Tagging English by Path Voting Constraints

Gökhan Tür and Kemal Oflazer

Department of Computer Engineering and Information Science
Bilkent University, Bilkent, Ankara, TR-06533, TURKEY
{tur,ko}@cs.bilkent.edu.tr

Abstract: We describe a constraint-based tagging approach where individual constraint rules vote on sequences of matching tokens and tags. Disambiguation of all tokens in a sentence is performed at the very end by selecting tags that appear on the path that receives the highest vote. This constraint application paradigm makes the outcome of the disambiguation *independent* of the rule sequence, and hence relieves the rule developer from worrying about potentially conflicting rule sequencing. The approach can also combine statistically and manually obtained constraints, and incorporate negative constraint rules to rule out certain patterns. We have applied this approach to tagging English text from the Wall Street Journal and the Brown Corpora. Our results from the Wall Street Journal Corpus indicate that with 400 statistically derived constraint rules and about 800 hand-crafted constraint rules, we can attain an *average accuracy of 97.89%* on the training corpus and an *average accuracy of 97.50%* on the testing corpus. We can also relax the single tag per token limitation and allow ambiguous tagging which lets us trade recall and precision.

1 Introduction

Part-of-speech tagging is one of the preliminary steps in many natural language processing systems in which the proper part-of-speech tag of the tokens comprising the sentences are disambiguated using either statistical or symbolic local contextual information. Tagging systems have used either a statistical approach where a large corpora is employed to train a probabilistic model which then is used to tag unseen text, (e.g., Church (1988), Cutting et al. (1992), DeRose (1988)), or a constraint-based approach which employs a large number of hand-crafted linguistic constraints that are used to eliminate

impossible sequences or morphological parses for a given word in a given context, recently most prominently exemplified by the Constraint Grammar work (Karlsson et al., 1995; Voutilainen, 1995b; Voutilainen et al., 1992; Voutilainen and Tapanainen, 1993). Brill (1992; 1994; 1995) has presented a transformation-based learning approach.

This paper extends a novel approach to constraint-based tagging first applied for Turkish (Oflazer and Tür, 1997), which relieves the rule developer from worrying about conflicting rule ordering requirements and constraints. The approach depends on assigning votes to constraints via statistical and/or manual means, and then letting constraints vote on matching sequences on tokens, as depicted in Figure 1. This approach does not reflect the outcome of matching constraints to the set of morphological parses immediately as usually done in constraint-based systems. Only after all applicable rules are applied to a sentence, tokens are disambiguated in parallel. Thus, the outcome of the rule applications is *independent* of the order of rule applications.

Figure 1: Voting Constraint Rules

Figure 2: Representing sentences with a directed acyclic graph

2 Tagging by Path Voting Constraints

We assume that sentences are delineated and that each token is assigned all possible tags by a lexicon or by a morphological analyzer. We represent each sentence as a standard chart using a directed acyclic graph where nodes represent token boundaries and arcs are labeled with ambiguous interpretations of tokens. For instance, the sentence *I can can the can.* would be represented as shown in Figure 2, where bold arcs denote the correct tags.

We describe constraints on token sequences using rules of the sort $R = (C_1, C_2, \dots, C_n; V)$, where the C_i are, in general, feature constraints on a sequence of the ambiguous parses, and V is an integer denoting the vote of the rule. For English, the features that we use are: (1) **LEX**: the lexical form, and (2) **TAG**: the tag. It is certainly possible to extend the set of features used, by including features such as initial letter capitalization, any derivational information, etc. (see (Ofazer and Tür, 1997)). For instance, $([TAG=MD], [TAG=RB], [TAG=VB]; 100)$ is a rule with a high vote to promote modal followed by a verb with an intervening adverb. The rule $([TAG=DT, LEX=that], [TAG=NNS]; -100)$ demotes a singular determiner reading of *that* before a plural noun, while $([TAG=DT, LEX=each], [TAG=JJ, LEX=other]; 100)$ is a rule with a high vote that captures a collocation (Santorini, 1995).

The constraints apply to a sentence in the following manner: Assume for a moment that all possible paths from the start node to the end node of a sentence graph are explicitly enumerated, and that after the enumeration, each path is augmented by a vote component. For each path at hand, we apply each constraint to all possible sequences of token parses. Let $R = (C_1, C_2, \dots, C_n; V)$ be a constraint and let $w_i, w_{i+1}, \dots, w_{i+n-1}$ be a sequence of token parses labeling sequential arcs of the path. We

say rule R matches this sequence of parses, if $w_j, i \leq j \leq i + n - 1$ is subsumed by the corresponding constraint C_{j-i+1} . When such a match occurs, the vote of the path is incremented by V . When all constraints are applied to all possible sequences in all paths, we select the path with the maximum vote. If there are multiple paths with the same maximum vote, the tokens whose parses are different in those paths are assumed to be left ambiguous.

Given that each token has on the average more than 2 possible tags, the procedural description above is very inefficient for all but very short sentences. However, the observation that our constraints are localized to a window of a small number of tokens (say at most 5 tokens in a sequence), suggests a more efficient scheme originally used by Church (1988). Assume our constraint windows are allowed to look at a window of at most size k sequential parses. Let us take the first k tokens of a sentence and generate all possible paths of k arcs (spanning $k + 1$ nodes), and apply all constraints to these “short” paths. Now, if we discard the first token and consider the $(k + 1)^{st}$ token, we need to consider and extend only those paths that have accumulated the maximum vote among the paths whose last $k - 1$ parses are the same. The reason is that since the first token is now out of the context window, it can not influence the application of any rules. Hence only the highest scoring (partial) paths need to be extended, as lower scoring paths can not later accumulate votes to surpass the current highest scoring paths.

In Figure 3 we describe the procedure in a more formal way where $w_1, w_2, \dots, w_s$ denotes a sequence of tokens in a sentence, $amb(w_i)$ denotes the number of ambiguous tags for token w_i , and k denotes the maximum context window size (determined at run time).

1. $P = \{ \text{all } \prod_{j=1}^{k-1} \text{amb}(w_j) \text{ paths of the first } k-1 \text{ tokens} \}$
2. $i = k$
3. **while** $i \leq s$
4. **begin**
 - 4.1) Create $\text{amb}(w_i)$ copies of each path in P and extend each such copy with one of the distinct tags for token w_i .
 - 4.2) Apply all constraints to the last k tokens of every path in P , updating path votes accordingly.
 - 4.3) Remove from P any path p if there is some other path p' such that $\text{vote}(p') > \text{vote}(p)$ and the last $k-1$ tags of path p are same as the last $k-1$ tags of p' .
 - 4.4) $i = i + 1$
- end**

Figure 3: Procedure for fast constraint application

3 Results from Tagging English

We evaluated our approach using 11-fold cross validation on the Wall Street Journal Corpus and 10-fold cross validation on a portion of the Brown Corpus from the Penn Treebank CD. We used two classes of constraints: (i) we extracted a set of tag k -grams from a training corpus and used them as constraint rules with votes assigned as described below, and (ii) we hand-crafted a set rules mainly incorporating negative constraints (demoting impossible or unlikely situations), or *lexicalized* positive constraints. These were constructed by observing the failures of the statistical constraints on the training corpus.

Rules derived from the training corpus

For the statistical constraint rules, we extract tag k -grams from the tagged training corpus for $k = 2$, and $k = 3$. For each tag k -gram, we compute a vote which is essentially very similar to the rule strength used by Tzoukermann et al. (1995) except that we do not use their notion of genotypes exactly in the same way. Given a tag k -gram $t_1, t_2, \dots, t_k$, let $n = \text{count}(t_1 \in \text{Ttags}(w_i), t_2 \in \text{Ttags}(w_{i+1}), \dots, t_k \in \text{Ttags}(w_{i+k-1}))$ for all possible i 's in the training corpus, be the number of possible places the tags sequence *can* possibly occur. $\text{footnoteTtags}(w_i)$ is the set of tags

associated with the token w_i . Let f be the number of times the tag sequence $t_1, t_2, \dots, t_k$ actually occurs in the tagged text, that is, $f = \text{count}(t_1, t_2, \dots, t_k)$. We smooth f/n by defining $p = \frac{f+0.5}{n+1}$ so that neither p nor $1-p$ is zero. The uncertainty of p is then given as $\sqrt{p(1-p)/n}$ (Tzoukermann et al., 1995). We then computed the vote for this k -gram as

$$\text{Vote}(t_1, t_2, \dots, t_k) = (p - \sqrt{p(1-p)/n}) * 100.$$

This formulation thus gives high votes to k -grams which are selected most of the time they are “selectable.” And, among the k -grams which are equally good (same f/n), those with a higher n (hence less uncertainty) are given higher votes.

After extracting the k -grams as described above for $k = 2$ and $k = 3$, we ordered each group by decreasing votes and conducted an initial set of experiments to select a small group of constraints performing satisfactorily. We selected the first 200 (with highest votes) of the 2-gram and the first 200 of the 3-gram constraints, as the set of statistical constraints. It should be noted that the constraints obtained this way are purely constraints on tag sequences and do not use any lexical or genotype information.

Hand-crafted rules In addition to these statistical constraint rules, we introduced 824 hand-crafted constraint rules. Most of the hand-crafted constraints imposed negative constraints (with large negative votes) to rule out certain tag sequences that we encountered in the Wall Street Journal Corpus. Another set of rules were lexicalized rules involving the tokens as well as the tags. A third set of rules for idiomatic constructs and collocations was also used. The votes for negative and positive hand-crafted constraints are selected to override any vote a statistical constraint may have.

Initial Votes To reflect the impact of lexical frequencies we initialize the total vote of each path with the sum of the lexical votes for the token and tag combinations on it. These lexical votes for the parse $t_{i,j}$ of token w_i are obtained from the training corpus in the usual way, i.e., as $\text{count}(w_i, t_{i,j})/\text{count}(w_i)$, and then are normalized to between 0 and 100.

Experiments on WSJ and Brown Corpora We tested our approach on two English Corpora

from the Penn Treebank CD. We divided a 5500 sentence portion of the Wall Street Journal Corpus into 11 different sets of training texts (with about 118,500 words on the average), and corresponding testing texts (with about 11,800 words on the average), and then tagged these texts using the statistical rules and hand-crafted constraints. The hand-crafted rules were obtained from only one of the training text portions, and not from all, but for each experiment the 400 statistical rules were obtained from the respective training set.

We also performed a similar experiment with a portion of the Brown Corpus. We used 4000 sentences (about 100,000 words) with 10-fold cross validation. Again we extracted the statistical rules from the respective training sets, but the hand-crafted rules were the ones developed from the Wall Street Journal training set. For each case we measured the *accuracy* by counting the correctly disambiguated tokens. The manual rules used for Brown Corpus were the rules derived from Wall Street Journal data. The results of these experiments are shown in Table 1.

Const. Set	WSJ		Brown	
	Tra. Acc.	Test Acc.	Tra. Acc.	Test Acc.
1	95.59	94.54	95.75	94.25
1+2	96.47	95.68	96.78	95.76
1+3	96.39	95.37	96.50	95.10
1+2+3	96.66	95.96	96.91	96.02
1+4	97.21	96.70	96.27	95.53
1+2+4	97.85	97.43	97.13	96.51
1+3+4	97.60	97.08	96.80	96.09
1+2+3+4	97.89	97.50	97.18	96.67

(1) Lexical Votes (2) 200 2-grams
(3) 200 3-grams (4) 824 Manual Constr.

Table 1: Results from tagging the WSJ and Brown Corpora.

We feel that the results in the last row of Table 1 are quite satisfactory and warrant further extensive investigation. On the Wall Street Journal Corpus, our tagging approach is on par or even better than stochastic taggers making closed vocabulary assumption. Weischedel et al. (1993) report a 96.7% accuracy with 1,000,000 words of training corpus. The performance of

p	Test Set		
	Recall	Precision	Ambiguity
0.99	97.94	96.70	1.012
0.98	98.27	95.29	1.031
0.97	98.48	93.63	1.052
0.96	98.65	91.63	1.076
0.95	98.78	90.21	1.095
0.94	98.98	88.92	1.113
0.93	99.03	88.05	1.124
0.92	99.10	87.19	1.136
0.91	99.13	86.68	1.143

Table 2: Recall and precision results on a WSJ test set with some tokens left ambiguous

our system with Brown corpus is very close to that of Brill's transformation-based tagger, which can reach 97.2% accuracy with closed vocabulary assumption and 96.5% accuracy with open vocabulary assumption with no ambiguity (Brill, 1995). Our tagging speed is also quite high. With over 1000 constraint rules (longest spanning 5 tokens) loaded, we can tag at about 1600 tokens/sec on a Ultrasparc 140, or a Pentium 200.

It is also possible for our approach to allow for some ambiguity. In the procedure given earlier, in line 4.3, if one selects all (partial) paths whose accumulated vote is within p ($0 < p \leq 1$) of the (partial) path with the largest vote, then a certain amount of ambiguity can be introduced, at the expense of a slowdown in tagging speed and an increase in memory requirements. In such a case, instead of accuracy, one needs to use *ambiguity*, *recall*, and *precision* (Voutilainen, 1995a). Table 2 presents the recall, precision and ambiguity results from tagging one of the Wall Street Journal test sets using the same set of constraints but with p ranging from 0.91 to 0.99. These compare quite favorably with the k-best results of Brill(1995), but reduction in tagging speed is quite noticeable, especially for lower p 's. Any improvements in single tag per token tagging (by additional hand crafted constraints) will certainly be reflected to these results also.

4 Conclusions

We have presented an approach to constraint-based tagging that relies on constraint rules vot-

ing on sequences of tokens and tags. This approach can combine both statistically and manually derived constraints, and relieves the rule developer from worrying about rule ordering, as removal of tags is not immediately committed but only after all rules have a say. Using positive or negative votes, we can promote meaningful sequences of tags or collocations, or demote impossible sequences. Our approach is quite general and is applicable to any language. Our results from the Wall Street Journal Corpus indicate that with 400 statistically derived constraint rules and about 800 hand-crafted constraint rules, we can attain an *average accuracy of 97.89%* on the training corpus and an *average accuracy of 97.50%* on the testing corpus. Our future work involves extending to open vocabulary case and evaluating unknown word performance.

5 Acknowledgments

A portion of the first author's work was done while he was visiting Johns Hopkins University, Department of Computer Science with a NATO Visiting Student Scholarship. This research was in part supported by a NATO Science for Stability Program Project Grant – TU-LANGUAGE.

References

- Eric Brill. 1992. A simple-rule based part-of-speech tagger. In *Proceedings of the Third Conference on Applied Natural Language Processing*, Trento, Italy.
- Eric Brill. 1994. Some advances in rule-based part of speech tagging. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*, Seattle, Washinton.
- Eric Brill. 1995. Transformation-based error-driven learning and natural language processing: A case study in part-of-speech tagging. *Computational Linguistics*, 21(4):543–566, December.
- Kenneth W. Church. 1988. A stochastic parts program and a noun phrase parser for unrestricted text. In *Proceedings of the Second Conference on Applied Natural Language Processing*, Austin, Texas.
- Doug Cutting, Julian Kupiec, Jan Pedersen, and Penelope Sibun. 1992. A practical part-of-speech tagger. In *Proceedings of the Third Conference on Applied Natural Language Processing*, Trento, Italy.
- Steven J. DeRose. 1988. Grammatical category disambiguation by statistical optimization. *Computational Linguistics*, 14(1):31–39.
- Fred Karlsson, Atro Voutilainen, Juha Heikkilä, and Arto Anttila. 1995. *Constraint Grammar-A Language-Independent System for Parsing Unrestricted Text*. Mouton de Gruyter.
- Kemal Oflazer and Gökhan Tür. 1997. Morphological disambiguation by voting constraints. In *Proceedings of ACL'97/EACL'97, The 35th Annual Meeting of the Association for Computational Linguistics*, June.
- Beatrice Santorini. 1995. Part-of-speech tagging guidelines fro the penn treebank project. Available at <http://www ldc.upenn.edu/>. 3rd Revision, 2nd Printing.
- Evelyne Tzoukermann, Dragomir R. Radev, and William A. Gale. 1995. Combining linguistic knowledge and statistical learning in french part-of-speech tagging. In *Proceedings of the ACL SIGDAT Workshop From Texts to Tags: Issues in Multilingual Language Analysis*, pages 51–57.
- Atro Voutilainen and Pasi Tapanainen. 1993. Ambiguity resolution in a reductionistic parser. In *Proceedings of EACL'93*, Utrecht, Holland.
- Atro Voutilainen, Juha Heikkilä, and Arto Anttila. 1992. *Constraint Grammar of English*. University of Helsinki.
- Atro Voutilainen. 1995a. Morphological disambiguation. In Fred Karlsson, Atro Voutilainen, Juha Heikkilä, and Arto Anttila, editors, *Constraint Grammar-A Language-Independent System for Parsing Unrestricted Text*, chapter 5. Mouton de Gruyter.
- Atro Voutilainen. 1995b. A syntax-based part-of-speech analyzer. In *Proceedings of the Seventh Conference of the European Chapter of the Association of Computational Linguistics*, Dublin, Ireland.
- Ralph Weischedel, 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.