

MAXIMAL CONSISTENT INTERPRETATIONS OF ERRORFUL DATA
IN HIERARCHICALLY MODELLED DOMAINS

Mark S. Fox and David Jack Mostow
Computer Science Department 1
Carnegie-Mellon University
Pittsburgh, Pa. 15213

ABSTRACT

A method is presented for constructing maximal consistent interpretations of errorful data. The method appears applicable to many tasks (speech understanding, natural language understanding, vision, medical diagnosis) requiring partial-matching of errorful data against complex, hierarchically defined patterns. The data is represented as symbolic structures (word sequences, line segment configurations, disease symptoms). Errors consist of missing data (unrecognized words, occluded lines, undetected symptoms) and extra (possibly inconsistent) data (incorrectly recognized words, visual noise, spurious symptoms). Data interpretations correspond to substructures of a hierarchy of predefined concepts. Constraints on consistent conceptual structures are embedded in the hierarchy. An implementation of the method has correctly interpreted errorful sets of sentence fragments recognized by the HEARSAY-II speech understanding system. The implementation has also correctly interpreted typed-in ungrammatical sentences. Detailed examples illustrate operation of the method on real data.

1. INTRODUCTION

The application of AI methods to complex domains (e.g., speech, vision, medical diagnosis) has expanded the dimensions of data interpretation to incorporate some novel phenomena. Two of these phenomena are data error and hierarchically defined data patterns.

Many complex domains are characterized by errorful data. Errors such as insertion, deletion, substitution, and repetition of information increase as the uncertainty of source data transduction and interpretation increases. Data may be mutually inconsistent in that two or more pieces of information cannot be explained consistently. Tolerating error and inconsistencies in the data requires robust methods that can not only find the best interpretation but are able to distinguish the inconsistent and errorful data from the consistent data.

Another aspect of data interpretation in complex domains is that interpretations represent complex, hierarchically defined concepts (ideas, rules, patterns) rather than simple, independent concepts (features). Often the concepts used in interpretations can be placed in a hierarchy where each concept is defined in terms of its subconcepts. This structure of concepts is called a conceptual hierarchy. A collection of data can then be interpreted by the highest concept in the hierarchy supported (validated) by the data. The interpretation of the data is defined by the concept's descendants (subconcepts, subsubconcepts, etc.) and the data which supports them. These descendants form a substructure of the conceptual hierarchy.

The general data interpretation problem can now be restated as a search for the concept in the conceptual hierarchy that explains (is supported by) the most data. The data supporting the structure underlying this maximal concept can be described as the maximal consistent subset of data.

In this paper we define conceptual

This work was supported in part by the Defense Advanced Research Projects Agency under contract no. F44620-73-C-0074 and monitored by the Air Force Office of Scientific Research. In addition, the first author was partially supported by a National Research Council of Canada Postgraduate Scholarship and the second author was partially supported by a National Science Foundation Graduate Fellowship.

hierarchies and maximal consistent interpretations. We then describe a method for interpreting data in such an environment, i.e., finding maximal consistent interpretations in a conceptual hierarchy. Examples illustrating the method are shown. Finally, we show the actual application of the method to the problem of interpreting errorful sentence fragments recognized by the HEARSAY-II speech understanding system (Erman, 1975).

2. A REAL EXAMPLE

The matching problem used as an example throughout this paper is taken from the HEARSAY-II speech understanding system. When HEARSAY-II is unable to completely recognize a spoken sentence (utterance), it generates a set of sentence fragments (Hayes-Roth et al, 1976c) which must be interpreted by the semantic interpretation module, named SEMANT. The generated fragments can be both errorful and mutually inconsistent (Example 2.1). A sentence fragment is a chunk of consistent data in that it consists of a grammatically plausible sequence of recognized words. HEARSAY-II mechanisms effective in identifying such chunks are not suited to combining them into an overall consistent interpretation of the utterance.

EXAMPLE 2.1

```
1: <0> [ WHAT HAS HERBERT <75>
2: <18> PAPER ABOUT PATTERN MATCHING ]<177>
3: <29> IN LEARNING OR PATTERN MATCHING ] <177>
4: <0> [ WHO <24>
```

Correct Sentence:

```
<0>[ WHO HAS WRITTEN ABOUT PATTERN MATCHING ]<177>
```

Example 2.1 shows four sentence fragments generated when HEARSAY-II was unable to recognize the sentence [WHO HAS WRITTEN ABOUT PATTERN MATCHING]. The square brackets denote the start and finish of the spoken utterance. The numbers enclosed in angle brackets specify, in centiseconds, how long after the start of the utterance each fragment begins and ends. Fragment 4 correctly matches the initial portion of the spoken sentence. Fragments 1-3 contain substitution errors. Fragments 1 and 2 are mutually inconsistent in that they provide different interpretations of the overlapping time period <18:75>. The fragment pairs 1 & 3, 1 & 4, and 2 & 3 are inconsistent for the same reason. Also, Fragment 1 specifies a WHAT question whereas fragment 4 specifies a WHO question. Thus Fragments 1 and 4 are semantically inconsistent, irregardless of their times. Each fragment is semantically described by a hierarchically structured collection of concepts. Figure 2.1 shows a portion of the conceptual hierarchy used by the SEMANT module in HEARSAY-II. Figure 2.2 shows the hierarchical description of the correct sentence.

The problem of interpreting these fragments illustrates the phenomena of data error and hierarchically-structured interpretations. The method used for solving this problem appears applicable to a significant class of problems exhibiting these two phenomena.

3. CONCEPTUAL HIERARCHIES

A conceptual hierarchy can be represented by a directed graph of concepts. This graph is tree structured in that it has a root at the top and leaf nodes at the bottom; however, cycles are permitted. The sons of a node define the subconcepts that compose the father. The root of the graph defines the highest level (most general) interpretation of all the concepts beneath it.

A given interpretation task has a set of

prespecified patterns, modelling possible data-generating events (e.g., utterances, scenes, diseases). Each pattern has its own underlying hierarchy. These hierarchies are collapsed into a single hierarchy (for computational purposes) by adding a new root concept. The sons of the root concept are called the primary concepts of the hierarchy. The primary concepts are the roots of the original disjoint hierarchies. A collection of data is interpreted by choosing one or more primary concepts matched by the data.

Whether a primary concept can be considered to be matched by the data depends on which (if any) of its subconcepts have been matched and on the relationship between the primary concept and its subconcepts. A concept that requires all its subconcepts (sons) to be matched as a necessary condition for itself to be matched is a conjunctive concept. The subconcepts are related to the conjunctive concept by the constituent (IS-PART-OF) relationship. A concept requiring any non-empty subset of its subconcepts to be matched is a union concept. The subconcepts are related to the union concept by the optional constituent (IS-OPTIONAL-PART-OF) relationship. A concept which requires one of its subconcepts to be matched is a disjunctive concept. The subconcepts are related to the disjunctive concept by the taxonomic (IS-A) relationship. Other, more complex relationships can be defined on the subconcepts of a concept by defining parameterized constraints on the data supporting the subconcepts. Figure 2.1 shows part of the conceptual hierarchy used to describe the types of sentences expected by the HEARSAY-II speech understanding system. The method used to match sentence fragments to concepts is parameterless but successful nonetheless. This method is described in Section 6.

4. MAXIMAL CONSISTENT INTERPRETATIONS

Based on the above definitions, a maximal consistent interpretation of the data is defined as the primary concept and the subtree underlying it that is matched by a maximal consistent subset of the data. A maximal consistent subset contains the greatest amount of domain information (measured by some function) that is mutually consistent. If the subconcepts of a concept are mutually consistent, it follows that the domain information that supports (matches) these subconcepts is mutually consistent, with an important qualification. This qualification is necessary because two or more competing (mutually inconsistent) pieces of domain data may support the same concept. A consistent interpretation must choose only one of these pieces of data to support the concept. Since there is a choice of which data to incorporate in the interpretation, there are many possible interpretations derivable from the data supporting the subtree. These interpretations can be ordered by the function that measures the quantity of data incorporated in (explained by) an interpretation. Thus once concept matching has been carried out, any subtree within the conceptual hierarchy defines consistent data sets. Furthermore, the distinction between conjunctive, union, and disjunctive concepts allows us to identify which information is missing in a particular interpretation. Missing information corresponds to unsupported sons of partially supported conjunctive concepts.

Data consistency must be defined relative to a particular application. A set of data is considered consistent if it satisfies some set of application-specific constraints. Some of these constraints can be incorporated in the structure of the conceptual hierarchy; a given hierarchy implicitly defines a class of permissible data combinations. For example, the data configurations supporting the sons of a node are mutually consistent if the node is conjunctive but not if it is disjunctive. Other constraints can be incorporated in the chunking process which generates configurations of data. Data chunks incorporate information about data consistency insofar as the data in a chunk is mutually consistent, and subconfigurations of the chunk are consistent with the chunk itself. This information is incomplete in that it doesn't specify whether different chunks are mutually consistent. Finally,

constraints not incorporated in the hierarchy or the chunking process must be satisfied by special tests on appropriate properties of the data supporting a potential interpretation. One such constraint in the HEARSAY-II example is temporal consistency between the various data fragments supporting the interpretation of an utterance. Fragments which assign different transcriptions to the same time interval are mutually inconsistent.

The measure of an interpretation must also be defined relative to a particular application. The measuring function should reflect the differing credibility of alternative interpretations. Several factors affect this credibility. One of them is the amount of data satisfactorily explained by a given interpretation. An interpretation which accounts for a large subset of the data may be more credible than an interpretation which accounts for only a small subset. Another factor affecting the credibility of an interpretation is the cogency of its conceptual structure. For example, an interpretation with many missing pieces (unsupported sons of conjunctive nodes) may be less credible than an interpretation with no missing pieces. An extensive interpretation (supported by many concepts) may be more credible than a limited interpretation (involving very few concepts). A third factor is the individual credibility of the data chunks supporting the interpretation. The number of consistency constraints satisfied by a chunk increases with its size. If these constraints are reasonably rigorous, larger chunks may be more credible than smaller chunks. Thus the credibility of a particular datum may be sensitive to the context (chunk) in which it occurs. The more accurately these various credibility factors are represented in the function which rates alternative interpretations, the more often the maximal (highest-rated) consistent interpretation will in fact be correct.

5. MATCHING AND INTERPRETING

As previously described, the conceptual hierarchy is a graph whose nodes are concepts. We allow each node to act as a repository of information during the matching and interpretation. Finding the maximal consistent interpretation is a three part process.

The first phase matches the concepts against the data. We assume that the initial part of this process is performed by some mechanism which structures the data by identifying chunks, i.e., local configurations of mutually consistent data. In the current example, this mechanism is HEARSAY-II, and the chunks are sentence fragments. This match may be full or partial in that conjunctive concepts may be completely or partially matched. When a concept is successfully matched, the domain information matching the concept is stored at the corresponding node. This information is said to directly support the concept.

The second phase integrates the chunks by finding concepts which explain combinations of chunks. This is accomplished by "notching" each matched concept and all its ancestors, i.e., increasing their credibility scores according to the amount of data supporting the matched concept. The notching process assigns a metric of how well each concept is supported by data. Various metrics are possible. The metric used in this paper is defined as follows. The score of a disjunctive concept is the size of the largest chunk directly supporting the concept plus the score of the concept's highest-rated son. The score of a conjunctive or union concept is the sum of the sizes of the chunks supporting it directly, plus the sum of the scores of its sons. The score of an unsupported concept is zero. Scoring is computed by a one-pass notching process which propagates scores bottom-up starting at the leaves of the hierarchy. The notching process can be viewed as a flow of support from the data through the conceptual hierarchy.

The last phase selects a consistent interpretation by walking top-down through the hierarchy starting at the root concept, and incorporating the visited concepts into the interpretation. The maximally supported subtree in the hierarchy is found by interrogating the score of each concept. When the walk encounters a disjunctive concept, only its highest-scored son is incorporated in the interpretation and

subsequently visited, since the sons of a disjunctive concept are mutually exclusive. When a conjunctive or union concept is encountered, all of its sons with non-zero scores are included in the interpretation and subsequently visited, since they are mutually consistent. Unsupported (zero-scored) sons of conjunctive concepts in the interpretation identify missing data. (More complex relationships between concepts and their sons would allow more complex deductions.) The subtree produced in this fashion represents an interpretation supported by a maximal consistent subset of the data. This consistent subset can be readily identified since the subtree points to the data that supports it.

The nature of the interpretation generated depends on whether the root of the conceptual hierarchy is disjunctive or union. If the root is disjunctive, only one of the primary concepts is incorporated in the interpretation. This property is useful when the purpose of matching is to classify an event according to which single concept best models it. If the root is union, the interpretation can integrate multiple primary concepts in order to explain the data. This capability is useful in domains such as medical diagnosis where the primary concepts model different events (diseases) which can occur simultaneously.

6. A DETAILED EXAMPLE

SEMANT's matching domain is composed of errorful, sometimes mutually inconsistent sentence fragments (chunks). A portion of the conceptual hierarchy is shown in Figure 2.1.

The initial process of matching the domain (fragments) to the concepts, i.e., interpreting individual chunks, is done by parsing. A parser called PPARSE (Erman, 1977), taken from the HEARSAY-II syntax and semantics module (Hayes-Roth, Mostow, and Fox, 1977; Hayes-Roth et al, 1976a), is used to parse each sentence fragment. PPARSE generalizes existing parsing techniques to parse connected subsequences of sentences generated by the grammar. Such a sequence may cross the boundaries of the grammatical hierarchy in that it may not be grammatically derivable from any single non-terminal. PPARSE produces all derivation trees for each fragment. (An ambiguous fragment has more than one derivation tree.) The grammar used by PPARSE is a semantic grammar (Hayes-Roth, Mostow, and Fox, 1977) in which some of the non-terminals have associated semantic meanings. These non-terminals, called semantic nodes, correspond to matched concepts in the hierarchy. In the present grammar, a semantic node has the same name as the corresponding concept. Thus the derivation tree for a sentence fragment points directly to the concepts it matches.

The matching process can be described as follows:

1) The data is chunked by HEARSAY-II into possibly overlapping sentence fragments.

2) The process of single-chunk interpretation determines how each chunk fits into the conceptual hierarchy:

2a) Each fragment is parsed by PPARSE.

2b) For each semantic node in the parse, the corresponding concept in the hierarchy is found, and a pointer to the semantic node is placed at the concept. Thus one can retrieve the word sequence(s) supporting any given concept.

3) The concept and its ancestors are notched by the number of words underlying the semantic node in the parse of the fragment. The details of the notching metric have already been discussed (Section 5).

Figure 6.1 shows the parse trees for the fragments "PAPERS ABOUT PATTERN MATCHING" and "ARTIFICIAL INTELLIGENCE". Nonterminals are distinguished by the "\$" prefix. Only the circled nodes are matched into the conceptual hierarchy: \$TOPIC because it is semantically meaningful and \$MENTION!TOPICS because it is the root node of a parse. When the root node of a derivation tree is not a semantic node, it matches the concept(s) corresponding to its nearest semantically meaningful ancestor(s) in the grammar. In this example, the nearest such ancestors of the root node \$MENTION!TOPICS are \$QUERY!TOPIC and \$QUERY!TOPIC!AUTHOR.

Figure 6.2 shows the matching of the TOPIC concept by the \$TOPIC node. (Note that concept names are not prefixed by "\$"). The score of the \$TOPIC node is 2 because the sub-fragment "PATTERN MATCHING" underlying it is two words long. This score contributes to the scores of all the ancestors of TOPIC in the conceptual hierarchy.

Figure 6.3 shows the matching of the QUERY!TOPIC and QUERY!TOPIC!AUTHOR concepts by the \$MENTION!TOPICS node. The score for this node is 4, since \$MENTION!TOPICS is supported by the 4-word sequence PAPERS ABOUT PATTERN MATCHING. All concepts supported by the \$MENTION!TOPICS node are accordingly notched by 4. Figure 6.4 shows the matching of the TOPIC concept by the \$TOPIC node supported by the fragment "ARTIFICIAL INTELLIGENCE." The TOPIC concept and all its ancestors are notched by 2. Note that while the fragment parse trees contain more than one \$TOPIC node, the conceptual hierarchy contains a single canonical TOPIC node.

The construction of the maximal consistent interpretation starts at the root of the hierarchy. At a disjunctive concept SEMANT chooses the highest-scored son to be in the interpretation. In Figure 6.4 the highest-scored primary concept is REQUEST and is therefore chosen instead of PRUNE. SEMANT next looks at the sons of the REQUEST concept, which is also disjunctive. The QUERY concept is chosen since it is the highest-scored son of REQUEST. The highest-scored sons of QUERY are QUERY!TOPIC and QUERY!TOPIC!AUTHOR (both are supported by \$MENTION!TOPICS). Either one can be chosen to be part of the interpretation. When SEMANT reaches the TOPIC concept, it must choose which supporting data to incorporate in the interpretation. Since SEMANT has traversed a concept supported by the node \$MENTION!TOPICS in order to reach the TOPIC concept, the choice of topic is carried out in the context of the \$MENTION!TOPICS. Hence PATTERN MATCHING is chosen since it is part of that context (i.e., is part of the fragment supporting \$MENTION!TOPICS). Figure 6.5 depicts a resulting interpretation and its corresponding support. (Note that choosing the concept QUERY!TOPIC!AUTHOR instead of QUERY!TOPIC yields an equally well-supported interpretation.)

Figure 6.6 shows the matching of fragments generated when HEARSAY-II was unable to recognize the sentence "LET'S RESTRICT OUR ATTENTION TO PAPERS SINCE NINETEEN SEVENTY FOUR". The fragments "TO PAPERS SINCE NINETEEN SEVENTY FOUR" and "LET'S RESTRICT OUR ATTENTION TO" are mutually consistent while the fragment "DESIGN IN THE ARTS" is not consistent with either of the other two fragments according to the structure of the conceptual hierarchy. The figure shows the state of the conceptual hierarchy after notching has taken place. The maximal consistent interpretation generated is shown in Figure 6.7. This interpretation is the same as that of the correct sentence. Thus the inconsistent information is ignored and the two consistent fragments are semantically combined to form a maximal consistent interpretation of the utterance.

SEMANT can use contextual information to discard the incorrect portion of a partially correct fragment. This capability is most clearly illustrated by a hypothetical problem. Suppose HEARSAY-II fails to recognize the utterance "DID REDDY WRITE ANY ARTICLES ABOUT LEARNING" but generates the fragments "DID REDDY WRITE ANY ARTICLES ABOUT" and "INTERESTED IN LEARNING." Figure 6.8 shows how these fragments are matched into the conceptual hierarchy. The first fragment supports AUTHOR and QUERY!AUTHOR!TOPIC. The second fragment supports TOPIC and SELECTION. The generated interpretation is shown in Figure 6.9. It incorporates the highest-scored primary concept REQUEST, in preference to the lower-scored SELECTION. The incorporated conjunctive concept QUERY!TOPIC!AUTHOR is supported by both AUTHOR and TOPIC. Since these two concepts are mutually consistent, they are both included in the interpretation, even though they are supported by different fragments. Consequently the interpretation incorporates the correct word "LEARNING" from the second fragment, but discards the incorrect sub-fragment "INTERESTED IN," since the SELECTION concept is not part of the interpretation.

Another feature of SEMANT is its ability to identify which data is missing. Suppose in the preceding example HEARSAY-II generates the fragment "PUBLISHED IN IJCAI" instead of the fragment "INTERESTED IN LEARNING." Figure 6.10 shows the matching of the generated fragments. This example differs from the preceding example in that the second fragment contains no information consistent with the first. The maximal consistent interpretation, shown in Figure 6.11, is supported only by the first fragment. It incorporates the conjunctive concept QUERY!TOPIC!AUTHOR, whose son TOPIC is unsupported. Thus SEMANT can predict that the missing data (unrecognized portion of the utterance) includes data which would support TOPIC. Such a semantic prediction could be used to guide further efforts by HEARSAY-II to recognize the utterance (Hayes-Roth et al, 1976b; Hayes-Roth, Mostow, and Fox, 1977). Alternatively, it could be used as grounds for asking the user to repeat the topic (Hayes-Roth, Gill, and Mostow, 1977).

7. COMPLICATIONS

The problem of finding a maximal consistent interpretation of the data is complicated by a conflict between maximality and consistency. Maximality is defined in terms of a scoring metric on concept support. A correct metric function will score the nodes in such a way that a simple top-down walk that selects the highest-rated son of every disjunctive node will in fact generate the maximal consistent interpretation. Ideally, the scoring process should require a single bottom-up pass which visits each node at most once.

Unfortunately, the context-sensitive nature of consistency may preclude the realization of this ideal. The incorporation of a chunk of data as support for a high-level concept in an interpretation creates a commitment to incorporate subparts of that data chunk as support for lower-level concepts. This idea is illustrated in Figure 7.1. The high-level disjunctive concept QUERY is supported by the 5-word fragment (chunk) "DO ANY ARTICLES MENTION LEARNING," and has score 8. QUERY has two sons: QUERY!TOPIC, which is supported by the 2-word sub-fragment "MENTION LEARNING," and QUERY!SOURCE, which is supported by the 3-word fragment "IN IJCAI PROCEEDINGS." Accordingly, QUERY!TOPIC has score 2 and QUERY!SOURCE has score 3. Consider the behavior of a top-down walk which selects the highest-rated son of every disjunctive node. Such an algorithm incorporates the fragment supporting QUERY into the interpretation, thereby morally committing itself to incorporate the subfragment supporting QUERY!TOPIC. The algorithm then violates its commitment by selecting QUERY!SOURCE over the lower-scored QUERY!TOPIC and consequently generates the inconsistent interpretation shown in Figure 7.2. The correct maximal consistent interpretation, shown in Figure 7.3, is constructed by fulfilling this commitment.

What exactly is the problem here? The inclusion of a chunk of data as support for a concept in an interpretation creates a commitment to include concepts supported by subchunks of that data. In short, the selection of support for a concept is context-sensitive, since it depends on the data chosen to support the concept's ancestors in the conceptual hierarchy. However, the scores assigned by a one-pass bottom-up notching algorithm are context-free. Consequently they do not always select the correct (maximal consistent) interpretation, as the preceding example illustrates. We see several possible approaches to solving this problem.

The first approach compensates for deficiencies of the context-free scoring function by introducing some search in the top-down selection of an interpretation. If incorporating the highest-scored son of a disjunctive node leads to an inconsistent interpretation, the next-highest node can be tried.

The second approach uses a context-sensitive scoring scheme so that a non-backtracking top-down walk will work correctly. One way to do this is to notch concepts using a bottom-up process, but under certain circumstances to reevaluate descendants of a concept in the context of its supporting data. Note that the context-free score of a concept is an upper bound on the size of the largest consistent set of data supporting that concept. The application of additional constraints such as

context can only decrease the size of this set. Thus one indication that a node may need to be reevaluated in context is a failure to support it with a data set as large as its score.

These approaches have a common defect: the "maximal consistent interpretations" they generate may fail to satisfy certain consistency constraints not represented in the structure of the conceptual hierarchy. For example, a consistent interpretation of a spoken utterance cannot be supported by two conflicting data fragments (word sequences) spanning the same temporal interval of the utterance. Similarly, a consistent interpretation of a scene cannot assign two conflicting labels to the same region. Representation of such constraints in the conceptual hierarchy appears to require the propagation of temporal or spatial information through the hierarchy and the parameterization of node relations (currently AND, XOR, UNION) to test such information for consistency.

The third approach, currently under development, uses a parameterized conceptual hierarchy. After data support is attached to appropriate nodes in the hierarchy, "notch tokens" are propagated up from the leaves of the hierarchy. Each token represents a particular set of data supporting (instantiating) a concept. A token is propagated upward from a node by passing copies of it to the node's parents. When tokens are passed to a conjunctive or union node from several of its subconcept nodes, a new token is formed representing the combined data supporting the concept. If this data is mutually inconsistent, it is split into maximal consistent subsets, each represented by a new token.

Such parameterized conceptual hierarchies provide stronger domain models by incorporating additional consistency constraints. However, experience with such parameterized hierarchies shows that they involve more computation than do unparameterized hierarchies, since the various instances (tokens) of each concept (distinguished by their different parameter values) must be processed (e.g., scored) separately (Hayes-Roth and Mostow, 1975; Mostow and Hayes-Roth, 1977). Thus it may be desirable to develop a hybrid matching scheme that tests as many constraints as possible in a parameterless conceptual hierarchy and only tests residual constraints afterwards. In this fourth approach, the parameterless hierarchy functions as a weak model of the domain. The matching process efficiently filters the data, generating maximal interpretations satisfying all the constraints embedded in the hierarchy. If such an interpretation fails to satisfy the residual constraints, the matcher finds the next highest-scored interpretation. This process continues until an interpretation satisfying all constraints is found. This is the desired maximal consistent interpretation.

8. DISCUSSION

Several points about the presented method should be emphasized.

8.1 Importance of Chunking

Chunking contributes to the success of our method in several ways. The chunking process identifies semantically meaningful configurations of data, i.e., configurations corresponding to (substructures of) known concepts. This structuring of the data is essential to the construction of a coherent interpretation. Chunking provides information about data consistency insofar as the data in a chunk is mutually consistent. This information is incorporated in the process of constructing an interpretation. Chunking also provides information about the contextual credibility of data insofar as the data in a chunk is mutually confirmatory. This information, represented by varying chunk size, is incorporated in the scoring metric and helps discriminate between alternative interpretations.

8.2 Importance of Hierarchy

Another important aspect of the method is its use of hierarchical structure to embed constraints on data consistency. Mutual exclusion, mutual necessity, and mutual consistency of subconcepts are modelled respectively by disjunctive, conjunctive, and union nodes. Any subgraph of the

hierarchy in which no disjunctive node has more than one son constitutes a consistent (possibly incomplete) conceptual structure. The data supporting such a structure consequently satisfies many constraints on data consistency.

The hierarchical structure also permits the identification of missing data. Mutual necessity of concept constituents is represented by conjunctive nodes. Unsupported sons of conjunctive concepts incorporated in an interpretation therefore represent missing constituents.

8.3 Constraint Reduction

The parameterless nature of the conceptual hierarchy precludes the embedding of certain types of constraints. In the speech understanding example, since temporal information is not propagated through the hierarchy, temporal constraints such as adjacency, ordering, and non-overlap are not represented in the hierarchy. In the vision domain, since location information is not propagated through the hierarchy, spatial constraints such as alignment, adjacency, proximity, ordering, and non-overlap are not represented. This reduction of constraint allows semantically consistent chunks to be incorporated in an interpretation even if they don't conform to a stronger (more constrained) model of the domain. This aspect of the representation permits increased flexibility in the matching process, in that the constraints on the integration of multiple chunks into an interpretation are weaker than the constraints on the local integration of data into individual chunks. Furthermore, the simplicity of the representation should make the matching process faster than methods which represent consistency constraints as tests on propagated parametric information. The disadvantage of the simpler representation is its greater potential for constructing inconsistent interpretations.

8.4 Power of the Method

The presented method interprets sets of hierarchically structured, possibly mutually inconsistent chunks of data. Although it exploits information incorporated in the chunk structure, the method is not restricted to accepting or rejecting chunks in an all-or-none fashion; the method can discard part of a chunk in order to construct a consistent interpretation which incorporates the remainder of the chunk. The constructed interpretation corresponds to a highly partial-matched substructure of the conceptual hierarchy. Unsupported constituents of the substructure identify missing data.

8.5 Applications of the Current Implementation

SEMANT was originally developed to interpret sentences and sentence fragments recognized by HEARSAY-II (Hayes-Roth et al, 1976b). In addition to this task, SEMANT has been applied to the interpretation of ungrammatical sentences. A sentence is chunked into its maximal grammatical subsequences, which are input to SEMANT as fragments. SEMANT then integrates the fragments into an interpretation of the sentence. This method has been used to correctly interpret sentences containing errors of insertion, deletion, substitution, repetition, and re-ordering.

9. CONCLUSIONS

We have designed and implemented a method for identifying and interpreting maximal consistent subsets of data in hierarchically modelled domains characterized by data error and inconsistency. The implementation has correctly interpreted spoken sentence fragments recognized by the HEARSAY-II speech understanding system. It has also been used successfully to interpret typed-in ungrammatical sentences.

The method appears applicable to many tasks (e.g., speech understanding, natural language understanding, scene analysis, medical analysis) requiring matching of errorful data against complex, hierarchically describable structures. When missing data or the inherent nature of the task causes the structures to be incompletely instantiated, partial-matching of these structures provides consistent, meaningful interpretations of the data.

The continuing progress of AI beyond toy problems will be characterized by intelligent

programs performing real-world tasks. Such programs will have to handle uncertain, inconsistent data corresponding only approximately to known concepts. The problem of identifying consistent subsets of data and integrating them into a hierarchically organized conceptual knowledge base can accordingly be expected to assume increasing importance.

10. BIBLIOGRAPHY

- Erman, L.D., 1975, Overview of the HEARSAY Speech Understanding Research. Computer Science Research Review (1974-1975), Computer Science Department, Carnegie-Mellon University, Pittsburgh, PA.
- Erman, L.D., 1977, A Functional Description of the HEARSAY-II Speech Understanding System, Proc. 1977 IEEE Inter. Conf. on Acoustics, Speech and Signal Processing (to appear).
- Hayes-Roth F., L.D. Erman, M.S. Fox, and D.J. Mostow, 1976a, Syntactic processing in HEARSAY-II, Speech Understanding Systems: Summary of Results of the Five-Year Research Effort, Department of Computer Science, Carnegie-Mellon University, Pittsburgh.
- Hayes-Roth F., M.S. Fox, G. Gill, and D.J. Mostow, 1976b, Semantics and pragmatics in the HEARSAY-II speech understanding system, Speech Understanding Systems: Summary of Results of the Five-Year Research Effort, Department of Computer Science, Carnegie-Mellon University, Pittsburgh.
- Hayes-Roth F., V. Lesser, D.J. Mostow, and L.D. Erman, 1976c, Policies for rating hypotheses, halting, and selecting a solution in the HEARSAY-II speech understanding system, "Speech Understanding Systems: Summary of Results of the Five-Year Research Effort, Department of Computer Science, Carnegie-Mellon University, Pittsburgh.
- Hayes-Roth F., and D.J. Mostow, 1975, An automatically compilable recognition network for structured patterns. Proceedings of the Fourth International Joint Conference on Artificial Intelligence. Cambridge: MIT.
- Hayes-Roth F., G. Gill, and D.J. Mostow, 1976, "Discourse analysis & task performance in the HEARSAY-II speech understanding system, " Speech Understanding Systems: Summary of Results of the Five-Year Research Effort, Department of Computer Science, Carnegie-Mellon University, Pittsburgh.
- Hayes-Roth F., D.J. Mostow, and M.S. Fox, 1977, Understanding Speech in the HEARSAY-II System, To appear: Natural Language Communication With Computers, L. Bolc (Ed.), Berlin: Springer-Verlag.
- Mostow D.J., and F. Hayes-Roth, 1977, A production system for speech understanding. In D. A. Waterman and F. Hayes-Roth (Eds.), Pattern Directed Inference Systems. New York: Academic Press, (in press).




FIGURE 2.1

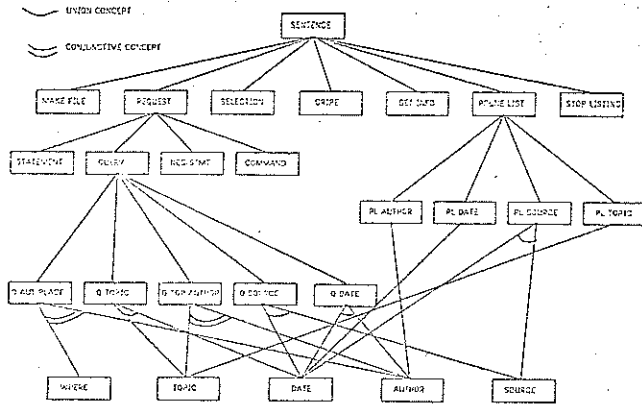


FIGURE 6.2 PORTION OF CONCEPTUAL HIERARCHY FOR ABSTRACT RETRIEVAL TASK

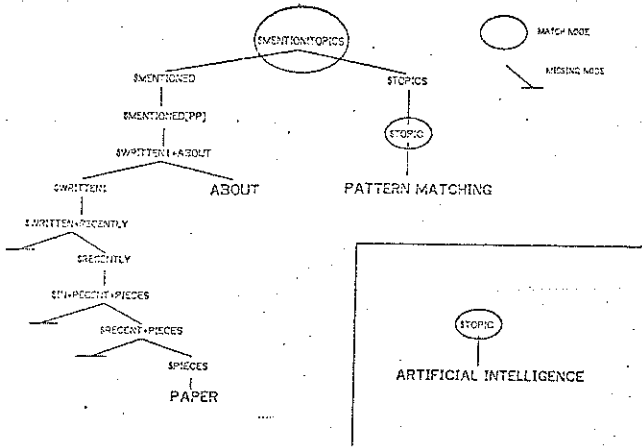


FIGURE 6.1

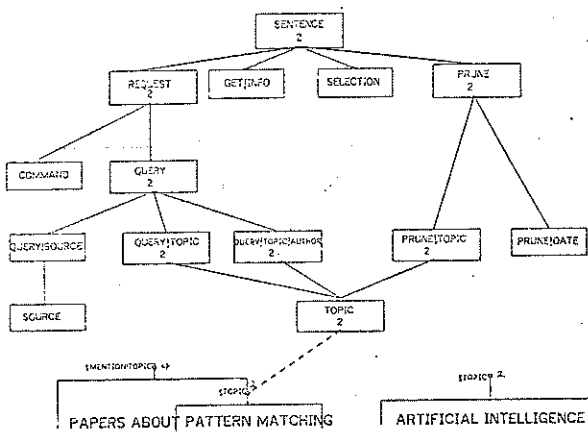


FIGURE 6.2

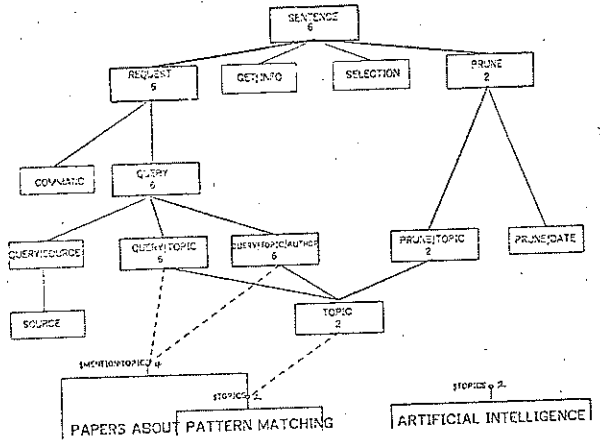


FIGURE 6.3

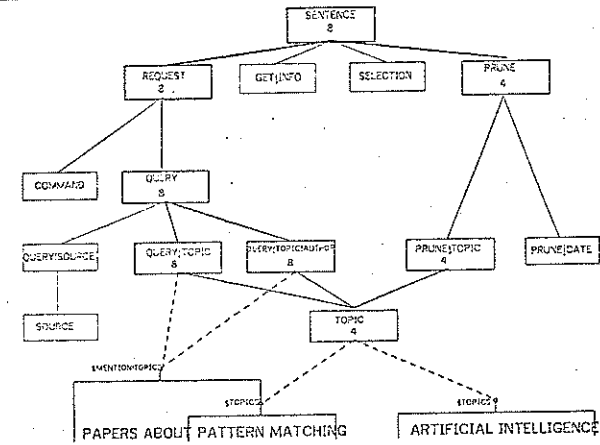


FIGURE 6.4

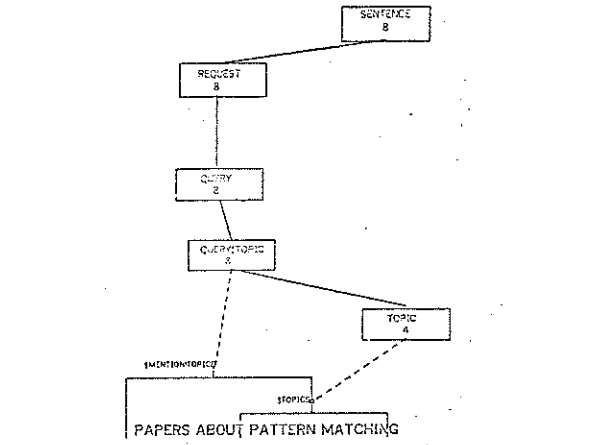


FIGURE 6.5

