each concept have not yet been chosen. In other words, a generation process (see (G1) for example) would be applied to each inference, but would be stopped short of the final lexical substitutions of word senses. By precharging all the senses of the various words which could be substituted in such a proto-sentence, the CM would have a word sense "set" which would be a function of the kind of restrictive inferential context which I feel is so vital to the process of analysis. This idea is obviously computationally exorbitant, but it might model a very real mechanism. We often catch ourselves subvocalizing what we expect to hear next (especially while listening to an annoyingly slow speaker), and this is tantalizing evidence that something like a proto-sentence generator is thrashing about upstairs.

EXAMPLE 1: (CONTEXT) John asked Mary which piece of fruit she wanted.  
 (SENSE) Mary picked the apple.  
 versus (CONTEXT) Mary climbed the apple tree.  
 (SENSE) Mary picked the apple.

EXAMPLE 2: (CONTEXT) John was in a meadow.  
 (SENSE) The grass smelled good.  
 versus (CONTEXT) John was looking forward to getting high  
 (SENSE) The grass smelled good.

EXAMPLE 3: (Riesbeck's example (R2))  
 John went on a hunting trip. He shot two bucks.  
 It was all he had!

**FIGURE 13a**  
 Examples of word sense promotion.



Mapping inferences back into proto-sentences, activating many word senses.

**FIGURE 13b**

Mapping inferences back into proto-sentences, activating many word senses.

## 9. Conclusion

Any theory of language must also be a theory of inference and memory. It does not appear to be possible to "understand" even the simplest of utterances in a contextually meaningful way in a system in which language fails to interact with a language-free belief system, or in a system which lacks a spontaneous inference reflex.

One very important theoretical issue concerns exactly how much "inference energy" is expended before the fact (prediction, expectation) versus how much is expended after the fact to clear up specific problems of how the utterance fits the context. My belief is that there is a great deal of exploratory, essentially undirected inferencing which is frequently overlooked and which cannot be repressed because it is the language-related manifestation of the much broader motivational structure of the brain. Rather than argue at an unsubstantiatable neurophysiological level, I have compiled evidence for this hypothesis within the domain of language. I believe, however, that spontaneity of inference pervades all other modes of perception as well, and that quantity --as much as quality -- of spontaneous inference is a necessary requirement for general intelligence.

APPENDIX A: CAUSAL.CHAIN EXPANSION COMPUTER EXAMPLE

WORKING "FORWARD", GENERATING  
RESULTATIVE INFERENCES FROM  
THE PROPEL UNDERLYING "HIT":

(CONTEXT) Bill swiped Mary's book.  
(CAUSAL) Mary kissed John because  
he hit Bill.

- resultative \* John propelled his hand toward Bill
- ↓
- resultative \* John's hand came into physical contact with Bill
- ↓
- resultative \* Because it was propelled, the physical contact was probably forceful
- ↓
- resultative \* Bill probably suffered a negative change in physical state
- ↓
- resultative \* Because Bill suffered a negative change, and Mary felt a negative emotion toward Bill at the time, Mary might have experienced a positive change in joy
- ↓
- \* Because Mary may have experienced this positive change, and because it was John whose action indirectly caused her positive change, she might feel a positive emotion toward John

WORKING "BACKWARD", GENERATING  
CAUSATIVE INFERENCES FROM THE  
PHYSCONT UNDERLYING "KISS":

POINT OF CONTACT:  
Mary probably feels a  
positive emotion toward  
John.

- \* Mary's placing her lips in contact with John was probably caused by Mary feeling a positive emotion toward John.
- ↑
- causative
- ↓
- \* Mary's lips were in contact with John

Figure 5-21. One explanation of why Mary's kissing was related to John's hitting.

(MARY KISSED JOHN BECAUSE HE HIT BILL)

```
((CON ((CON ((ACTOR (JOHN) <=>
(*PROPEL*) OBJECT (*PHYSOBJ* SPEC
(*U*)) FROM (JOHN) TO (BILL)) TIME
(TIM01)) <=> ((ACTOR (*PHYSOBJ* SPEC
(*U*)) <=> (*PHYSCONT* VAL (BILL)))
TIME (TIM01)))) <=> ((CON ((ACTOR
(MARY) <=> (*DO*) TIME (TIM02) SPEC
(*U*)) <=> ((ACTOR (*LIPS* PART (MARY))
<=> (*PHYSCONT* VAL (JOHN))) TIME
(TIM02)))))) TIME (TIM01))
```

```
(TIM00 ((VAL T-0))
(TIM01 ((BEFORE TIM02 X)))
(TIM02 ((BEFORE TIM00 X)))
```

This is the input sentence. Its underlying conceptual graph is shown next.

```
((CAUSE ((CAUSE ((*PROPEL* (#JOHN1)
(C0013) (#JOHN1) (#BILL1)) (TIME
(C0016))) ((*PHYSCONT* (C0013) (#BILL1))
(TIME (C0016)))))) ((CAUSE ((*DO*
(#MARY1) C0010) (UNSPECIFIED ). (TIME
(C0017))) ((*PHYSCONT* (C0021) (#JOHN1))
(TIME _ (C0017)))))) (TIME _ (C0016)))
```

This is the partially integrated memory structure, after references have been established. No reference ambiguity is assumed to exist for this example.

C0035

C0035 is the resulting memory structure for this utterance.

STARTING INFERENCE QUEUE:  
(X 1.0 C0035)

We suppress all but this structure on the starting inference queue.

.....

(We will be seeing about one-fourth of the original trace output for this example)

```
ABOUT TO APPLY @CAUSE1 TO C0035
C0035: (CAUSE (CAUSE (*PROPEL* #JOHN1
C0013 #JOHN1 #BILL1) (*PHYSCONT*
C0013 #BILL1)) (CAUSE (*DO* #MARY1
C0010) (*PHYSCONT* C0021
#JOHN1)))
INFERRING: C0028
```

Here, the CAUSE inference molecule is injecting the two subconceptualizations, A and B in Fig. 1 into the inference stream.

```
ABOUT TO APPLY @CAUSE2 TO C0035
C0035: (CAUSE (CAUSE (*PROPEL* #JOHN1
C0013 #JOHN1 #BILL1) (*PHYSCONT*
C0013 #BILL1)) (CAUSE (*DO* #MARY1
C0010) (*PHYSCONT* C0021 #JOHN1)))
INFERRING: C0034
```

RECORDING CAUSAL RELATION:  
(C0024 . C0032)

The causal structure of this conceptualization indicated that a path should be found relating structure C to structure D in Fig. 1. This is noted. C0024 corresponds to C, C0032 to D.

.....

```
ABOUT TO APPLY @PHYSCONT1 TO C0032
C0032: (*PHYSCONT* C0021 #JOHN1)
INFERRING: (*FEEL* #MARY1 #POSEMOTION
#JOHN1)
ALSO GENERATING: (TIME C0039 C0017)
```

Here, the causative inference that Mary's kissing was probably caused by her feeling a positive emotion toward John is made.

.....

```
ABOUT TO APPLY @PROPEL1 TO C0024
C0024: (*PROPEL* #JOHN1 C0048 #JOHN1
#BILL1)
INFERRING: (*FORCECONT* C0048 #BILL1)
ALSO GENERATING: (TS C0052 C0016)
```

Because the PHYSCONT of John's hand and Bill was caused by a PROPEL, MEMORY here makes the inference that it was a forceful contact.

.....

```
ABOUT TO APPLY FORCECONT2 TO C0052
C0052: (*FORCECONT* C0048 #BILL1)
INFERRING: (NEGCHANGE #BILL1 #PSTATE)
ALSO GENERATING: (TIME C0055 C0016)
```

Since one of the objects involved in the FORCECONT was a person, MEMORY predicts a small NEGCHANGE on his part. The degree of the NEGCHANGE is dependent upon the type of object which came into contact with him.

.....

ABOUT TO APPLY @NEGCHANGE2 TO C0055  
C0055: (NEGCHANGE #BILL1 #PSTATE)  
INFERRING: (POCHANGE #MARY1 #JOY)  
ALSO GENERATING: (TIME C0061 C0016)

.....

ABOUT TO APPLY @POCHANGE1 TO C0061  
C0061: (POCHANGE #MARY1 #JOY)  
INFERRING: (\*MFEEL\* #MARY1 #POSEMOTION  
#JOHN1)  
ALSO GENERATING: (TS C0068.C0016)

CAUSAL EXPANSION ACHIEVED:  
(C0024 C0032)  
CONTACT POINTS ARE: (C0068 C0039)  
MERGING:  
C0068: (\*MFEEL\* #MARY1 #POSEMOTION  
#JOHN1)  
C0039: (\*MFEEL\* #MARY1 #POSEMOTION  
#JOHN1)

.....

\*!EXPANDED\_CAUSALS  
(C0024 . C0032)  
\*(CAUSAL\_PATH @C0024 @C0032)  
(C0024 C0052 C0055 C0061 C0068 C0032)

-----  
C0024: (\*PROPEL\* #JOHN1 C0048 #JOHN1  
#BILL1)  
ASET:  
C0054: (CAUSE # C0052)  
C0028: (CAUSE # C0026)  
C0025: (TIME # C0016)  
REGENCY: 9900  
TRUTH: T, STRENGTH: 1.0  
REASONS:  
C0028: (CAUSE C0024 C0026)  
OFFSPRING:  
C0070: (CAUSE C0061 C0068)  
C0068: (\*MFEEL\* #MARY1 #POSEMOTION  
#JOHN1)  
C0065: (CAUSE C0055 C0063)  
C0063: (\*MFEEL\* #BILL1 #NEGEMOTION  
#JOHN1)  
C0054: (CAUSE C0024 C0052)  
C0053: (TS C0052 C0016)  
C0052: (\*FORCECONT\* C0048 #BILL1)  
ISEEN: (@PROPEL1)

Here, because Mary was feeling a negative emotion toward Bill at the time, when Bill underwent a small NEGCHANGE; the prediction can be made that Mary may have experienced a degree of joy.

Looking back the causal path which lead to Mary's likely change in joy, the POSCHANGE inference molecule discovers that it was an action on John's part which was most directly responsible for her joy. The inference that Mary might have started feeling a positive emotion toward John is made.

As this last inference is made, the inference evaluator notices that the same information exists elsewhere in the memory. This is a point of contact in inference space. It is furthermore noticed that the two MFEEL structures join a causal path between two structures which have been related causally by language. The two MFEEL structures are merged into one, and this event is noted as a causal chain expansion. To the left, C0068 and C0039 are the contact points, C0024 and C0032 are the two structures which have now been causally related.

Inference proceeds, and finally stops. At that point, we took a look at the structures lying along this explained causal path. C0024 is the original PROPEL structure, C0032 is the PHYSCONT-lips structure. The service function CAUSAL\_PATH will track down the causal linkage for us. The causal chain consists of the six structures to the left.

This is the original PROPEL. During the process, but not shown, C0048 was detected as unspecified, and filled in as John's hand. Notice on the REASONS and OFFSPRING sets the results of other inferencing which was not discussed above.

```

-----
C0052: (*FORCECONT* C0048 #BILL1)
ASET:
  C0077: (WANT #JOHN1 #)
  C0057: (CAUSE # C0055)
  C0054: (CAUSE C0024 #)
  C0053: (TS # C0016)
REGENCY: 18416
TRUTH: T, STRENGTH: 0.89999999
REASONS:
  C0024: (*PROPEL* #JOHN1 C0048.#JOHN1
          #BILL1)
OFFSPRING:
  C0078: (TS C0077 C0016)
  C0077: (WANT #JOHN1 C0052)
  C0057: (CAUSE C0052 C0055)
  C0056: (TIME C0055 C0016)
  C0055: (NEGCHANGE #BILL1 #PSTATE)
ISEEN: (@FORCECONT2)

```

Here is the FORCECONT which was inferred from the PROPEL.

```

-----
C0055: (NEGCHANGE #BILL1 #PSTATE)
ASET:
  C0079: (WANT #JOHN1 #)
  C0067: (CAUSE # C0059)
  C0066: (CAUSE # C0061)
  C0065: (CAUSE # C0063)
  C0057: (CAUSE C0052 #)
  C0056: (TIME # C0016)
REGENCY: 19833
TRUTH: T, STRENGTH: 0.85500000
REASONS:
  C0052: (*FORCECONT* C0048 #BILL1)
  I0002: (ISA #BILL1 #PERSON)
OFFSPRING:
  C0080: (TIME C0079 C0016)
  C0079: (WANT #JOHN1 C0055)
  C0067: (CAUSE C0055 C0059)
  C0066: (CAUSE C0055 C0061)
  C0065: (CAUSE C0055 C0063)
  C0064: (TS C0063 C0016)
  C0063: (*MFEEL* #BILL1 #NEGEMOTION
          #JOHN1)
  C0062: (TIME C0061 C0016)
  C0061: (POSCHANGE #MARY1 #JOY)
  C0060: (TS C0059 C0016)
  C0059: (WANT #BILL1 C0058)
ISEEN: (@NEGCHANGE3 @NEGCHANGE2
        @NEGCHANGE1)

```

This is Bill's likely (small) change in PSTATE which resulted from the FORCECONT.

```

-----
C0061: (POSCHANGE #MARY1 #JOY)
ASET:
  C0070: (CAUSE # C0068)
  C0066: (CAUSE C0055 #)
  C0062: (TIME # C0016)
REGENCY: 24616
TRUTH: T, STRENGTH: NIL
REASONS:
  C0055: (NEGCHANGE #BILL1 #PSTATE)
  I0137: (MFEEL #MARY1 #NEGEMOTION
          #BILL1)
OFFSPRING:

```

This is the important inference that Bill's NEGCHANGE may have cause a small degree of happiness in Mary. Notice that one of the REASONS was assumed to be the case beforehand. (I0137).

C0070: (CAUSE C0061 C0068)  
C0069: (TS C0068 C0016)  
C0068: (\*MFEEL\* #MARY1 #POSEMOTION  
#JOHN1)  
ISEEN: (@POSCHANGE1)

-----  
C0068: (\*MFEEL\* #MARY1 #POSEMOTION  
#JOHN1)  
ASET:  
C0085: (WANT #JOHN1 #)  
C0040: (TIME # C0017)  
C0044: (\*MLOC\* # C0041)  
C0047: (CAUSE # C0032)  
C0070: (CAUSE C0061 #)  
C0069: (TS # C0016)  
REGENCY: 27366  
TRUTH: T, STRENGTH: 0.95000000  
REASONS:  
C0061: (POSCHANGE #MARY1 #JOY)  
C0024: (\*PROPEL\* #JOHN1 C0048 #JOHN1  
#BILL1)  
C0044: (\*MLOC\* C0068 C0041)  
OFFSPRING:  
C0087: (TS C0085 C0016)  
C0086: (TIME C0085 C0017)  
C0085: (WANT #JOHN1 C0068)  
ISEEN: NIL

-----  
C0032: (\*PHYSCONT\* C0021 #JOHN1)  
ASET:  
C0088: (WANT #JOHN1 #)  
C0071: (WANT #MARY1 #)  
C0047: (CAUSE C0068 #)  
C0046: (CAUSE # C0044)  
C0034: (CAUSE C0029 #)  
C0033: (TIME # C0017)  
REGENCY: 12016  
TRUTH: T, STRENGTH: 1.0  
REASONS:  
C0034: (CAUSE C0029 C0032)  
OFFSPRING:  
C0089: (TIME C0088 C0017)  
C0088: (WANT #JOHN1 C0032)  
C0072: (TIME C0071 C0017)  
C0071: (WANT #MARY1 C0032)  
C0047: (CAUSE C0068 C0032)  
C0046: (CAUSE C0032 C0044)  
C0045: (TS C0044 C0017)  
C0044: (\*MLOC\* C0068 C0041)  
C0040: (TIME C0068 C0017)  
ISEEN: (@PHYSCONT2 @PHYSCONT1)

Here, Mary is feeling a positive emotion toward John, whose action indirectly caused her joy. This structure is the point of contact, and is the structure which resulted from the merge. Notice that its STRENGTH has assumed the higher STRENGTH of the two structures which were merged.

This is the original PHYSCONT-lips structure which lead, via a causative inference to the prediction that Mary may have felt a positive emotion toward John.

This WANT is a prediction that one reason Mary may have kissed John is so that he would know she felt a positive emotion toward him.  
This MLOC represents the inference that John probably now knows that Mary MFEELS a positive emotion toward him.



## APPENDIX B

### INFERENCE-REFERENCE RELAXATION CYCLE, COMPUTER EXAMPLE

This computer example illustrates reference-inference, reference-inference interaction (two inference passes). Hearing the input "Bill saw John kiss Jenny.", MEMORY is unable to decide upon the referent of "Jenny": it could be Jenny Jones or Jenny Smith. MEMORY therefore creates a temporary token having as features all the common features of Jenny Jones and Jenny Smith. By inference, MEMORY is able to decide upon Jenny Jones. At that point, the temporary token is merged into the concept for Jenny Jones, and a second pass of inferencing is initiated. However, on the second pass a new inference arises: because Bill loves Jenny Jones, and he saw John kiss her, he (probably) became angry at John. This inference was not triggered on the first inference pass because being loved by Bill was not a common feature of both Jennys, and hence not accessible then (ie. it had not been copied to the temporary token's occurrence set).

The example begins with a few lines to set the scene for MEMORY. Inferencing on these setup lines (which is normally spontaneous) has been suppressed for the sake of simplicity in this example.

```
-----  
JOHN WAS IN PALO ALTO YESTERDAY  
((*LOC* (#JOHN1) (#PALOALTO))  
(TIME _ (C0001)))  
C0002
```

```
-----  
JENNY JONES WAS IN PALO ALTO YESTERDAY  
((*LOC* (#JENNY2) (#PALOALTO))  
(TIME _ (C0004)))  
C0005  
-----
```

This example illustrates reference-inference, reference-inference interaction. That is, MEMORY is unable to establish a reference, so it creates a temporary token, and proceeds with inference. Inferencing generates new information which solves the reference, so more inferencing can be undertaken. However, because features of the referent are accessible on the second inference pass, new inferences are possible.

To the left, MEMORY is reading in some

JENNY SMITH WAS IN FRANCE YESTERDAY.  
((\*LOC\* (#JENNY1) (#FRANCE))  
(TIME - (C0007)))  
C0008

-----  
BILL LOVES JENNY JONES  
((\*MFEEL\* (#BILL1) (#LOVE) (#JENNY2)))  
C0010

-----  
BILL SAW JOHN KISS JENNY YESTERDAY.

COPYING COMMON FEATURES TO C0015  
FROM (#JENNY2 #JENNY1)

(((\*MTRANS\* (#BILL1) ((CAUSE ((\*DO\*  
(#JOHN1) (#UNSPECIFIED)) (TIME  
(C0011))) ((\*PHYSCONT\* (C0012) (C0015))  
(TIME - (C0011)))))) (C0018) (C0021))  
(TIME - (C0011)) (INST - ((\*LOOK\_AT\*  
(#BILL1) (C0015 #JOHN1)) (TIME -  
(C0011))))))

C0031

-----  
#JENNY1: NIL

ASET:  
I0019: (SURNAME # SMITH)  
I0018: (ISA # #PERSON)  
I0017: (NAME # JENNY)  
REGENCY: NIL

-----  
#JENNY2: NIL

ASET:  
I0022: (SURNAME # JONES)  
I0021: (ISA # #PERSON)  
I0020: (NAME # JENNY)  
REGENCY: NIL

-----  
C0015: NIL

ASET:  
C0029: (\*LOOK\_AT\* #BILL1 #)  
C0026: (\*PHYSCONT\* C0012 #)  
C0017: (ISA # #PERSON)  
C0016: (NAME # JENNY)  
REGENCY: 9866

-----  
STARTING INFERENCE QUEUE:  
(X 1.0 C0031) (X 1.0 C0017)  
(X 1.0 C0016).

.....

information which is relevant to this demonstration. Each of these inputs would normally produce inferences as it is processed, but inferencing has been suppressed for the first four sentences of this example. The four sentences are shown with their partial integrations and final structures: C0002, C0005, C0008, C0010.

The synopsis of this short plot is as follows: There are two Jennys: Jenny Jones and Jenny Smith. Bill loves Jenny Jones, John and Jenny Jones were in Palo Alto yesterday, Jenny Smith was in France yesterday. The climax comes when Bill sees John kiss Jenny. It is MEMORY's job to figure out which Jenny. MEMORY will decide upon Jenny Jones, then re-inference and infer that Bill probably got angry at John-- something which wouldn't have happened if Bill had seen John kiss Jenny Smith.

To the left, the climax line is in the process of being read and internalized. Its final structure is C0031. Notice that C0015 was created to stand for some Jenny, and that all common features of the two Jenny candidates were copied to it.

-----  
We interrupt MEMORY at this point to have a look at the two Jennys and C0015, the token representing one of these Jennys.

-----  
This is the person named Jenny who Bill saw yesterday, and who John kissed. C0012 is the token representing John's lips, which were in \*PHYSCONT\* with this person named Jenny (C0015) at time C0011.

-----  
MEMORY begins inferencing from this input. The starting inference queue consists of the main structure for the sentence, together with all other facts known about C0015. In 378's case, these are simply that C0015 is

APPLYING INF MOLECULE \*MLOC\* TO  
C0037: (\*MLOC\* (CAUSE (\*DO\* #JOHN1  
#UNSPECIFIED) (\*PHYSCONT\* C0012  
C0015)) C0021)

ABOUT TO APPLY @MLOC1 TO C0037  
INFERRING: (\*MLOC\* C0028 C0040)  
ALSO GENERATING: (TS C0043 C0011)

.....

APPLYING INF MOLECULE \*PHYSCONT\* TO  
C0026: (\*PHYSCONT\* C0012 C0015)

ABOUT TO APPLY @PHYSCONT1 TO C0026  
INFERRING: (\*MFEEL\* #JOHN1 #POSEMOTION-  
C0015)  
ALSO GENERATING: (TIME C0049 C0011)

ABOUT TO APPLY @PHYSCONT2 TO C0026  
INFERRING: (\*MLOC\* C0049 C0051)  
ALSO GENERATING: (TS C0054 C0011)

ABOUT TO APPLY @PHYSCONT3 TO C0026  
INFERRING: (\*LOC\* C0015 #PALOALTO)  
ALSO GENERATING: (TIME C0056 C0011)

.....

APPLYING @POSTSCAN TO C0043: . . .  
(\*MLOC\* (CAUSE (\*DO\* #JOHN1  
#UNSPECIFIED) (\*PHYSCONT\* C0012 C0015))  
C0040)

.....

INFERRING: (\*MLOC\* C0049 C0040)  
COPYING TIMES FROM C0043 TO C0086

.....

-----  
C0015: NIL

ASET:  
C0056: :(\*LOC\* # #PALOALTO)  
C0053: (PART C0051 #)  
C0049: (\*MFEEL\* #JOHN1 #POSEMOTION #)  
C0029: (\*LOOK\_AT\* #BILL1 #)  
C0026: (\*PHYSCONT\* C0012 #)  
C0017: (ISA # #PERSON)

a person, and that its name is Jenny. These will not be of use in this example. All other subpropositions have been suppressed from the starting inference queue for this example.

One inference from Bill's seeing this event is that he knows that the event occurred. That is, the event went from his eyes to his conscious processor, C0021. To the left, the inference that Bill knows about John's kissing Jenny is being generated; information in Bill's CP (C0021) will also enter his LTM, C0040. This fact will be of use during the second pass of inferencing (after MEMORY decides that C0015 is Jenny Jones).

Another inference arises from John's lips being in PHYSCONT with C0015: that John feels a positive emotion toward C0015. The structure representing this inference is C0049.

Another inference from John's kissing action is that C0015 knows that John feels a positive emotion toward C0015, C0051 is C0015's LTM. This inference will be of no direct consequence in this example.

MEMORY also infers from John's kissing C0015 that John and C0015 had the same location at the event time, C0011 (yesterday). Since MEMORY knows that John was in Palo Alto, and has no information concerning C0015's location yesterday, MEMORY infers that C0015 was also in Palo Alto yesterday. This information will solve the reference ambiguity.

During the postscan inferencing, the fact that Bill saw John kiss C0015 leads to the inference that Bill knows that John feels a positive emotion toward C0015. This inference type implements the principle that if a person knows X, he also is likely also to know the inferences which can be drawn from X. That is, MEMORY assumes that other people possess the same inference powers as MEMORY does.

-----  
Inferencing eventually ceases. We interrupt processing at this point to examine C0015, the unknown Jenny. Notice the new information which has been built up about C0015.

C0051 is C0015's LTM.

C0012 is John's lips.

C0016: (NAME # JENNY)  
RECENCY: 9350

-----  
C0056: (\*LOC\* C0015 #PALOALTO)

ASET:

C0078: (\*MLOC\* # C0040)

C0057: (TIME # C0011)

RECENCY: 42533

TRUTH: T, STRENGTH: 0.90250000

REASONS:

C0002: (\*LOC\* #JOHN1 #PALOALTO)

C0026: (\*PHYSCONT\* C0012 C0015)

OFFSPRING:

C0101: (\*MLOC\* C0024 C0051)

ISEEN: NIL

-----  
C0086: (\*MLOC\* C0049 C0040)

(\*MLOC\* (\*MFEEL\* #JOHN1 #POSEMOTION  
C0015) C0040)

ASET:

C0087: (TS # C0011)

RECENCY: 25750

TRUTH: T, STRENGTH: 0.95000000

REASONS:

C0043: (\*MLOC\* C0028 C0040)

ISEEN: (\*MLOC2)

-----  
RETRYING REFERENCE:

(C0015 #JENNY2 #JENNY1)

REFERENCE AMBIGUITY SOLVED.

OLD: (C0015 #JENNY2 #JENNY1)

NEW: #JENNY2

MERGING:

#JENNY2: #JENNY2

C0015: C0015

PURGING (\*LOC\* C0015 #PALOALTO)

PURGING: (\*MLOC\* (\*LOC\* C0015 #PALOALTO)  
C0040)

PURGING: (TS (\*MLOC\* (\*LOC\* C0015  
#PALOALTO) C0040) C0011)

PURGING: (TIME (\*LOC\* C0015 #PALOALTO)  
C0011)

PURGING: (ISA C0015 #PERSON)

PURGING: (NAME C0015 JENNY)

-----  
#JENNY2: NIL

ASET:

C0117: (IDENTIFIES # C0015)

C0026: (\*PHYSCONT\* C0012 #)

C0029: (\*LOOK\_AT\* #BILL1 #)

C0049: (\*MFEEL\* #JOHN1 #POSEMOTION #)

C0053: (PART C0051 #)

Since it will settle the reference ambiguity, we have a closer look at the structure which represents C0015's being in Palo Alto yesterday (C0011). C0078 represents Bill's knowledge of C0015's location yesterday (but has no direct relevance to this example).

Notice that the reasons for MEMORY believing that C0015 was in Palo Alto at time C0011 are twofold: that John was in Palo Alto at that time, and that a body part of John was in PHYSCONT with C0015 then.

We also examine the structure which represents the inference that Bill knows that John feels a positive emotion toward C0015. This information will come into play after C0015's identity is solved (on the second inference pass). C0087 indicates when Bill started knowing this fact (C0040 is his LTM).

-----  
The first pass of inferencing is now finished. We allow MEMORY to proceed. It notices that a reference decision is pending, and attempts to decide between #JENNY1 and #JENNY2 as the referent of C0015 by using newly-inferred information about C0015 (from the first pass). It succeeds, because #JENNY2 was known to be in Palo Alto yesterday, and this matches new C0015 information, C0056.

MEMORY merges C0015 into #JENNY2, purging old information which is not used to augment #JENNY2. Recall that the merge replaces occurrence set pointers, so that every MEMORY structure which referenced C0015 now references #JENNY2.

We have another look at #JENNY2 before the second inference pass begins.

C0010: (\*MFEEL\* #BILL1 #LOVE #)  
C0005: (\*LOC\* # #PALOALTO)  
I0019: (SURNAME # JONES)  
I0018: (ISA # #PERSON)  
I0017: (NAME # JENNY)  
REGENCY: 8950

RE-INFERRING...

.....

APPLYING INF MOLECULE \*MLOC\* TO  
C0086: (\*MLOC\* (\*MFEEL\* #JOHN1  
#POSEMOTION #JENNY2) C0040)

ABOUT TO APPLY @MLOC3 TO C0086  
INFERRING: (\*MFEEL\* #BILL1 #ANGER  
#JOHN1)  
ALSO GENERATING: (TS C0119 C0011)

.....

-----  
C0119: (\*MFEEL\* #BILL1 #ANGER #JOHN1)

ASET:  
C0121: (CAUSE C0086 #)  
C0120: (TS # C0011)  
REGENCY: 107600  
TRUTH: T, STRENGTH: 0.90250000  
REASONS:  
C0086: (\*MLOC\* C0049 C0040)  
C0010: (\*MFEEL\* #BILL1 #LOVE #JENNY2)  
ISEEN: NIL

-----  
MEMORY begins the second pass of inferencing. This consists of subjecting each inference which arose from the first pass to inference again. The ISEEN property prevents duplication of inferences during second and subsequent passes.

One new inference which was not possible on the first pass is that Bill probably became angry at John. This inference arises from Bill's knowing that John feels a positive emotion toward #JENNY2, someone Bill loves. C0119 is the structure representing Bill's incipient anger toward John. The crucial point is that this inference became possible only after #JENNY2's features became available after a reference decision, which was in turn made possible through first-pass inferencing.

-----  
Finally, we have a look at this second pass inference.

C0121 represents the cause of Bill's anger as being C0086, his knowing about the kissing event. C0049.

Notice the reasons MEMORY believes that Bill became angry at John: he knew John kissed #JENNY2 (this structure is C0049), and he loves #JENNY2.

## REFERENCES

- (A1) Abelson, R. P., "The Structure of Belief Systems," in Schank and Colby (eds.) Computer Models of Thought and Language, W. H. Freeman, San Francisco, 1973
- (A2) Anderson, J., and Bower, G., Human Associative Memory, Wiley & Sons, New York, 1973
- (C1) Charniak, E., "Toward a Model of Children's Story Comprehension," Doctoral Dissertation, MIT, 1972
- (E1) Eccles, J., "The Physiology of Imagination," *Scientific American*, Sept. 1958
- (G1) Goldman, N., "Computer Generation of Natural Language from a Deep Conceptual Base," Doctoral Dissertation, Stanford, 1974
- (H1) Hewitt, C., "Procedural Embedding of Knowledge in PLANNER," Proc. 2nd Joint Conf. on Artificial Intelligence, 1971
- (L1) Levy, D., "Contextual Guidance of Inference," working paper, Stanford, 1974
- (M1) Marcus, M., "Wait-and-See Strategies for Parsing Natural Language," working paper, M.I.T., AI Lab, August 1974
- (Q1) Quillian, R., "Semantic Memory," in Minsky (ed.) Semantic Information Processing, MIT Press, Cambridge Mass., 1968
- (R1) Rieger, C., "Conceptual Memory: A Theory and Computer Program for Processing the Meaning Content of Natural Language Utterances," Doctoral Dissertation, Stanford, 1974
- (R2) Riesbeck, C., "Computational Understanding: Analysis of Sentences and Context," Doctoral Dissertation, Stanford, 1974
- (S1) Schank, R., "Conceptual Dependency: A Theory of Natural Language Understanding," *Cognitive Psychology*, 3 (4), 1972
- (S2) Schank, R., Goldman, N., Rieger, C., and Riesbeck, C., "Primitive Concepts Underlying Verbs of Thought," Stanford AI Memo #162, 1972

- (S3) Schank, R., Goldman, N., Rieger, C., and Riesbeck, C., "MARGIE: Memory, Analysis, Response Generation and Inference on English," in Proc. 3rd Joint Conf. of Artificial Intelligence, 1973
- (W1) Winograd, T., "Procedures as a Representation for Data in a Computer Program for Understanding Natural Language," Doctoral Dissertation, MIT, 1971
- (Z1) Zadeh, L., "The Concept of a Linguistic Variable and Its Application to Approximate Reasoning," working paper, U.C. Berkeley, 1973

# END

