

ALGORITHMS THAT LEARN TO EXTRACT INFORMATION — BBN: TIPSTER PHASE III

*Scott Miller, Michael Crystal, Heidi Fox, Lance Ramshaw, Richard Schwartz,
Rebecca Stone, and Ralph Weischedel*

BBN Technologies
70 Fawcett Street
Cambridge, MA 02138
weischedel@bbn.com

ABSTRACT

All of BBN's research under the TIPSTER III program has focused on doing extraction by applying statistical models trained on annotated data, rather than by using programs that execute hand-written rules. Within the context of MUC-7, the SIFT system for extraction of template entities (TE) and template relations (TR) used a novel, integrated syntactic/semantic language model to extract sentence level information, and then synthesized information across sentences using in part a trained model for cross-sentence relations. At the named entity (NE) level as well, in both MET-1 and MUC-7, BBN employed a trained, HMM-based model.

The results in these TIPSTER evaluations are evidence that such trained systems, even at their current level of development, can perform roughly on a par with those based on rules hand-tailored by experts. In addition, such trained systems have some significant advantages:

- They can be easily ported to new domains by simply annotating fresh data.
- The complex interactions that make rule-based systems difficult to develop and maintain can here be learned automatically from the training data.

We believe that improved and extended versions of such trained models have the potential for significant further progress toward practical systems for information extraction.

INTRODUCTION

We believe that trained statistical models offer significant advantages for information extraction tasks. In this report on BBN's research under the TIPSTER III program, we describe a number of research efforts that developed fully-trained

systems whose extraction performance was close to the highest levels achieved by carefully optimized systems based on hand-written rules.

SIFT, the first system described, extracts entities and relations from text. On the sentence level, it combines syntactic and semantic knowledge in a novel way, thus taking advantage of the significant recent progress in statistical parsing and leveraging those techniques for information extraction. Knowledge of English syntax extracted from the Penn Treebank is automatically combined with semantically annotated training material in the target domain that identifies how the entities and relations of interest in the domain are signaled in text. At the message level, the local entities and relations identified within each sentence are then merged, and cross-sentence relations are identified using an additional trained model. The resulting system achieved the second-best score of those participating in the MUC-7 evaluation.

The second system described here is the *IdentiFinder*TM system for locating named entities. This system is a fully-trained, HMM-based model that learns from examples the contextual clues that help to identify names in the text.

STATISTICAL EXTRACTION OF ENTITIES AND RELATIONS

The SIFT system ("Statistically-derived Information From Text") combines a sentence-level model with message-level processing to merge elements and identify cross-sentence relations.

At the sentence level, SIFT employs a unified statistical process to map from words to semantic structures. That is, part-of-speech determination, name-finding, parsing, and relationship-finding

Report Documentation Page

*Form Approved
OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE OCT 1998		2. REPORT TYPE		3. DATES COVERED 00-00-1998 to 00-00-1998	
4. TITLE AND SUBTITLE Algorithms that Learn to Extract Information - BBN: TIPSTER Phase III				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) BBN Technologies,70 Fawcett Street,Cambridge,MA,02138				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES TIPSTER TEXT PROGRAM PHASE III: Proceedings of a Workshop held at Baltimore, Maryland, October 13-15, 1998. Sponsored by the Defense Advanced Research Projects Agency.					
14. ABSTRACT All of BBN's research under the TIPSTER III program has focused on doing extraction by applying statistical models trained on annotated data, rather than by using programs that execute hand-written rules. Within the context of MUC- 7, the SIFT system for extraction of template entities (TE) and template relations (TR) used a novel, integrated syntactic/semantic language model to extract sentence level information, and then synthesized information across sentences using in part a trained model for cross-sentence relations. At the named entity (NE) level as well in both MET-1 and MUC-7, BBN employed a trained, HMM-based model. The results in these TIPSTER evaluations are evidence that such trained systems, even at their current level of development, can perform roughly on a par with those based on rules handtailored by experts. In addition, such trained systems have some significant advantages ??? They can be easily ported to new domains by simply annotating fresh data. ??? The complex interactions that make rulebased systems difficult to develop and maintain can here be learned automatically from the training data. We believe that improved and extended versions of such trained models have the potential for significant further progress toward practical systems for information extraction.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

all happen as part of the same process. This allows each element of the model to influence the others, and avoids the assembly-line trap of having to commit to a particular part-of-speech choice, say, early on in the process, when only local information is available to inform the choice.

The SIFT sentence-level model was trained from two sources:

- General knowledge of English sentence structure was learned from the Penn Treebank corpus of one million words of Wall Street Journal text.
- Specific knowledge about how the target entities and relations are expressed in English was learned from about 500 K words of on-domain text annotated with named entities, descriptors, and semantic relations.

In the on-domain training data, the names and descriptors of relevant items (persons, organizations, locations, and artifacts) are marked, as well as the target relationships between them that are signaled syntactically. For example, in the phrase “GTE Corp. of Stamford”, the annotation would record a “location-of” connection between the company and the city. The model can thus learn the structures that are typically used in English to convey the target relationships. Doing extraction in a new domain would require fresh semantically annotated training data appropriate to the new domain, but the general syntactic knowledge acquired from the Penn Treebank would still be applicable.

After the sentence-level model has identified names, descriptors, and relationships that are syntactically signaled within each sentence, further message-level processing is required to link up entities mentioned more than once or in different sentences, and to try to identify cross-sentence relationships or those not syntactically signaled. After the names, descriptors, and local relationships have been extracted from the sentence-level decoder’s output, a merging process is applied to link multiple occurrences of the same name or of alternative forms of the name from different sentences. A second, cross-sentence model is then invoked to try to identify relationships that were not picked up by the decoder, such as when the two entities do not occur in the same sentence. Finally, some

additional fields required by the MUC answer specification are filled in using heuristic tests and a gazetteer database, and output filters are applied to select which of the proposed internal structures should be included in the output. We are actively exploring ways of integrating this message-level processing more closely with the sentence-level model, since an integrated statistical model is the only way in which to make every choice in a nuanced way, based on all the available information.

The following sections describe the sentence-level and message-level processing of the SIFT system in more detail.

SIFT’s Sentence-Level Model

Figure 1 is a block diagram of the sentence-level model showing the main components and data paths. Two types of annotations are used to train the model: syntactic annotations for learning about the general structure of English, and semantic annotations for learning about the target entities and relations. From these annotations, the training program estimates the parameters of a unified statistical model that accounts for both syntax and semantics. Later, when presented with a new sentence, the search program explores the statistical model to find the most likely combined semantic and syntactic interpretation.

Training Data

Our source for syntactically annotated training data was the Penn Treebank (Marcus *et al.*, 1993). Significantly, we do not require that syntactic annotations be from the same source, or cover the same domain, as the target task. For example, while the Penn Treebank consists of Wall Street Journal text, the target source for this evaluation was New York Times newswire. Similarly, although the Penn Treebank domain covers general and financial news, the target domain for the MUC-7 evaluation was space technology. The ability to use syntactic training from a different source and domain than the target is an important feature of our model.

Since the Penn Treebank serves as our syntactically annotated training corpus, we need only create a semantically annotated corpus. Stated generally, semantic annotations serve to denote the entities and relations of interest in the

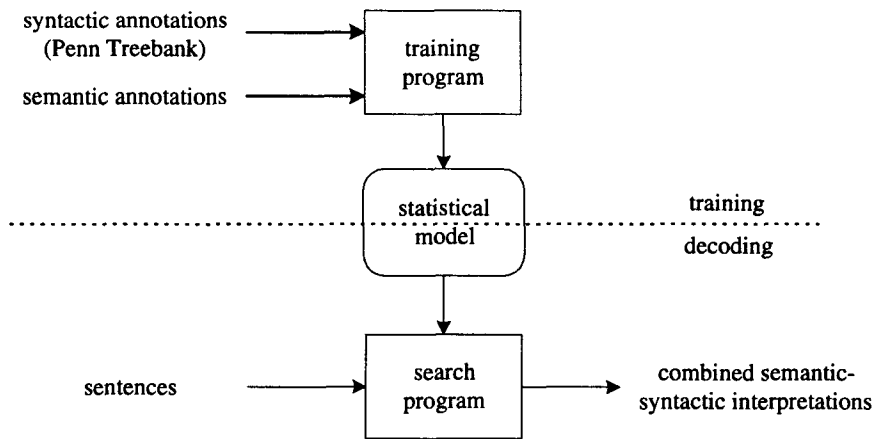


Figure 1: Block diagram of sentence-level model.

target domain. More specifically, entities are marked as either names or descriptors, with co-reference between entities marked as well. Figure 2 shows a semantically annotated fragment of a typical sentence.

From only these simple semantic annotations, the system can be trained to work in a new domain. To train SIFT for MUC-7, we annotated approximately 500,000 words of New York Times newswire text, covering the domains of air disasters and space technology. (We have not yet run experiments to see how performance varies with more/less training data.)

Semantic/Syntactic Structure

While our semantic annotations are quite simple, the internal model of sentence structure is substantially more complicated, since this combined model must account for syntactic structure as well as for entities and semantic relations. Our underlying training algorithm requires examples of these internal structures in order to estimate the parameters of the unified semantic/syntactic model. However, we do not

wish to incur the high cost of annotating parse trees. Instead, we use the following multi-step training procedure, exploiting the Penn Treebank:

- 1) Train the sentence-level model on the purely syntactic parse trees in the Treebank. Once this step is complete, the model will function as a state-of-the-art statistical parser.
- 2) For each sentence in the semantically annotated corpus:
 - a) Apply the sentence level model to syntactically parse the sentence, constraining the model to produce only parses that are consistent with the semantic annotation.
 - b) Augment the resulting parse tree to reflect semantic structure as well as syntactic structure.
- 3) Retrain the sentence-level model on the augmented parse trees produced in step 2. Once this step is complete, we have an integrated model of semantics and syntax.

Details of the statistical model will be discussed

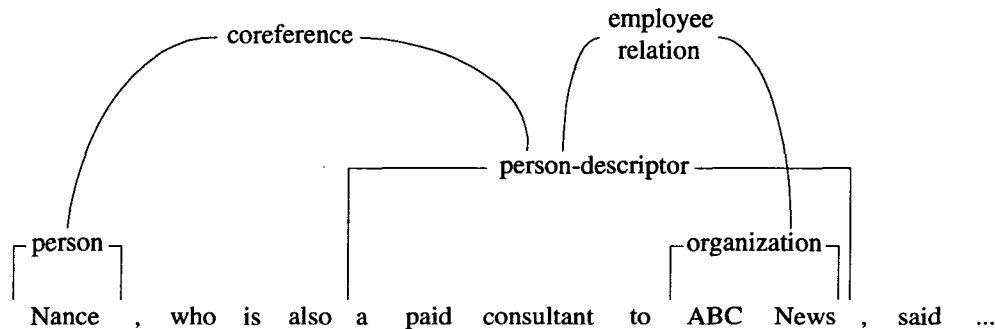


Figure 2: An example of semantic annotation.

later. For now, we turn our attention to (a) constraining the decoder and (b) augmenting the parse trees with semantic structure.

Constraints are simply bracketing boundaries that may not be crossed by any parse constituent. There are two types of constraints: hard constraints that cannot be violated under any conditions, and soft constraints, that may be violated only if enforcing them would result in no plausible parse. All named entities and descriptors are treated as hard constraints; the model is prohibited from producing any constituents that overlap either edge of the span of these elements. In addition, we attempt to keep possible appositives together through soft constraints. Whenever there is a co-referential relation between two entities that are either adjacent or separated by only a comma, we posit an appositive and introduce a soft constraint to encourage the parser to keep the elements together.

Once a constrained parse is found, it must be augmented to reflect the semantic structure. Augmentation is a five step process.

- 1) Nodes are inserted into the parse tree to distinguish names and descriptors that are not bracketed in the parse. For example, the parser produces a single noun phrase with no internal structure for "Lt. Cmdr. David Edwin Lewis". Additional nodes must be inserted to distinguish the descriptor, "Lt. Cmdr.," and the name, "David Edwin Lewis."
- 2) Semantic labels are attached to all nodes that correspond to names or descriptors. These labels reflect the entity type, such as person, organization, or location, as well as whether the node is a proper name or a descriptor.
- 3) For relations between entities, where one entity is not a syntactic modifier of the other, the lowermost parse node that spans both entities is identified. A semantic tag is then added to that node denoting the relationship. For example, in the sentence "Mary Fackler Schiavo is the inspector general of the U.S. Department of Transportation," a co-reference semantic label is added to the *S* node spanning the name, "Mary Fackler Schiavo," and the descriptor, "the inspector general of the U.S. Department of Transportation."

- 4) Nodes are inserted into the parse tree to distinguish the arguments to each relation. In cases where there is a relation between two entities, and one of the entities is a syntactic modifier of the other, the inserted node serves to indicate the relation as well as the argument. For example, in the phrase "Lt. Cmdr. David Edwin Lewis," a node is inserted to indicate that "Lt. Cmdr." is a descriptor for "David Edwin Lewis."
- 5) Whenever a relation involves an entity that is not a direct descendant of that relation in the parse tree, semantic *pointer labels* are attached to all of the intermediate nodes. These labels serve to form a continuous chain between the relation and its argument.

Figure 3 shows an augmented parse tree corresponding to the semantic annotation in Figure 2. Note that nodes with semantic labels ending in "-r" mark MUC reportable names and descriptors.

Statistical Model

In SIFT's statistical model, augmented parse trees are generated according to a process similar to that described in Collins (1996, 1997). For each constituent, the head is generated first, followed by the modifiers, which are generated from the head outward. Head words, along with their part-of-speech tags and features, are generated for each modifier as soon as the modifier is created. Word features are introduced primarily to help with unknown words, as in Weischedel *et al.* (1993).

We illustrate the generation process by walking through a few of the steps of the parse shown in Figure 3. At each step in the process, a choice is made from a statistical distribution, with the probability of each possible selection dependent on particular features of previously-generated elements. We pick up the derivation just after the topmost *S* and its head word, *said*, have been produced. The next steps are to generate in order:

1. A head constituent for the *S*, in this case a *VP*.
2. Pre-modifier constituents for the *S*. In this case, there is only one: a *PER/NP*.
3. A head part-of-speech tag for the *PER/NP*, in this case *PER/NNP*.

of the modifier word , t_m , the part-of-speech tag of the head word , t_h , and the head word itself, w_h :

$$P(w_m | c_m, t_m, t_h, w_h), \text{ e.g.} \\ P(\text{nance} | \text{per} / \text{np}, \text{per} / \text{nnp}, \text{vbd}, \text{said})$$

Finally, word features, f_m , for modifiers are predicted based on the modifier, c_m , the part-of-speech tag of the modifier word , t_m , the part-of-speech tag of the head word , t_h , the head word itself, w_h , and whether or not the modifier head word, w_m , is known or unknown.

$$P(f_m | c_m, t_m, t_h, w_h, \text{known}(w_m)), \text{ e.g.} \\ P(\text{cap} | \text{per} / \text{np}, \text{per} / \text{nnp}, \text{vbd}, \text{said}, \text{true})$$

The probability of a complete tree is the product of the probabilities of generating each element in the tree. If we generalize the tree components (constituent labels, words, tags, etc.) and treat them all as simply elements, e , and treat all the conditioning factors as the history, h , we can write:

$$P(\text{tree}) = \prod_{e \in \text{tree}} P(e | h)$$

Training the Model

Maximum likelihood estimates for all model probabilities are obtained by observing frequencies in the training corpus. However, because these estimates are too sparse to be relied upon, they must be smoothed by mixing in lower-dimensional estimates. We determine the mixture weights using the Witten-Bell smoothing method.

For modifier constituents, the mixture components are:

$$P'(c_m | c_p, c_{hp}, c_{m-1}, w_p) = \\ \lambda_1 P(c_m | c_p, c_{hp}, c_{m-1}, w_p) \\ + \lambda_2 P(c_m | c_p, c_{hp}, c_{m-1})$$

For part-of-speech tags, the mixture components are:

$$P'(t_m | c_m, t_h, w_h) = \lambda_1 P(t_m | c_m, w_h) \\ + \lambda_2 P(t_m | c_m, t_h) \\ + \lambda_3 P(t_m | c_m)$$

For head words, the mixture components are:

$$P'(w_m | c_m, t_m, t_h, w_h) = \lambda_1 P(w_m | c_m, t_m, w_h) \\ + \lambda_2 P(w_m | c_m, t_m, t_h) \\ + \lambda_3 P(w_m | c_m, t_m) \\ + \lambda_4 P(w_m | t_m)$$

Finally, for word features, the mixture components are:

$$P'(f_m | c_m, t_m, t_h, w_h, \text{known}(w_m)) = \\ \lambda_1 P(f_m | c_m, t_m, w_h, \text{known}(w_m)) \\ + \lambda_2 P(f_m | c_m, t_m, t_h, \text{known}(w_m)) \\ + \lambda_3 P(f_m | c_m, t_m, \text{known}(w_m)) \\ + \lambda_4 P(f_m | t_m, \text{known}(w_m))$$

Searching the Model

Given a sentence to be analyzed, the search program must find the most likely semantic and syntactic interpretation. More concretely, it must find the most likely augmented parse tree. Although mathematically the model predicts tree elements in a top-down fashion, we search the space bottom-up using a chart based search. The search is kept tractable through a combination of CKY-style dynamic programming and pruning of low probability elements.

Dynamic Programming: Whenever two or more constituents are equivalent relative to all possible later parsing decisions, we apply dynamic programming, keeping only the most likely constituent in the chart. Two constituents are considered equivalent if:

1. They have identical category labels.
2. Their head constituents have identical labels.
3. They have the same head word.
4. Their leftmost modifiers have identical labels.
5. Their rightmost modifiers have identical labels.

Pruning: Given multiple constituents that cover identical spans in the chart, only those constituents with probabilities within a threshold of the highest scoring constituent are maintained; all others are pruned. For purposes of pruning, and only for purposes of pruning, the prior probability of each constituent category is multiplied by the generative probability of that

constituent (Goodman, 1997). We can think of this prior probability as an estimate of the probability of generating a subtree with the constituent category, starting at the topmost node. Thus, the scores used in pruning can be considered as the product of:

1. The probability of generating a constituent of the specified category, starting at the topmost node.
2. The probability of generating the structure beneath that constituent, having already generated a constituent of that category.

The outcome of the search process is a tree structure that encodes both the syntactic and semantic structure of the sentence, so that the TE entities and local TR relations can be directly extracted from these sentential trees.

SIFT's Message-Level Processing

The sentence-level model in SIFT predicts names, descriptors, and relationships that are cued by the local sentence structure, but it considers each sentence in isolation. Merging such information between sentences is an important and difficult problem in information extraction. The information that indicates the presence of a template relation is often distributed across multiple sentences, and this merging problem would naturally become even more severe when trying to extract more complex structures like full scenario templates. We have explored various approaches to this merging problem in our TIPSTER research.

Our overall goal is to use trained and integrated models where possible, particularly for all of the language understanding. For some portions of SIFT's message-level processing, we used hand-written rules combined with external sources like gazetteers. The MUC-7 deadlines caused us to use an existing alias process for merging names rather than implementing a statistical alias procedure. In the current system, simple heuristic code handles the filling of the type and country fields that are required by the MUC specification, and the distinction between substantial and non-substantial descriptors. (The MUC guidelines call for ignoring certain descriptors like "the company".)

A trained cross-sentence relation model is used to identify template relations that link entities across different sentences. This model was

trained on 200 articles annotated with full MUC answer keys, so that even non-local relations were marked. (That level of semantic annotation was available for only a small subset of the data used to train the sentence-level model.) The model applies a set of structural and contextual features that help to indicate when such a relation might be present. Feature counts from the training data are used to estimate the probability of a relationship between each possible pair of entities mentioned in separate sentences in the text.

While the cross-sentence model is currently applied as a separate step after the sentence-level decoding is complete, we are exploring various approaches toward integrating the two models more closely, and also toward doing more of the named entity merging and type field prediction by means of trained models.

Merging Named Entities

The first step in merging the results of the sentence-level model is to group together the different mentions of the same named entity. In SIFT, a set of heuristic rules were used for this. Different mentions of the same name (say, different mentions of "IBM") would be grouped, as would strings that were related in certain predictable ways, for example, by initials (linking "IBM" with "International Business Machines") or by the addition of a corporate designator (linking "International Business Machines" with "International Business Machines, Inc."). This merging process also tested whether one name was a prefix of the other, linking "Legg Mason Wood Walker, Inc." with "Legg Mason".

The Cross-Sentence Relation Model

The cross-sentence model then uses structural and contextual clues to hypothesize template relations between two elements that are not mentioned within the same sentence. Since 80-90% of the relations found in the answer keys connect two elements that are mentioned in the same sentence, the cross sentence model has a narrow target to shoot for. Very few of the pairs of entities seen in different sentences turn out to be actually related. This model uses features extracted from related pairs in training data to try to identify those cases.

It is a classifier model that considers all pairs of entities in a message whose types are compatible with a given relation; for example, a person and an organization would suggest a possible employment relation. For the three MUC-7 relations, it turned out to be somewhat advantageous to build in a functional constraint, so that the model would not consider, for example, a possible employment relation for a person already known from the sentence-level model to be employed elsewhere.

Given the measured features for a possible relation, the probability of a relation holding or not holding can be computed as follows:

$$p(rel | feats) = \frac{p(feats | rel)p(rel)}{p(feats)}$$

$$p(\sim rel | feats) = \frac{p(feats | \sim rel)p(\sim rel)}{p(feats)}$$

If the ratio of those two probabilities, computed as follows, is greater than 1, the model predicts a relation:

$$\frac{p(rel | feats)}{p(\sim rel | feats)} = \frac{p(feats | rel)p(rel)}{p(feats | \sim rel)p(\sim rel)}$$

We approximate this ratio by assuming feature independence and taking the product of the contributions for each feature.

$$\frac{p(rel | feats)}{p(\sim rel | feats)} \approx \frac{p(rel) \prod_i p(feats_i | rel)}{p(\sim rel) \prod_i p(feats_i | \sim rel)}$$

The cross-sentence feature model applies to entities found by the sentence-level model, which is run over all of the sentence-like portions of the text. An initial heuristic procedure checks for sections of the preamble or trailer that look like sentential material, that should be treated like the body text. There is also a separate handwritten procedure that searches the preamble text for any byline, and, if one is found, instantiates an appropriate employee relationship.

Model Features

Two classes of features were used in this model: structural features that reflect properties of the text surrounding references to the entities involved in the suggested relation, and content

features based on the actual entities and relations encountered in the training data.

Structural Features

The structural features exploit simple characteristics of the text surrounding references to the possibly-related entities. The most powerful structural feature, not surprisingly, was distance, reflecting the fact that related elements tend to be mentioned in close proximity, even when they are not mentioned in the same sentence. Given a pair of entity references in the text, the distance between them was quantized into one of three possible values:

Code	Distance Value
0	Within the same sentence
1	Neighboring sentences
2	More remote than neighboring sentences

For each pair of possibly-related elements, the distance feature value was defined as the minimum distance between some reference in the text to the first element and some reference to the second.

A second structural feature grew out of the intuition that entities mentioned in the first sentence of an article often play a special topical role throughout the article. The "Topic Sentence" feature was defined to be true if some reference to one of the two entities involved in the suggested relation occurred in the first sentence of the text-field body of the article.

Other structural features that were considered but not implemented included the count of the number of references to each entity.

Content Features

While the structural features learn general facts about the patterns in which related references occur and the text that surrounds them, the content features learn about the actual names and descriptors of entities seen to be related in the training data. The three content features in current use test for a similar relationship in training by name or by descriptor or for a conflicting relationship in training by name.

The simplest content feature tests using names whether the entities in the proposed relationship have ever been seen before to be related. To test

this feature, the model maintains a database of all the entities seen to be related in training, and of the names used to refer to them. The “by name” content feature is true if, for example, a person in some training message who shared at least one name string with the person in the proposed relationship was employed in that training message by an organization that shared at least one name string with the organization in the proposed relationship.

A somewhat weaker feature makes the same kind of test for a previously seen relationship using descriptor strings. This feature fires when an entity that shares a descriptor string with the first argument of the suggested relation was related in training to an entity that shares a name with the second argument. Since titles like “General” count as descriptor strings, one effect of this feature is to increase the likelihood of generals being employed by armies. Observing such examples, but noting that the training didn’t include all the reasonable combinations of titles and organizations, the training for this feature was seeded by adding a virtual message constructed from a list of such titles and organizations, so that any reasonable such pair would turn up in training.

The third content feature was a kind of inverse of the first “by name” feature which was true if some entity sharing a name with the first argument of the proposed relation was related to an entity that did *not* share a name with the second argument. Using the employment relation again as an example, it is less likely (though still possible) that a person who was known in another message to be employed by a different organization should be reported here as employed by the suggested one.

Training

Given enough fully annotated data, with both sentence-level semantic annotation and message-level answer keys recorded along with the connections between them, training the features would be quite straightforward. For each possibly-related pair of entities mentioned in a document, one would just count up the 2x2 table showing how many of them exhibited the given structural feature and how many of them were actually related. The training issues that did arise stemmed from the limited supply of answer keys and that the keys were not connected to the sentence-level annotations.

The government training and dry run data provided 200 messages’ worth of TE and TR answer keys. Those answer keys, however, contained strings without recording where in the text they were found. In order to train structural features from that data, we needed the locations of references within the text. A heuristic string matching process was used to make that connection, with a special check to ensure for names that the shorter version of a name did not match a string in the text that also matched a longer version of the same name.

Training the content features, on the other hand, did not require positional information about the references. The plain answer keys could be used in combination with a database of the name and descriptor strings for entities related in training to count up the feature probabilities for actually related and non-related pairs. The string database was collected first, and one-out training was then used, so that the rest of the training corpus provided the string database for training the feature counts on each particular message. The additional training data that was semantically annotated for training the sentence-level model but for which answer keys were not available could still also be used in building up the string database for the content features.

The probabilities based on the final feature counts were smoothed by mixing them with 0.01% of a uniform model.

Other Message Level Processing

After the cross sentence model has been applied, some further heuristic message-level processing is done before generating the answers in MUC template form. In one step, those portions of the preamble of the message, which includes the title and by-line, that are not English sentences are searched for a possible employment relation between the article author and the organization holding the copyright. A limited form of voting was also applied across messages, so that if the same name was identified by the sentence-level model as, say, an organization in one case and a person in another, only the plurality type is actually output. Heuristic models are used to fill in some additional required fields, distinguishing, for instance, between civilian, military, and government organizations; this could have been trained, but time did not permit this. Identifying the type and country of locations

is a simple process, benefiting greatly from gazetteer lookup.

Finally, a heuristic choice is made whether or not to output each element. For example, a descriptor that was not paired by the sentence-level processing with any named entity could either actually be an isolated descriptor or it could be one where the true link with a named entity was missed by the sentence-level model. Lacking at this point any trained model to distinguish those two cases, SIFT plays it safe by not outputting such entities.

SIFT System Examples

The main determinant of SIFT's performance is the sentence-level model, and the semantic structures that it produces. Secondary but still significant effects on performance come from the message-level processing steps that derive TE and TR output from the sentence-level decoder tree:

- Extracting elements and relations

```
(SINV
  ...
  (VP
    (VBD said))
  (PER/NP
    (PER/NPA
      (PER/NPP
        (NNP Eric)
        (NNP Stallmer)))
    (, ,)
    (PER-DESC-OF/NP-LINK
      (PER-DESC/NP-R
        (PER-DESC/NPA
          (NN spokesman))
        (ORG-OF/NP-PP-LINK
          (ORG-PTR/PP
            (IN for)
            (ORG/NP
              (ORG/NPA
                (DT the)
                (ORG/NPP
                  (NNP Space)
                  (NNP Transportation)
                  (NNP Association))))
          (LOC-OF/NP-PP-LINK
            (LOC-PTR/PP
              (IN of)
              (LOC-PTR/NPA
                (LOC/NPP
                  (LOC/NPP
                    (NNP Arlington))
                    (, ,)
                    (LOC/NPP
                      (NNP Virginia)))))))))))))
```

Example 1

- Merging TE elements
- Searching for additional relations with the cross-sentence model
- Filtering candidate entities and relations for output

This section will present examples from the output for one of the MUC-7 test messages, demonstrating the different effects that applied.

Example 1 shows a case where everything worked as planned.

Here the decoder correctly recognized a person name (PER/NPA) bound to a person descriptor (PER-DESC/NP-R). That descriptor contains an organization (ORG/NP) which in turn is linked to a location. The LINK and PTR nodes connect the descriptor with the person, the organization with the person descriptor (and thus indirectly with the person), and the location with the organization. In the post-processing, the person name is extracted, with the descriptor text is linked to it, the organization name is extracted, and the employment relationship noted. The organization is also linked to the nested location;

of the two location elements in the LOC phrase, the first is taken as the LOCALE field filler, while the second is looked up in the gazetteer to identify a country in which the locale value is then looked up.

Example 2 shows the effect of a decoder error.

```
(ORG/NP
 (ORG/NPA
  (ORG/NPP
   (NNP Bloomberg)
   (NNP Information)
   (NNP Television)))
 (, ,)
 (ORG-DESC-OF/NP-LINK
  (ORG-DESC/NP-R
   (ORG-DESC/NPA
    (DT a)
    (NN unit))
   (PP
    (IN of)
    (ORG/NPA
     (ORG/NPP
      (NNP Bloomberg)
      (NNP L.P.))))))
 (, ,)
 (ORG-DESC-OF/NP-LINK
  (ORG-DESC/NP-R
   (ORG-DESC/NPA
    (DT the)
    (NN parent))
   (PP
    (IN of)
    (ORG/NPA
     (ORG/NPP
      (NNP Bloomberg)
      (NNP Business)
      (NNP News))))))
 (, ,))
```

Example 2

Here the sentence-level decoder linked both organization descriptors back to the top-level named organization, while the correct reading would have attached the second descriptor to the nested "Bloomberg L.P.". The post-processing also therefore links both descriptor phrases to "Bloomberg Information Television" internally. Only the longest descriptor, however, is actually output, which in this case results in output of only the mistaken value.

Not surprisingly, a number of the decoder errors that affected output stemmed from conjunctions. In another paragraph, for example, the manufacturer organization name "Lockheed Space and Strategic Missiles" was incorrectly broken at the conjunction, causing the location relation with Bethesda to be missed.

The cross sentence model is the system component that tries to find further relations

beyond those identified by the sentence-level model. In the walk-through article, that component did not happen to succeed in finding any such relations. Example 3 shows the sort of relation that we would like that model to be able to get. There the sentence-level decoder did link Rubenstein to the organization descriptor "company", but since that descriptor was never linked to "News Corporation", the employee relation was missed. However, since News Corporation is mentioned both in that sentence and the following sentence, an improved cross sentence model would be one way of attacking such examples.

```
(PER-DESC/NP
 (PER-DESC/NP
  (PER-DESC/NPA-R
   (ORG-DESC-OF/NP-LINK
    (ORG-DESC/NP-R
     (NN company)))
    (NN spokesman))
  (PER-OF/NPA-LINK
   (PER-PTR/NPA
    (PER/NPP
     (NNP Howard)
     (NNP J.)
     (NNP Rubenstein))))))
```

Example 3

The last step in processing is the output filter, which heuristically determines whether a proposed constituent should be included in the output. Example 4 shows two examples where this filter overrode correct decoder structure.

```
(S
 (ART-DESC/NP-R
  (ART-DESC/NPA
   (DT A)
   (JJ Chinese)
   (NN rocket))
 (ART-PTR/VP
  (VBG carrying)
  (ART-DESC/NPA-R
   (DT an)
   (ORG/NPP
    (NNP Intelsat))
    (NN satellite))))
 (VP
  (VBD exploded)
  ...
```

Example 4

Here the decoder correctly identified both the artifact descriptors "A Chinese rocket" and "an Intelsat satellite", but the output filter chose not to include them. That choice was made because of frequent cases where an indefinite artifact descriptor not linked to any named artifact should not be output; an example from elsewhere in this message is "the last rocket I'd

recommend". But this example shows that this decision not to output such cases sometimes cost the system points.

SIFT System Results and Summary

The SIFT system worked by first applying the sentence-level model to each sentence in the message and then extracting entities, descriptors, and relations from the resulting trees, heuristically merging TE elements, applying the cross-sentence model to identify non-local relations, and finally filtering and formatting TE and TR templates for output. In the MUC-7 evaluation, the system's score on the TE task was 83% recall with 84% precision, for an F of 83.49%. Its score on TR was 64% recall with 81% precision, for an F of 71.23%.

Because most of the relations in the answer keys were locally signaled, the cross sentence model in this application adds only a small boost to the performance of the sentence-level model. When measured before the evaluation on 10 randomly-selected messages from the airplane crash domain training, the cross sentence model improved TR scores by 5 points. It proved a bit less effective on the 100 messages of the MUC-7 test set, improving scores there by only 2 points. (The F score on the formal test set with the cross sentence model component disabled was 69.33%.)

A STATISTICAL NAME-FINDER

Overview of the Identifinder HMM Model

For identifying named entities in text, BBN has developed the Identifinder™ trained named entity extraction system (Bikel, et. al., 1997), which utilizes an HMM to recognize the entities present in the text.

The HMM labels each word either with one of the desired classes (e.g., person, organization, etc.) or with the label NOT-A-NAME (to represent "none of the desired classes"). The states of the HMM fall into regions, one region for each desired class plus one for NOT-A-NAME. (See Figure 4.) The HMM thus has a model of each desired class and of the other text. Note that the implementation is not confined to the seven name classes used in the NE task; the particular classes to be recognized can be easily changed via a parameter.

Within each of the regions, we use a statistical bigram language model, and emit exactly one word upon entering each state. Therefore, the number of states in each of the name-class regions is equal to the vocabulary size, $|V|$. Additionally, there are two special states, the START-OF-SENTENCE and END-OF-SENTENCE states. In addition to generating the word, states may also generate features of that word. Features used in the MUC-7 version of the system include several features pertaining to numeric expressions, capitalization, and membership in lists of important words (e.g.

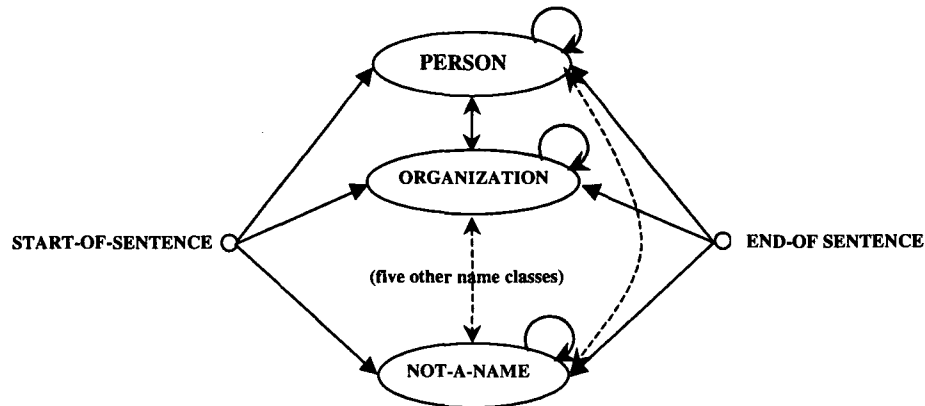


Figure 4: Pictorial representation of conceptual model

known corporate designators).

The generation of words and name-classes proceeds in the following steps:

1. Select a name-class NC , conditioning on the previous name-class and the previous word.
2. Generate the first word inside that name-class, conditioning on the current and previous name-classes.
3. Generate all subsequent words inside the current name-class, where each subsequent word is conditioned on its immediate predecessor.
4. If not at the end of a sentence, go to 1.

Whenever a person or organization name is recognized, the vocabulary of the system is dynamically updated to include possible aliases for that name. Using the Viterbi algorithm, we search the entire space of all possible name-class assignments, maximizing $\Pr(W, F, NC)$, the joint probability of words, features, and name classes.

This model allows each type of “name” to have its own language, with separate bigram probabilities for generating its words. This reflects our intuition that:

- *There is generally predictive internal evidence regarding the class of a desired entity.* Consider the following evidence: Organization names tend to be stereotypical for airlines, utilities, law firms, insurance companies, other corporations, and government organizations. Organizations tend to select names to suggest the purpose or type of the organization. For person names, first person names are stereotypical in many cultures; in Chinese, family names are stereotypical. In Chinese and Japanese, special characters are used to transliterate foreign names. Monetary amounts typically include a unit term, e.g., Taiwan dollars, yen, German marks, etc.
- *Local evidence often suggests the boundaries and class of one of the desired*

expressions. Titles signal beginnings of person names. Closed class words, such as determiners, pronouns, and prepositions often signal a boundary. Corporate designators (Inc., Ltd., Corp., etc.) often end a corporation name.

While the number of word-states within each name-class is equal to $|V|$, this “interior” bigram language model is ergodic, *i.e.*, there is a non-zero probability associated with every one of the $|V|^2$ transitions. As a parameterized, trained model, for transitions that were never observed, the model “backs off” to a less-powerful model which allows for the possibility of unknown words.

Training

The model as used for the MUC-7 NE evaluation was trained on a total of approximately 790,000 words of NYT newswire data, annotated with approximately 65,500 named entities. In order to increase the size of our training set beyond the 90,000 words of training of airline crash documents provided by the Government, we selected additional training data from the North American News Text corpus. We annotated full articles before discovering a more effective annotation strategy. Since the test domain was to be similar to the dry-run domain of air crashes, we used the University of Massachusetts INQUERY system to select 2000 articles which were similar to the 200 dry run training and test documents. About half of our training data consisted of full messages; this portion included the 200 messages provided by the Government as well as 319 messages from the 2000 retrieved by INQUERY. The second half of the data consisted of sample sentences selected from the remainder of the 2000 messages with the hope of increasing the variety of training data. This sampling strategy proved more effective than annotating full messages. Improvement in performance as measured on the (dry run) airline crash test set is shown in Figure 5.

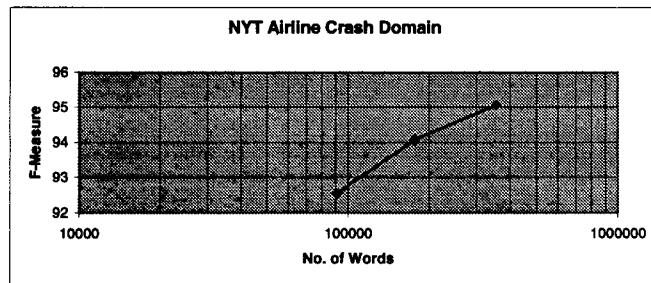


Figure 5: F-Measure Increases With Size of Training Set

IdentiFinder Results under Varying Test and Training Conditions

Our F-measure for the official MUC-7 test, 90.44, is shown as “Text Baseline” in Figure 6. In addition to this baseline condition, we performed some unofficial experiments to measure the accuracy of the system under more difficult conditions. Specifically, we evaluated the system on the test data modified to remove all case information (“Upper Case” in Figure 6), and also on the test data in SNOR (Speech Normalized Orthographic Representation) format (“SNOR” in Figure 6). By converting the text to all upper case characters, information useful for recognizing names in English is removed. Automatically transcribed speech, even with no recognition errors, is harder due to the lack of punctuation, spelling numbers out as words, and upper case in SNOR format.

The degradation in performance from mixed case to all upper case is somewhat greater than that previously observed in similar tests run on generic newswire data (about 2 points). One possible explanation is that case information is more useful in instances where the test domain is different than the domain of the training set. The degradation from all upper case to SNOR is similar to that previously observed.

We also measured the effect of the training set size on the performance of the system in the air crash domain of the dry run. As is to be expected, increasing the amount of training data results in improved system performance.

Figure 5 shows an almost two point increase in F-measure as the training set size was doubled from 91,000 words to 176,000 words. However, the next doubling of the number of words in the training set only resulted in a one point increase

in F-measure. This is most likely due to the fact that as training set size increases, the likelihood of seeing a unique name or construction decreases. Though performance might not have peaked, adding more training data will have a progressively smaller effect since the system will not be seeing many constructions which it has not already seen in previous training.

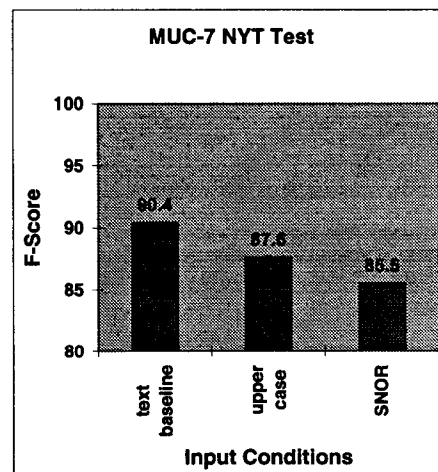


Figure 6: IdentiFinder Named Entity Results

CONCLUSIONS

Throughout its extraction research under the TIPSTER III program, BBN’s goal has been to apply statistical models trained from data in as integrated a fashion as possible. We believe that this approach is fully capable of matching the performance of systems based on rules handwritten by experts, and that it further offers significant advantages in applicability to new problems and new domains, and to degraded input (e.g., from a speech recognizer, from OCR, or from sources less polished than newspaper text).

The SIFT system successfully uses an integrated syntactic/semantic model to extract entities and relations. It employs the Penn Treebank as its source of syntactic information, and thus requires for its training data only the semantic annotation of entities, descriptors, and relationships. Its sentence-level model determines parts of speech, parses, finds names, and identifies semantic relationships in a single, integrated process, with a separate merging model then used to connect information between sentences. Given the current early stage of development of the SIFT system, we believe that significant performance improvements are still possible. We are also interested in measuring performance as a function of training set size, and have begun applying SIFT to the broadcast news domain.

IdentiFinder is BBN's trained system for identifying named entities. Its performance in the MUC-7 evaluation demonstrates the robustness of the learning algorithm used, even when the testing is in a different though similar domain to that of the training material. Further tests also showed its robustness to all upper case input, and input with no punctuation. Our future plans for IdentiFinder include:

- evaluation in the broadcast news domain, which requires speech input in a much broader domain,
- applying IdentiFinder to unsegmented languages, and
- working on performance improvements and improvements in the training process.

ACKNOWLEDGEMENTS

The work reported here was supported in part by the Defense Advanced Research Projects Agency. Technical agents for part of this work were Fort Huachucha and AFRL under contract numbers DABT63-94-C-0062, F30602-97-C-0096, and 4132-BBN-001. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the United States Government.

We appreciate the contributions of the Annotation Group at BBN: Ann Albrect, Elizabeth Arentzen, Rachel Bers, Ada Brunstein, Georgina Garcia, Maia Mesnil, and Hugh Walsh.

We thank Michael Collins of the University of Pennsylvania for his valuable suggestions.

REFERENCES

Bikel, Dan; S. Miller; R. Schwartz; and R. Weischedel. (1997) "NYMBLE: A High-Performance Learning Name-finder." In *Proceedings of the Fifth Conference on Applied Natural Language Processing*, Association for Computational Linguistics, pp. 194-201.

Collins, Michael. (1996) "A New Statistical Parser Based on Bigram Lexical Dependencies." In *Proceedings of the 34th Annual Meeting of the Association for Computational Linguistics*, pp. 184-191.

Collins, Michael. (1997) "Three Generative, Lexicalised Models for Statistical Parsing." In *Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics*, pp. 16-23.

Marcus, M.; B. Santorini; and M. Marcinkiewicz. (1993) "Building a Large Annotated Corpus of English: the Penn Treebank." *Computational Linguistics*, 19(2):313-330.

Goodman, Joshua. (1997) "Global Thresholding and Multiple-Pass Parsing." In *Proceedings of the Second Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, pp. 11-25.

Weischedel, Ralph; Marie Meteer; Richard Schwartz; Lance Ramshaw; and Jeff Palmucci. (1993) "Coping with Ambiguity and Unknown Words through Probabilistic Models." *Computational Linguistics*, 19(2):359-382.