An Efficient Two Stage Approach to Robust Language Interpretation

Carolyn Penstein Rosé
Carnegie Mellon University
Baker Hall 135F
Pittsburgh, PA 15213
cprose@cs.cmu.edu

1 Introduction

The most basic task of a natural language interface is to map the user's utterance onto some meaning representation which can then be used for further processing. The three biggest challenges which continue to stand in the way of accomplishing even this most basic task are extragrammaticality, ambiguity, and recognition errors. The system presented here, ROSE[1]: RObustness with Structural Evolution, repairs extragrammatical input in two stages. The first stage, Repair Hypothesis Formation, is responsible for assembling a set of hypotheses about the meaning of the ungrammatical utterance. This stage is itself divided into two steps, Partial Parsing and Combination. In the Combination step, the fragments from a partial parse are assembled into a set of meaning representation hypotheses. In ROSE's second stage, Interaction with the user, the system generates a set of queries and then uses the user's answers to these queries to narrow down to a single best meaning representation hypothesis.

2 Comparison to Alternative Approaches

Rather than placing the full burden of robustness on the parser itself, I argue that it is more economical for Partial Parsing and Combination to be separate steps in the Hypothesis Formation stage. Efforts towards solving the problem of extragrammaticality have primarily been in the direction of building flexible parsers. In principle, Minimum Distance Parsers (Lehman, 1989; Hipp, 1992) have the greatest flexibility. They fit an extragrammatical sentence to the parsing grammar through a series of insertions, deletions, and transpositions. Since any string can be mapped onto any other string through a series of insertions, deletions, and transpositions, this approach makes it possible to perform any desired repair. The underlying assumption behind the MDP approach is that the analysis of the string which deviates the least from the input string is most likely

to be the best analysis. Thus, Minimum Distance Parsing appears to be a reasonable approach.

In practice, however, Minimum Distance Parsing has only been used successfully in very small and limited domains. Lehman's core grammar, described in (Lehman, 1989), has on the order of 300 rules, and all of the inputs to her system can be assumed to be commands to a calendar program. Hipp's Circuit Fix-It Shop system, described in (Hipp, 1992), has a vocabulary of only 125 words and a grammar size of only 500 rules. Flexible parsing algorithms introduce a great deal of extra ambiguity. This in turn may deem certain approaches impractical for systems of realistic scale. Therefore, an important question one must ask is whether the MDP approach can scale up to a larger system and/or domain.

An example of a less powerful parsing algorithm is Lavie's GLR* skipping parser described in (Lavie, 1995). This parser is capable of skipping over any portion of an input utterance that cannot be incorporated into a grammatical analysis and recover the analysis of the largest grammatical subset of the utterance. Partial analyses for skipped portions of the utterance are also returned by the parser. Thus, whereas MDP considers insertions and transpositions in addition to deletions, GLR* only considers deletions. The weakness of this and other partial parsing approaches (Abney, 1997; Nord, 1997; Srinivas et al., 1997; Federici, Montemagni, and Pirrelli, 1997) is that part of the original meaning of the utterance may be thrown away with the portion(s) of the utterance which are skipped if only the analysis for the largest subset is returned, or part of the analysis will be missing if the parser only attempts to build a partial parse. These less powerful algorithms trade coverage for speed. The idea is to introduce enough flexibility to gain an acceptable level of coverage at an acceptable computational expense.

The goal behind ROSE and other two stage approaches (Ehrlich and Hanrieder, 1997; Danieli and Gerbino, 1995) is to increase the coverage possible at a reasonable computational cost by introducing a post-processing repair stage, which constructs a complete meaning representation out of the frag-

---

[1] ROSE is pronounced Rosé, like the wine.

ments of a partial parse. Since the input to the second stage is a collection of partial parses, the additional flexibility that is introduced at this second stage can be channeled just to the part of the analysis that the parser does not have enough knowledge to handle straightforwardly. This is unlike the MDP approach, where the full amount of flexibility is unnecessarily applied to every part of the analysis. Therefore, this two stage process is more efficient since the first stage is highly constrained by the grammar and the results of this first stage are then used to constrain the search in the second stage. Additionally, in cases where the limited flexibility parser is sufficient, the second stage can be entirely bypassed, yielding an even greater savings in time.

3 A Simple Example

The heart of the ROSE approach is the Combination Mechanism, a genetic programming (Koza, 1992; Koza, 1994) environment in which programs are evolved which combine the fragments of a partial parse into a complete meaning representation structure. I present a simple example in Figure 1 for the sake of clarity. This should not be taken to be an indication of the full potential of this approach.

Chunks:

1. Thursday
((frame *simple-time)
 (day-of-week thursday))

2. I am out
((frame *busy)
 (who ((frame *i))))

Ideal Repair Hypothesis:

(my-comb
 ((frame *busy) (who ((frame *i))))
 ((frame *simple-time) (day-of-week thursday))
 when)

An Alternative Repair Hypothesis:

(my-comb
 ((frame *busy) (who ((frame *i))))
 ((frame *simple-time) (day-of-week thursday))
 why)

Result of Ideal Hypothesis:

((frame *busy)
 (who ((frame *i)))
 (when ((frame *simple-time)
        (day-of-week thursday))))

Figure 1: Combination Example

The ideal repair hypothesis for this example is one

that specifies that the temporal expression should be inserted into the when slot in the *busy frame. Other hypotheses are also evolved and tested as the genetic programming system runs, such as the alternative example included in Figure 1. A fitness function ranks hypotheses, narrowing down on a small set. The final result is selected through interaction with the user.

References

Abney, S. 1997. Partial parsing via finite-state cascades. In *Proceedings of the Eight European Summer School In Logic, Language and Information, Prague, Czech Republic.*

Danieli, M. and E. Gerbino. 1995. Metrics for evaluating dialogue strategies in a spoken language system. In *Working Notes of the AAAI Spring Symposium on Empirical Methods in Discourse Interpretation and Generation.*

Ehrlich, U. and G. Hanrieder. 1997. Robust speech parsing. In *Proceedings of the Eight European Summer School In Logic, Language and Information, Prague, Czech Republic.*

Federici, S., S. Montemagni, and V. Pirrelli. 1997. Shallow parsing and text chunking: a view on underspecification in syntax. In *Proceedings of the Eight European Summer School In Logic, Language and Information, Prague, Czech Republic.*

Hipp, D. R. 1992. *Design and Development of Spoken Natural-Language Dialog Parsing Systems.* Ph.D. thesis, Dept. of Computer Science, Duke University.

Koza, J. 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection.* MIT Press.

Koza, J. 1994. *Genetic Programming II.* MIT Press.

Lavie, A. 1995. *A Grammar Based Robust Parser For Spontaneous Speech.* Ph.D. thesis, School of Computer Science, Carnegie Mellon University.

Lehman, J. F. 1989. *Adaptive Parsing: Self-Extending Natural Language Interfaces.* Ph.D. thesis, School of Computer Science, Carnegie Mellon University. CMU-CS-89-191.

Nord, G. Van. 1997. Robist parsing with the head-corner parser. In *Proceedings of the Eight European Summer School In Logic, Language and Information, Prague, Czech Republic.*

Srinivas, B., C. Doran, B. Hockey, and A. Joshi. 1997. An approach to robust partial parsing and evaluation metrics. In *Proceedings of the Eight European Summer School In Logic, Language and Information, Prague, Czech Republic.*