



DESIGN AND PERFORMANCE OF HARC, THE BBN SPOKEN LANGUAGE UNDERSTANDING SYSTEM

*Madeleine Bates, Robert Bobrow, Pascale Fung, Robert Ingria,
Francis Kubala, John Makhoul, Long Nguyen, Richard Schwartz, David Stallard*

BBN Systems and Technologies
Cambridge, MA 02138, USA

ABSTRACT

In this paper we describe the design and performance of a complete spoken language understanding system currently under development at BBN. The system, dubbed HARC (Hear And Respond to Continuous speech), successfully integrates state-of-the-art speech recognition and natural language understanding subsystems. The system has been tested extensively on a restricted airline travel information (ATIS) domain with a vocabulary of over 1000 words. In this application, the system functions as an electronic airline guide, searching a database to answer questions posed by the user.

HARC is implemented in portable, high-level software that runs in real time on today's workstations to support interactive online human-machine dialogs at a very comfortable pace. No special purpose hardware is required other than an A/D converter to digitize the speech.

The system works well for any native speaker of American English and does not require any enrollment data from the users. HARC has shown consistently high performance in formal evaluations on the ATIS domain. At the February 1992 DARPA-sponsored tests administered by the National Institute of Standards and Technology (NIST), the HARC system gave the highest speech recognition and overall understanding performance of all participating systems.

1. OVERVIEW

The BBN HARC spoken language system weds two technologies, speech recognition and natural language understanding, into a deployable human-machine interface. The problem of understanding goal-directed spontaneous speech is harder than recognizing and understanding read text, due to greater variety in the speech and language produced. We have made minor modifications to our speech recognition and understanding methods to deal with these variabilities. The speech recognition uses a novel multipass search strategy that allows great flexibility and efficiency in the application of powerful knowledge sources. The natural language system is based on formal linguistic principles with extensions to deal with speech errors and to make it robust to natural variations in language. The result is a very usable system for domains of moderate complexity.

While the techniques used here are general, the most complete test of the whole system thus far was made using the ATIS corpus, which is briefly described in Section 2. Section 3 describes the techniques used and the results obtained for speech recognition, and Section 4 is devoted to natural language. The methods for combining speech recognition and language understanding, along with results for the combined system are given in Section 5. Finally, in Section 6, we describe a real-time implementation of the system that runs entirely in software on a single workstation.

More details on the specific techniques used, the makeup of the corpus, and the results can be found in the papers presented at the 1992 DARPA Workshop on Speech and Natural Language [1, 2, 3, 4].

2. THE ATIS DOMAIN AND CORPUS

2.1. ATIS

The Air Travel Information Service (ATIS) is a system for getting information about flights. The information contained in the database is similar to that found in the Official Airline Guide (OAG) but is for a small number of cities. The ATIS corpus consists of spoken queries by

a large number of users who were trying to solve travel related problems. The ATIS0 and ATIS1 corpora contain about 4,000 utterances, most of which were read sentences, mostly by speakers with dialects from the southern U.S. The ATIS2 training corpus consists of 10,411 spontaneous utterances from 286 subjects, who were using simulated or real speech understanding systems in order to obtain realistic speech and language. The data originated from 5 collection sites using a variety of strategies for eliciting and capturing spontaneous queries from the subjects [4], with nearly half the data coming from MIT.

About 1,200 utterances were truncated or contained word fragments due to stuttering. Many more contained various nonspeech sounds. There were also frequent long pauses and hesitations in the data. Each sentence in the corpus was classified as class A (self contained meaning), class D (referring to some previous sentence), or class X (impossible to answer for a variety of reasons). The speech recognition system was tested on all three classes, although the results for classes A and D were given more importance. The natural language system and combined speech understanding systems were scored only on classes A and D, although they were presented with all of the test sentences in their original order.

2.2. Formal Evaluation Conditions

The February '92 evaluation test set has data from 37 speakers: 20 female and 17 male. The number of utterances per speaker varied from 5 to 64, but the number of utterances from each of the 5 data-collection sites was carefully balanced. All results given were collected with the Sennheiser microphone (same as the training data). The recognition mode was speaker-independent – the test speakers were not in the training set and every sentence was treated independently.

By committee decision there was no common baseline control condition for the training data or speech grammar to be used for the February '92 ATIS tests. The only constraint was that the single common test set must be used.

3. BYBLOS – SPEECH RECOGNITION

BYBLOS is a state-of-the-art, phonetically-based, continuous speech recognition system that has been under development at BBN for over six years. This system introduced an effective strategy for using context-dependent phonetic hidden Markov models (HMM) and demonstrated their feasibility for large vocabulary, continuous speech applications [5]. Over the years, the core algorithms have been refined primarily on artificial applications using read speech for training and testing. These same basic algorithms, with small extensions, have proven to be remarkably suited to the recognition of completely spontaneous speech produced in a goal-directed task, such as ATIS.

3.1. New Extensions for Spontaneous Speech

Spontaneous queries spoken in a problem-solving dialog exhibit a wide variety of disfluencies. There were three very frequent effects that we attempted to solve – excessively long segments of waveform with no speech, poorly transcribed training utterances, and a variety of nonspeech sounds produced by the user.

When background noise is present, the HMM is not a particularly reliable discriminator of speech vs. silence, and many insertion errors result. We chose to find and truncate long regions of nonspeech with a very reliable energy-based speech detector that can deal with noise bursts near the speech. The speech detector uses several simple adaptive SNR-dependent detection thresholds.

Typically, there are many untranscribed short segments of background silence remaining in the waveforms after truncating the long ones. We found them to be so numerous that they measurably degraded the performance gain usually derived from using cross-word-boundary triphone HMMs. We devised a procedure to automatically mark the missing silence locations by running the recognizer on the training data constrained to the correct word sequence, but allowing optional silence between each word. Then we retrained the model using the output of the recognizer as *corrected* transcriptions.

Spontaneous data from naive speakers exhibits a large number and variety of nonspeech events, such as pause fillers (um's and uh's), throat clearings, coughs, laughter, and heavy breath noise. We attempted to model a dozen broad classes of nonspeech sounds that were both prominent and numerous. However, when we allowed the decoder to find nonspeech models between words, we found that there were more false detections than correct ones. Because our silence model had little difficulty dealing with breath noises, lip smacks, and other noises, our best results were achieved by making the nonspeech models very unlikely in the grammar.

3.2. Forward-Backward N-best Search Strategy

The BYBLOS speech recognition system uses a novel multi-pass search strategy designed to use progressively more detailed models on a correspondingly reduced search space. It produces an ordered list of the *N* top-scoring hypotheses which is then reordered by several detailed knowledge sources. This N-best strategy [6, 7] permits the use of otherwise computationally prohibitive models by greatly reducing the search space to a few ($N=20-100$) word sequences. It has enabled us to use cross-word-boundary triphone models and trigram language models with ease. The N-best list is also a robust interface between speech and natural language that provides a way to recover from speech errors.

We use a 4-pass approach to produce the N-best lists for natural language processing.

1. A forward pass with a bigram grammar and discrete HMM models saves the top word-ending scores and times [8].
2. A fast time-synchronous backward pass produces an initial N-best list using the Word-Dependent N-best algorithm [9].
3. Each of the *N* sentence hypotheses is rescored with cross-word-boundary triphones and semi-continuous density HMMs and the N-best list is reordered.
4. The N-best list is rescored with a trigram grammar and reordered again.

The top choice in the final list constitutes the speech recognition results reported below. Then the entire list is passed to the language understanding component for further reordering and interpretation.

3.3. Training Conditions

Although there were no restrictions on the training data to be used, we used speech data from the ATIS2 subcorpus exclusively to train the parameters of the acoustic model. We did this because we felt that the ATIS2 subset best represented the conditions of the test and because we felt that simply adding more training data to achieve incremental improvements is scientifically uninteresting.

However, we filtered the training data for quality in several ways. We removed from the training any utterances that were marked as truncated, containing a word fragment, or containing rare nonspeech events. Our forward-backward training program also automatically rejects any input that fails to align properly, thereby discarding many sentences with incorrect transcriptions. These steps removed about 1,200 utterances from consideration. Another 600 utterances were removed due to name conflicts between a number of subjects from AT&T and MIT that were assigned identical speaker codes, thus making it difficult to match the speech utterances with the transcriptions.

After holding out 890 utterances as a development test set, we were left with a total of 7670 utterances from 237 speakers for training the HMMs. For statistical language model training we used all available (14,500) sentence texts from ATIS0, ATIS1, and ATIS2 (excluding the development test sentences from the language model training during the development phase).

The lexicon used for recognition was initialized by including all words observed in the complete grammar training texts. Common closed-class words such as days of the week, months, numbers, plane types, etc., were completed by hand. Similarly, we included

derivations (mostly plurals and possessives) of many open-class words in the domain. We also added about 400 concatenated word tokens for commonly occurring sequences such as WASHINGTON_D.C., SAN_FRANCISCO, and D.C.TEN. The final size of the lexicon was 1801 words. For the February '92 evaluation test set only 42 word tokens, covering 35 unique words, were out-of-vocabulary (OOV) for this lexicon. This is only a 0.4% OOV word occurrence rate over the whole test set.

We estimated the parameters of our statistical bigram and trigram grammars using a backing-off procedure similar to Katz [10]. The *n*-grams were computed on word classes in order to share the very sparse training. A total of 1054 semantic classes were defined (most words remained singletons in their class).

3.4. Speech Recognition Results

Table 1 shows the official results for BYBLOS on this evaluation, broken down by utterance class. We also provide the perplexity of the bigram and trigram language models as measured on the evaluation test set (ignoring out-of-vocabulary words).

Sentence Class	# Sents	Bigram Perplex	Trigram Perplex	# Word Errors	% Word Errors
A+D	687	17	12	501	6.2
A+D+X	971	20	15	1015	9.4
A	402	15	10	305	5.8
D	285	20	14	196	7.0
X	284	35	28	514	17.2

Table 1: BYBLOS Official SPREC results on the Feb '92 test set.

The 6.2% word error rate on class A+D sentences and the combined A+D+X error rate of 9.4% were significantly better than any other speech system reporting on this data. The recognition performance was well correlated with the measured perplexities. The trigram language model consistently, but rather modestly, reduced perplexity across all three classes. (We observed that word error was reduced by 40% on classes A+D with the trigram model.) More striking are the differences between the perplexities of the A and D sentences and the Class X sentences (those which are unevaluable with respect to the database).

Performance on the class X utterances is markedly worse than either class A or D utterances. In fact, more than half of the speech errors occur on these utterances. Since these utterances are not evaluable by the natural language component, it does not seem profitable to try to improve the speech performance on these utterances for a spoken language system.

The individual speaker results varied widely from 0.0% word error to 30% error with the median at about 7.5%. The female speakers got significantly better results than the male speakers. Further examination showed that the males in this corpus tended to speak faster and stray from the domain more frequently, resulting in utterances with higher perplexity.

The error rates varied significantly depending on the site where the data was collected. A close examination of the results showed that this was due to two effects: the number of serious spontaneous speech effects, and the amount of training from each site. We found (based on comparison of each error with the detailed transcriptions) that 22% of the word errors were directly caused by spontaneous speech effects. If we removed the errors due to spontaneous effects, then the remaining error rate was found to vary inversely with the square root of the amount of training data, as expected. Thus, the data from MIT, which had the lowest number of spontaneous effects and by far the most training data, resulted in the lowest error rate.

4. DELPHI - NATURAL LANGUAGE UNDERSTANDING

The natural language (NL) component of HARC is the DELPHI system. DELPHI uses a definite clause grammar formalism, augmented by the use of constraint nodes [11] and a labelled argument formalism [3]. Our initial parser used a standard context-free parsing algorithm, extended to handle a unification-based grammar. It was then modified to integrate semantic processing with parsing, so that only semantically coherent structures would be placed in the syntactic chart. The speed and robustness were enhanced by switching to an agenda-based chart-parser, with scheduling depending on the measured statistical likelihood

of grammatical rules [12]. This greatly reduced the search space for the best parse.

The most recent version of DELPHI includes changes to the syntactic and semantic components that maintain the tight syntactic/semantic coupling characteristic of earlier versions, while allowing the system to provide semantic interpretations of input which has no valid global syntactic analysis. This included the development of a "fallback component" [2], in which statistical estimates play an important role. This component allows DELPHI to deal effectively with linguistically ill-formed inputs that are common in spontaneous speech, as well as with the word errors produced by the speech recognizer.

4.1. Parsing as Transduction – Grammatical Relations

The DELPHI parser is not a device for constructing syntactic trees, but an information transducer. Semantic interpretation is a process operating on a set of messages characterizing local "grammatical relations" among phrases, rather than as a recursive tree walk over a globally complete and coherent parse tree. The grammar has been reoriented around local grammatical relations such as deep-structure subject and object, as well as other adjunct-like relations. The goal of the parser is to make these local grammatical relations (which are primarily encoded in ordering and constituency of phrases) as readily available to the semantic interpreter as information explicitly encoded in the words themselves.

From the point of view of a syntactic-semantic transducer, the key point of any grammatical relation is that it licenses a small number of semantic relations between the "meanings" of the related constituents. Sometimes the grammatical relation constrains the semantic relation in ways that cannot be predicted from the semantics of the constituents alone (e.g. Given "John", "Mary", and "kissed", only the grammatical relations or prior world knowledge determine who gave and who received). Other times the grammatical relation simply licenses the only plausible semantic relation (e.g., "John", "hamburger", and "ate"). Finally, in sentences like "John ate the fries but rejected the hamburger", semantics would allow the hamburger to be eaten, but our knowledge of its destiny is mediated by its lack of any grammatical relation to "ate".

Grammatical relations are expressed in the grammar by giving each element of the right hand side of a grammar rule a grammatical relation as a label. A typical rule, in schematic form, is:

(NP ...) → :HEAD (NP ...) :PP-COMP (PP :PREP ...)

which says that a noun phrase followed by a prepositional phrase provides evidence for the relation PP-COMP between the PP and HEAD of the NP.

One of the right-hand elements must be labeled the "head" of the rule, and is the initial source of information about the semantic and syntactic "binding state" which controls whether other elements of the right-hand side can "bind" to the head via their labeled relation.

This view made it possible to both decrease the number of grammar rules (from 1143 to 453) and increase syntactic coverage. Most attachments can be modelled by simple binary adjunction, and since the details of the syntactic tree structure are not central to a transducer, each adjunct can be seen as being "logically attached" to the "head" of the constituent.

In effect, we factored single grammar rules (that produced ordered sequences of constituents) into smaller binary adjunction rules that can be combined together in various ways; ordering constraints on these adjuncts are provided by lexical semantic well-formedness rules. For example, rather than using subcategorization features to name sets of categories that appear together as complements of a verb, we have defined approximately 15 verb phrase rules that list the possible constituents that may appear as complements to a verb. These may embed within each other freely, so long as the results are semantically interpretable by the head. Similar recursive binary complement and adjunct rules are provided for noun phrases. This recursive scheme allows the adjunction rules of the grammar to be combined together in novel ways, governed by the lexical semantics of individual words. The grammar writer does not need to foresee all possible combinations.

4.2. "Binding rules" – the Semantics of Grammatical Relations

The interface between parsing and semantics is a dynamic process structured as two coroutines in a cascade. The input to the semantic interpreter is a sequence of messages, each requesting the semantic "binding" of some constituent to a head. A set of "binding rules" for each grammatical relation licenses the binding of a constituent to a head via that relation by specifying the semantic implications of binding. These rules specify features of the semantic structure of the head and bound constituent that must be true for binding to take place, and may also specify syntactic requirements. Rules may also allow certain semantic roles (such as time specification) to have multiple fillers, while other roles may allow just one filler.

As adjuncts are added to a structure, the binding list is conditionally extended as long as semantic coherence is maintained. When a constituent is syntactically complete (i.e., no more adjuncts are to be added), DELPHI evaluates rules that check for semantic completeness and produce an "interpretation" of the constituent.

4.3. Robustness Based on Statistics and Semantics

Unfortunately, simply having a transduction system with semantics based on grammatical relations does not deal directly with the key issue of robustness – the ability to make sense of an input even if it cannot be assigned a well-formed global syntactic analysis. The difficulty with standard syntactic techniques is that local syntactic evidence is not enough to accurately determine grammatical relations. A NP followed by a verb may be the subject of that verb ("John flew to Boston") or may be unrelated ("The man I introduced to John flew to Boston"). The standard solution is to find a global parse, which provides the necessary confirming evidence for the local relations it contains.

In DELPHI we view standard global parsing as merely one way to obtain evidence for the existence of the grammatical relations in an input string. DELPHI's strategy is based on two other sources of information. DELPHI applies semantic constraints incrementally during the parsing process, so that only semantically coherent grammatical relations are considered. Additionally, DELPHI has statistical information on the likelihood of various word senses, grammatical rules, and grammatical-semantic transductions. Thus DELPHI can rule out many locally possible grammatical relations on the basis of semantic incoherence, and can rank alternative local structures on the basis of empirically measured probabilities. The net result is that even in the absence of a global parse, DELPHI can quickly and reliably produce the most probable local grammatical relations and semantic content of various fragments.

DELPHI first attempts to obtain a complete syntactic analysis of its input, using its agenda-based best-first parsing algorithm. If it is unable to do this, it uses the parser in a fragment-production mode, which produces the most probable structure for an initial segment of the input, then restarts the parser in a top down mode on the first element of the unparsed string whose lexical category provides a reasonable anchor for top-down prediction. This process is repeated until the entire input is spanned with fragments. Experiments have shown that the combination of statistical evaluation and semantic constraints produces chunks of the input that are very useful for interpretation by non-syntactically-driven strategies.

4.4. Advantages of This Approach

The separation of syntactic grammar rules from semantic binding and completion rules greatly facilitates fragment parsing. While it allows syntax and semantics to be strongly coupled in terms of processing (parsing and semantic interpretation) it allows them to be essentially decoupled in terms of notation. This makes the grammar and the semantics considerably easier to modify and maintain.

We believe, however, that in the long term the most important advantage is that separating syntactic rules from semantic binding leads us to a new kind of language model, based on grammatical relations, in which syntactic, semantic and lexical knowledge can be acquired by largely automatic means. We view the role of the grammar as codifying the way that tree structure provides evidence for grammatical relations. The separation between rule types will allow us for the first time to consider the effect of grammatical relations on meaning, independently of the way that evidence for these relations is produced by the parser.

One effect of this approach is to make it possible to use a hypoth-

esized semantic interpretation of a set of tree fragments to generate a new syntactic rule. Thus, in normal operation, the primary evidence for a grammatical relation is the result of actually parsing part of an input. However, since grammatical relations between constituents entail semantic relations, if we can make an estimate of the likelihood of certain semantic relations based on domain knowledge, pragmatics, and task models, etc., it is in principle possible to use abductive reasoning to suggest likely grammatical relations, and thereby automatically propose new grammar rules.

5. COMBINED SPOKEN LANGUAGE SYSTEM

The basic interface between BYBLOS and DELPHI in HARC is the N-best list. In the most basic strategy, we allowed the NL component to search arbitrarily far down the N-best list until it either found a hypothesis that produced a database retrieval or reached the end of the N-best list. However, we noticed that while it was beneficial for NL to look beyond the first hypothesis in an N-best list, the answers obtained by NL from speech output tended to degrade the further down in the N-best list they were obtained.

For this evaluation, we optimized both the depth of the search that NL performed on the N-best output of speech and how we used a number of fall-back strategies for NL text processing [2]. We found that, given the current performance of all the components, the optimal number of hypotheses to consider was $N=5$. Furthermore, we found that rather than applying the fall-back mechanism to each of these hypotheses in turn, it was better to make one pass through the N-best hypotheses using the full parsing strategy, and then, if no sentences were accepted, make another pass using the fall-back strategy.

5.1. Results

In Table 2 we show the official performance on the February '92 evaluation data as calculated by NIST. The percent correct and the weighted error rate is given for the DELPHI system operating on the transcribed text (NL) and for the combined HARC system (SLS). The results are shown for classes A+D combined and separately. The weighted error measure weights incorrect answers twice as much as no answer.

weighted error (WE) = $2 \times \% \text{incorrect answers} + \% \text{no answers}$

Corpus	NL Cor	NL WE	SLS Cor	SLS WE
A+D	76.7	33.9	71.8	43.7
A	80.1	26.4	74.9	35.8
D	71.9	44.6	67.4	54.7

Table 2: %Correct and Weighted error on the February '92 test set.

The weighted error on context-dependent sentences (D) is clearly higher than on sentences that stand alone (A). This higher error rate results from two phenomena. First, it is often difficult to resolve references correctly and to know how much of the previous constraints are to be kept. Second, in order to understand a context-dependent sentence correctly, we must correctly understand at least two sentences, which is less likely.

The weighted error from speech input is consistently 10% higher than from text. However, the difference is lower than might be expected. Even though the BYBLOS system misrecognized at least one word in 35% of the utterances, the DELPHI system was able to recover from most of these errors through the use of the N-best list and fallback processing.

The official SLS result for HARC was a weighted error of 43.7%. This was the best overall result for a spoken language system in the February '92 DARPA evaluation.

6. REAL-TIME IMPLEMENTATION

A real-time demonstration of the entire spoken language system described above has been implemented. The speech recognition was performed using BBN RUBYTM, a commercially available product for continuous speech recognition of medium-sized vocabularies (about 1,000 words). RUBY has essentially the same recognition accuracy as BYBLOS but can run in real-time entirely in software on a workstation with a built-in A/D converter (e.g., SGI Indigo or SUN Sparc II) without any additional hardware.

The speech recognition displays an initial answer as soon as the user stops speaking, and a refined (rescored) answer within 1–2 seconds. The natural language system chooses one of the N-best answers, interprets it, and computes and displays the answers, along with a paraphrase of the query so the user can verify what question the system answered. The total response cycle is typically 3–4 seconds, making the system feel extremely responsive. The error rates for knowledgeable interactive users appears to be much lower than those reported above for naive noninteractive users.

7. SUMMARY

We have described the HARC spoken language understanding system. HARC consists of a modular integration of the BYBLOS speech recognition system with the DELPHI natural language understanding system. The two components are integrated using the N-best paradigm. We have shown that this paradigm is a modular, efficient, and effective method for integrating speech recognition and language understanding components. In addition, the N-best strategy was shown to be useful within the speech recognition system as a means of applying expensive knowledge sources, such as cross-word acoustic models and trigram language models.

The official BYBLOS speech recognition results for the February '92 DARPA evaluation were 6.2% word error for the Class A+D subset of the test and 9.4% overall. Both of these results were significantly better than any other speech system tested. The official SLS result for HARC was a weighted error of 43.7%, which was the best overall result for a spoken language system.

Finally, the entire system has been implemented to run in real time on a standard workstation without the need for any additional hardware.

Acknowledgement

This work was supported by the Defense Advanced Research Projects Agency and monitored by the Office of Naval Research under Contract Nos. N00014-89-C-0008, N00014-91-C-0115, and N00014-92-C-0035.

REFERENCES

- [1] F. Kubala, C. Barry, M. Bates, R. Bobrow, P. Fung, R. Ingria, J. Makhoul, L. Nguyen, R. Schwartz, D. Stallard, "BBN BYBLOS and HARC February 1992 ATIS Benchmark Results", *Proc. of the DARPA Speech and Natural Language Workshop*, Morgan Kaufmann Publishers, Feb. 1992.
- [2] Bobrow R., D. Stallard, "Fragment Processing in the DELPHI System", *Proc. of the DARPA Speech and Natural Language Workshop*, Morgan Kaufmann Publishers, Feb. 1992.
- [3] Bobrow, R., R. Ingria, and D. Stallard, "Syntactic/Semantic Coupling in the DELPHI System", *Proc. of the DARPA Speech and Natural Language Workshop*, Morgan Kaufmann Publishers, Feb. 1992.
- [4] MADCOW, "Multi-Site Data Collection for a Spoken Language Corpus", *Proc. of the DARPA Speech and Natural Language Workshop*, Morgan Kaufmann Publishers, Feb. 1992.
- [5] Chow, Y., M. Dunham, O. Kimball, M. Krasner, G.F. Kubala, J. Makhoul, P. Price, S. Roucos, and R. Schwartz (1987) "BYBLOS: The BBN Continuous Speech Recognition System," *IEEE ICASSP-87*, pp. 89-92
- [6] Chow, Y-L. and R.M. Schwartz, "The N-Best Algorithm: An Efficient Procedure for Finding Top N Sentence Hypotheses", *Proc. of the DARPA Speech and Natural Language Workshop*, Morgan Kaufmann Publishers, Oct. 1989, pp. 199–202. Also in ICASSP-90, April 1990, Albuquerque S2.12, pp. 81-84.
- [7] Schwartz, R., S. Austin, Kubala, F., and J. Makhoul, "New Uses for the N-Best Sentence Hypotheses Within the BYBLOS Speech Recognition System", ICASSP92, San Francisco, CA.
- [8] Austin, S., Schwartz, R., and P. Placeway, "The Forward-Backward Search Algorithm", ICASSP91, Toronto, Canada, pp. 697-700.
- [9] Schwartz, R. and S. Austin, "A Comparison Of Several Approximate Algorithms for Finding Multiple (N-Best) Sentence Hypotheses", ICASSP91, Toronto, Canada, pp. 701-704.
- [10] Katz, S., "Estimation of Probabilities from Sparse Data for the Language Model Component of a Speech Recognizer", *IEEE Trans. on ASSP*, Mar. 1987, Vol. 35, No. 3.
- [11] Stallard, D., "Unification-Based Semantic Interpretation in the BBN Spoken Language System", *Proc. of the DARPA Speech and Natural Language Workshop*, Morgan Kaufmann, Publishers, Oct. 1989, pp. 39–46.
- [12] Bobrow, R., "Statistical Agenda Parsing", *Proc. of the DARPA Speech and Natural Language Workshop*, Morgan Kaufmann Publishers, Feb. 1991, pp. 222–224.