ITP: DESCRIPTION OF THE INTERPRETEXT SYSTEM AS USED FOR MUC-3

Kathleen Dahlgren

Carol Lord

Hajime Wada

Joyce McDowell

Edward P. Stabler, Jr.

Intelligent Text Processing, Inc.

1310 Montana Avenue, Suite 201

Santa Monica, CA 90403

213-576-4910

Internet: 72550,1670@compuserve.com

The ITP System for MUC3 is diagrammed in Figure 1. The three major modules handle different units of processing: the Message Handler processes a message unit; the ITP NLU Module processes a sentence and builds a Cognitive Model of the message; and the MUC3 Template Reasoning Module processes a segment of discourse.

Figure 1. ITP Text Understanding System for MUC3

The Message Handler identifies an individual MUC message unit and controls the flow of operations of the whole system for one message. It reads sentences, and sends one sentence at a time to the ITP Natural Language Understanding Module. It prunes temporal and locative expressions, which are added back later to modify events. The ITP NLU Module parses one sentence, and maps its parse tree onto a Discourse Representation Structure (DRS) [8],[1]. This DRS is added to the DRS for the whole segment. Temporal/locative reasoning, anaphora resolution and discourse segmentation operate upon this larger DRS. When reading sentences for one message is completed, the Message Handler receives the resulting analysis from the ITP NLU Module. The resulting Cognitive Model is a segmented discourse structure for the whole message. The Message Handler sends one discourse segment at a time to the MUC3 Template Reasoning Module. It fills all of the possible templates based on the information in the segment.

Natural Language Understanding Module

Intelligent Text Processing (ITP) accepts at face value the complexity of natural language, and applies deep natural language understanding techniques. Natural language is both highly ambiguous (the same pattern means many different things), and redundant (the same meaning can be expressed with many different patterns). ITP's Natural Language Understanding Module, which had been prototyped before MUC-3, analyzes this complex structure, and unravels its meaning layer by layer. It builds a Cognitive Model of the text content which is an unambiguous logical form. The Cognitive Model tracks entities and events across the discourse, reflects the connections between events, and separates discourse segments. The result is a refined, precise interpretation of text content.

Text cannot be disambiguated or interpreted without world knowledge. The construction of knowledge bases for natural language understanding has been the main bottleneck in the field [9]. ITP's approach, founded upon Naive Semantic Theory [4], associates a shallow layer of world knowledge with every word. The form of the knowledge derives from cognitive psychological research [10],[3],[7]. Constraints on the amount of knowledge have been discovered in the process of building a computational text understanding system. The Naive Semantic lexicon contains two types of information, ontological and generic. Ontological information classifies objects and events into the main distinctions people recognize, such as real vs abstract and natural vs social. Generic information includes properties of objects *embassies play a role in diplomacy* and implications of events *bombing typically results in damage*. Word senses are distinguished, making the interpretation very precise. This word-based lexicon makes ITP's knowledge engineering effort completely transportable. All of the lexicon moved over unchanged into MUC-3, and the MUC-3 knowledge engineering effort required the addition of only 240 new words. The Naive Semantic Lexicon informs algorithms at all levels of analysis: structural disambiguation, word sense disambiguation, formal semantic translation, and discourse-level anaphora resolution, coherence reasoning and segmentation.

The flow of a message through the module is depicted in Figure 2. Sentences of a message are passed one at a time to the parser. It assigns a single structure to the sentence. ITP's disambiguation modules reform the structure and select word senses. This prevents the combinatorial explosion which slows down competitor systems. Next the Sentential Formal Semantic Module translates each sentence into a partial DRS. This step takes the English into a form very close to first order logic, and incorporates the effects of operators such as negation which vitally alter the interpretation. Next the Discourse Formal Semantic Module adds each sentence to the overall DRS for the text. At this point the entities and events in the new sentence are linked up with those in prior sentences. The time course of events in the text is traced by temporal reasoning. Only the sentences within a given segment are considered in the anaphoric and temporal reasoning. A more sophisticated method under development would extract a discourse tree in which dominating segments contribute to anaphoric reasoning.

Figure 2. ITP Natural Language Understanding Module

Syntactic Parser

For MUC-3 ITP used a loaner parser. (A new wide coverage parser based on Government and Binding Theory (Chomsky 1981) is under development. It will disambiguate words and structure during parsing, solving the combinatorial explosion problem.) The loaner parser ITP used for MUC-3 assigns a single X-bar type parse to the sentence. The ITP Prepositional Phrase Attachment Algorithm reforms the parse to select the correct structure in sentences with post-verbal PP's. A post-verbal PP introduces three possible parses, as in Figure 3. In *The guerrillas attacked the outpost with fury*, the PP modifies the sentence. In *The guerrillas attacked the outpost with grenades*, the PP modifies the verb phrase. In *The guerrillas attacked the outpost with a sentry*, the PP modifies the noun phrase *the outpost*. The ITP Prepositional Phrase Disambiguation Algorithm selects the correct attachment using preposition-specific rules and the Naive Semantic Lexicon [6]. Next, the ITP Word Sense Disambiguation Algorithm selects the appropriate senses of words. It takes into account fixed and frequent phrases, syntactic context and semantic context. For example, in the first sentence of message 99, the verb *bomb* has two readings, one which refers to a physical act, and another which refers to the social process of failure. The second reading is intransitive, but the verb *bomb* in S2 has an object. Thus the algorithm selects the first reading.

Figure 3. PP Attachment

Sentential Formal Semantics

The sentential formal semantic module translates the parse into a logical formula which conveys truth conditions for the sentence. It expresses the conditions in the actual world (or some possible world) which would have to exist in order for the sentence to be true. Pruned S1 of message 99, *Terrorists bombed the embassies of the PRC and the Soviet Union*, contains *terrorists*, *embassies*, *PRC*, *Soviet Union* and an event of bombing which relates the terrorists and the embassies. A translation function converts the parse into a DRS as in Figure 4. The graphical representation of a DRS is a box with a list of indexed entities ascribing objects and events which are introduced into the discourse (e1,g1...) and a list of conditions which are first-order representations of properties and relations expressed in sentences. Thus we see in the top portion of the DRS an event e1 in which the entity g1 (terrorists) relates to the entity the1 (embassies) in an bombing event in which entity g1 does the bombing and entity the1 is bombed. Similarly, the event e1 occurred tonight (the same date as the message dateline).

Figure 4. Discourse Representation Structure

Negations and modals introduce opaque, embedded subDRSs. For example, in S7 of message 99, there is a modal *could*: *Police sources, however, have said the attacks could have been carried out by the Maoist Shining Path group or the Guevarist Tupac Amaru Revolutionary Movement group*. ITP's Sentential Formal Semantic Module produces a DRS as in Figure 5.

Figure 5. Treatment of Modal

With this DRS, it is easy to infer that a sentence *the Maoist Shining Path group or the Guevarist Tupac Amaru Revolutionary Movement group carried out the attacks* is not asserted as true in the message. Its corresponding proposition is embedded under the modal operator which introduces a possible world. Therefore, this representation allows ITP to correctly provide an empty PERPETRATOR slot for the second bombing incident.

Interpretation of message 99 requires analyses of coordination and negation. When negations appear in verb phrases, they are treated as sentential operators. Any propositional content under the scope of negation is considered as part of a negated world, not true in the current world. When an NP is quantified with *no* or *none*, we treat the quantified expression as a normal term with cardinality *none*. For example, *no injuries* in S2 of message 99 is interpreted as a case in which there is an injury whose cardinality is *none*. Therefore, a statement *there are many injuries* is not true in the world represented by the DRS segment for S2.

Coordination in the ITP treatment is represented by a predicate called 'collection'. The coordinate NP in S1, *embassies of the PRC and the Soviet Union*, is represented with a collection c1 consisting of the2 (the PRC) and the3 (the Soviet Union), as shown in the lower portion of Figure 3.

The Discourse Formal Semantic Module

The Discourse Formal Semantic Module includes 1) the temporal/locative information processing, 2) anaphor resolution, 3) coherence reasoning (not used in MUC-3), and 4) discourse segmentation. Each sentence is added to the overall DRS for a message to produce a Cognitive Model for the entire message.

Temporal/Locative Information Processing

Temporal and locative information processing determines a reporting time and a reporting location for each message, a reference time and a current location for each discourse segment, and an event time and a current location for each event and state in a sentence as basic parametric information. At the beginning of processing one message, the unit determines a reporting time and a reporting location for the message. In the processing of individual sentences, the unit checks whether new temporal/locative information is stored in the database by the Pruner in the Message Handler. If so, it retrieves the information, and converts it into appropriate predicates representing the original expressions. For example, *San Isidro* in S3 of message 99 is disambiguated from *San Isidro* in El Salvador due to the fact that the current reporting place is Lima, Peru, and is represented in the DRS by a predicate 'current_location(e3, '[SAN ISIDRO : DISTRICT [LIMA : CITY [PERU : COUNTRY]]])'. The locative information is attributed to an exploding event ascribed by an event marker e3.

Temporal expressions are processed similarly and are represented by a predicate 'event time', 'reference time', and 'reporting time.' These predicates have four arguments: 1) an event marker for which this temporal information is attributed; 2) a tense operator representing a precedence relation between two events such as <, =<, >, >=, ^; 3) a universal representation of Year, Month, Day, and Day Name; 4) an absolute temporal value. We make three way distinctions among various temporal expressions: absolute time expressions such as *July 15, 1989*, relative time expressions such as *two days earlier*, and relative to reporting time expressions such as *3 months ago*. Parallel to the explicit knowledge representation for locative information, temporal expressions are converted to absolute temporal values which represent a number of hours since 00:00 January 1, 1. For example, *October 25, 1989* has an absolute temporal value of 17433936, as shown in Figure 4. The absolute value and the universal representation of time is used in computation of precise ordering relations among event times.

The Discourse Segmentation Module

This module detects a shift of discourse focus when there is a change of time or place, or when an explicit segmenting connective such as *meanwhile* or *in sum* occurs. The shift is indicated by a segmentation marker in the DRS. We assume that each discourse segment involves topic events, participants, time, location, etc. News reports such as MUC3 texts are typically organized in this way. It simplifies the reasoning to recognize these segments, as described in our theory (Dahlgren [5]) because the events and participants within the segment all relate to each other and to the SAME template. The following figure illustrates the resulting segmentation structure for message 99.

Figure 6. Segmentation Structure for 99

The above segmentation structure clearly distinguishes five separate segments. These happen to correspond to five separate events. (The Template Reasoning Module can also handle cases in which more than one terrorist incident is reported inside the same segment). Segment 1 focuses on the first bombing event on the night of Oct. 25, 1989. Segment 2 is introduced by a transitional adverb, *meanwhile* in S4, and contains another bombing event. S8 is considered to introduce a new segment, Segment 3, since it names a new location *in Peru*, and contains a group of similar events in the past. *some 3 years ago* in S10 triggers a discourse focus shift. S14 introduces new temporal information, *today*, which introduces a new segment containing an event occurring on October 24, 1989. Note that each segment has its own temporal/locative locus for the topic event. The ITP Discourse Segmentation Module captures this fact by checking its reference time and current location predicates in the segmented DRS.

Among these five segments, reference times of Segment 3, *July, 1989*, and Segment 4, *some 3 years ago*, are out of the range of the criteria for a **RECENT** incident (within two months from the reporting date). The reference time for Segment 3 and 4 provides sufficient information for the Template Reasoning Module to recognize these segments as irrelevant for template filling. The ITP Segmentation Module enables the Template Reasoning Module to recognize five possible templates, discard Segment 3 and Segment 4 as irrelevant incidents, and precisely choose the correct three templates from Segments 1, 2, and 5. There is no problem of "template merging" which occurred in competitor systems which looked for a terrorism word in each individual sentence, posited a separate template for each and then tried to decide which of them should be merged under the same incident.

Anaphora Resolution

The Anaphora Resolution algorithm resolves anaphoric links for referential expressions appearing in the current sentence. Events and objects which are found to refer to the same entities are given the same reference marker. The ITP Anaphor Resolution algorithm resolves both individual (*he, she*) and event (*attacks*) anaphora. Clues are sought at all levels: syntax, formal semantics and naive semantics. As for syntax, the algorithm prefers as the antecedent of a subject pronoun a prior noun phrase which is a sentence subject as well. As for formal semantics, in the resolution of event anaphora, it looks for event type predicates as antecedents of event type nouns. Type information is directly displayed in the DRS reference markers. S5, *police said the attacks were carried out almost simultaneously and that the bombs broke windows and destroyed the two vehicles*, contains the anaphoric noun phrase *the attacks*. It is recognized as an event type noun phrase because it has been assigned an event type reference marker. The anaphora resolution algorithm looks for prior events as antecedents, not objects such as *bombs*. Naive Semantics is used in anaphora resolution. In message 01 a portion of the text reads *Spokesman Isaacs said that*

the attack, which resulted in the death of a civilian ... was not against the farm. He added that ... In seeking an antecedent for *he*, the anaphora resolution algorithm uses Naive Semantics to infer that *say that* and *add that* are both verbs of saying. It therefore chooses *Isaacs* as antecedent rather than *civilian*, because an agent is more likely to continue a like action in the next sentence of a discourse. *the attack* is excluded because it is an event, not a person.

In the process of resolution, the algorithm observes constraints imposed by the segmentation and logical structures of the DRS. In particular, antecedents are sought only within the current segment. For example, the definite description *the bombs* in S5 should be resolved with *two bombs* in the previous sentence S4, *meanwhile, two bombs were thrown at a ussr embassy vehicle ...* Acceptable antecedents are in the same segment, or in a dominating segment. Since *two bombs* is in the same segment as *the bombs*, it is an acceptable antecedent. The bombing event in S1 and the bombs in S2 are not a possible antecedents, because they are in Segment 1.

The example of *attacks* in S5 illustrates a limitation of the current segmentation algorithm, which does not provide the dominating segment (Segment 1) as input to anaphora resolution, so that it does not find the other attack (on the PRC embassy) to get both attacks as a collective antecedent.

Template Reasoning Module

The Template Reasoning Module is MUC-3 specific. However, much of it consists of very simple, short algorithms which inspect elements of the Cognitive Model and reason about them using Naive Semantics. To find a terrorist incident, the code looks for a event in the Cognitive Model named by a verb which has as consequence in the Naive Semantic Lexicon some typical consequence of terrorism, such as *injure*, *damage*, or *destroy*. This includes events introduced by nominals like *destruction* and *bombing*. If no such event is found in a segment, the code looks for a noun which is attached to *weapon* in the ontology, such as *bomb* and *gun*.

The type of incident is determined by certain implications of verbs (and nominals). Any verb which has an explosion as consequence indicates a bombing, for example.

Perpetrators are the first argument (agent) in an event condition. The perpetrator entity must be sentient in the ontology. Sentients are any thinking beings introduced into the discourse by name (where the name is not a location), a noun attached to "role" in the ontology, or some other noun which inherits "sentient". (Our poor performance on perpetrators was due to failure of our code for handling proper names, and one other bug.) The Template Reasoning Module could easily find the relationships among empty head nouns such as *group*, and modifiers like *maoist* and *shining path*, as they were all given the same reference marker in the DRS, and thereby determine that "group" was clandestine in the Naive Semantics for *shining path*. A problem was the reconstruction of a surface string pattern like "maoist shining path group".

In determining targets, the Template Reasoning Module looks for second argument (object) of event conditions identified as terrorist. The difference between human and physical targets is determined by the ontological attachments of the nouns describing the target. Physical targets had to be buildings, vehicles, etc. Generalized reasoning code looked for all descriptors of a particular reference marker. Target type is determined by finding attachments of the descriptors in the ontology such as "office" or "vehicle". Refinements like "diplomat" were extracted from generic knowledge in the Naive Semantic lexicon. For example, the function of a bus is *transportation*, and an embassy plays a role in *diplomacy*.

Conclusion

The Naive Semantic lexicon and all of its feature types, the Sentential Formal Semantic module and the Discourse Semantic module were all brought to bear on the MUC-3 task without change (except addition of vocabulary). This means that ITP's deep natural language approach proved useful for a very difficult task. Given that the parser only returned half of the sentences in the messages, the ITP Natural Language Understanding Module and the Template Reasoning Module performed very well. The MUC-3 task uncovered bugs in the NLU Module, some of which were fixed for MUC-3, others not. ITP expects far better performance with our new parser and bug-free code.

References

- [1] Asher, N. 1987. A Typology for Attitude Verbs and Their Anaphoric Properties. *Linguistics and Philosophy*. 10:125-198.
- [2] Chomsky, N. (1981). *Lectures on Government and Binding*. Foris, Dordrecht.
- [3] Dahlgren, K. 1985. The Cognitive Structure of Social Categories. *Cognitive Science* 9:379-398.
- [4] Dahlgren, K. (1988). *Naive Semantics for Natural Language Understanding*. Kluwer Academic Publishers, Norwell, Mass.
- [5] Dahlgren, K. (1989). "Coherence Relation Assignment," in *Proceedings of the Cognitive Science Society*, pp.588-596.
- [6] Dahlgren, K., J.P. McDowell and E.P. Stabler, Jr. 1989. "Knowledge Representation for Commonsense Reasoning with Text," in *Computational Linguistics*. 15:149-170.
- [7] Graesser, A. and Clark, L. 1985. *Structure and Procedures of Implicit Knowledge*. Norwood, NJ: Ablex.
- [8] Kamp, H. (1981). "A Theory of Truth and Semantic Representation," in Gronendijk, J.; T. Janssen; and M. Stokhof, editors, *Formal Methods in the Study of Language*, Mathematisch Centrum, Amsterdam.
- [9] Maybury, Mark T. 1990. "Future Directions in NLP: The BBN Natural Language Symposium," in *AI Magazine*.
- [10] Rosch, E., C. B. Mervis, W. D. Gray, D. M. Johnson, and P. Boyes-Braem. 1976. "Basic Objects in Natural Categories," in *Cognitive Psychology*. 8:382-439.