# Alignment of Shared Forests for Bilingual Corpora 

Adam Meyers, Roman Yangarber, Ralph Grishman<br>New York University<br>715 Broadway, 7th lloor, NY, NY 10003, USA<br>meyers/roman/grishman@cs.nyu.edu


#### Abstract

Research in example-bascd machine translation (IBM'I) has been hampered by the lack of efficient tree alignment algorithms for bilingual corpora. This paper clescribes an alignment algorithm for EBMT whose rumning time is quadratic in the size of the input parse trees. The algorithm uses dynamic programming to score all possible matching nodes between structure-sharing trees or forests. We describe the algorithm, various optimizations, and our implementation.


## 1 Introduction

The development of a machine translation (MT) system recuires the lengthy manual preparation of bilingual lexicons and transfer rules. Research over the past fow years using parallel sentencealigned bilingual corpora suggests ways in which this manual effort can be partly replaced by corpus-based training. Some of this research has treated the sentences as unstructured word sequences to be aligned; this work has primarily involved the acquisition of bilingual lexical correspondences (Chen, 1993), although there has also been an attempt to create a full M'I' system based on such treatment (Brown et al., 1993). Recently, several groups have been exploring the possibility of aligning parallel syntactically analyzed sentences from the source and target languages (cf. (Sato and Nagao, 1990), (Klavans and Tzoukermann, 1990), (Grishman and Kosaka, 1992), (Kaji ct al., 1992), (Matsumoto et al., 1993) and (Grishman, 1994)). This offers the potential for acquiring not just lexical but also structural correspondences between the two languages. The specific goal in aligning syntax trees is to identify the corresponding tree fragments in the source and target trees. By processing a substantial corpus, a large set of such
corresponding fragments can be collected. These can then serve as the example base for a form of example-based M'I' (cf. (Nagao, 1984), (Sato and Nagao, 1990), (Kaji et al., 1992), (Matsumoto ot al., 1993) and (Furuse and Iida, 1994)). This approach requires a fast tree alignment technique; research has been hampered by the lack of efficiont algorithms. This paper deseribes an efficient algorithm for bilingual tree alignment.

## 2 Our Approach

For each input sentence our parser produces a set of trees, corresponding to each possible syntactic analysis. Our parse trees are transformed into a "regularized" format, to represent the PredicatioArgument structure. For each sentence, the output of the parser is a structure-sharing forest. An example of structure sharing between two parse trees of the same input sentence is shown in ligure 1. We apply the parser to the source and target sentences, using a Spanish and an Finglish grammar, respectively. The resulting sets of structuresharing parse trees form the input to the alignment procedure.

Our alignment program employs dynamic programming ${ }^{1}$ algorithms, which are described in detail in later sections. The program begins at the roots of the source and target trees, and proceeds top-down recursively, filling a matrix of scores. Given $N$ nodes in the source tree $T_{s}=T\left(V_{s}, R_{s}\right)^{2}$ and $M$ nodes in the target tree $T_{t}^{\prime}=T\left(V_{l}, l_{t}^{\prime}\right)$, the score matrix is an $N \times M$ matrix. For each pair of $\operatorname{nodes} x_{i}, i=1, \ldots N \in V_{s}$ and $y_{j}, j=1, \ldots M \in V_{t}$, the corresponding entry in the score matrix is a measure of how well $x_{i}$ matches $y_{j}$. The score for cach pair of nodes depends only on the closeness of the lexical entries associated with the nodes and

[^0]

Figure 1: 'I'wo parses for: lil tenis es wh deporte fanorito del muchacho rico.
on the seores of the best matehing pairs of their descendants. Dynamic programmuing assures that each entry in the matrix, (i.e. the score for the corresponding pair of nodes) is computed only once.

We have implemented this approach for 185 sentences from two sources: Lil C'amino Real (a Sipanish textbook from which we used 73 sentences with their English transtations), and Curious George or Jorge ol Curioso (the Jinglisha and Spanish versions of II. A. Ray's popular children's book, from which we used all 111 Anglish sentences aurl 112 Spanish Sentences). Of the total 185 sentences, 57 from Ih Camino Real and 55 from Curious Gcorge produced at least one parse tree in both languages. 'I'he aligmment procedure was applied only to these pairss of sentences. ${ }^{3}$

## 3 Data Structures

Regularized parses are similar to the F structure of Lexical I'unctional (irammar (If(i), except that a dependency type structure is assumed. ${ }^{4}$ Here a relation Role( $P$ hRIM), AR(i) is represented by an are, the PRLiD) of LIPG at the tail of the are, and the ARG (or the role recipiont) at the head

[^1]of the are. For example, ef tenis is the subject of the predicate ser in ligure 1. In a regular. ized parse, certain closed syntactic classes, such as prepositions and subordinate conjunctions, are represented ats are labels demoting roles, (e.g., the preposition de in Figure 1) rather than as nodes in l'sistructure.

Structure sharing anong the trees in the parse forest allows us to reduce the momber of computed scores. We compute the score for a given pair of subtrees only once, regardless of the number of trees which share these subtrees because the score of a pair of nodes depends only on the seores of their desecmedanss (and not of their ancestors). Currently our parser records structure sharing only between NI's. Jixperiments in which all com mon structure is shared, as in I'igure 1, suggest, that extencling structure sharing to other typess of nodes would further improve performance. This structure-sharing approach is based on previous work in optimizing Peature Structure-based parsing. (See for example, (Karthonen, 1985) and (Percira, 1985)).

## 4 The LCA-Preserving Algorithms

We first, disenss the format aspects of the alignment problem and introduce iemminology.

### 4.1 The Maximal Tree Alignment Problem

The ohjective is to find a maximmo-score correspondence between nodes in a pair of trees. The statement of the problem of aligning two trees ${ }^{5} T$ s and $T_{t}$, corresponds closely to that found in (Mat.. sumoto et al., 1993). (Our algorithms are based on those presented in (Steol and Warnow, 1993), (larach ot al., 19951) and (1995a).

We say that a node $x$ is a common ancestor of nodes $a$ and $b$ in a tree ' $t$ ' if there exist, pathis of length $\geq 0$ from $x$ to $a$ and from $x$ to $b$. The least common ancestor (lea) of two nodes $a$ and $b$ is the mode, $x_{0}=\operatorname{lc} a(a, b)$, such that

1. $x_{0}$ is a common ancestor of $a$ and $b$, and
2. For any other common ancestor $a$ of $a$ and $b$, $x_{0}$ is a descendant of $x$.

An aligmment botween two trees $T^{\prime}=\left(V, l^{\prime}\right)$ and $T^{\prime \prime}=\left(V^{\prime}, l^{\prime}\right)$ is a correspondence (a one-to-one mapping) $I: S+S^{\prime}$, where $S S^{-} V$ and $S^{\prime \prime} G V^{\prime}$, which mantains the following relationship:

[^2]

Figure 2: There is no lca-preserving alignment between $a, b$ and $c$, and $a^{\prime}, b^{\prime}$ and $c^{\prime}$.

> If nodes $a \in V$ and $b \in V$ map into nodes $$
a^{\prime}=f(a) \in V^{\prime} \text { and } b^{\prime}=f(b) \in V^{\prime}
$$ then $f(l c a(a, b))=\operatorname{lca}(f(a), f(b))=$ $\quad l c a\left(a^{\prime}, b^{\prime}\right)$

To illustrate, in Figure 2 there is no lca-preserving alignment of the two trees which maps all three of the leaf nodes $a, b$ and $c$ into the nodes $a^{\prime}, b^{\prime}$ and $c^{\prime}$. Lca-preserving alignments are possible which map any two of the leaves.

The algorithm assumes that least common ancestors are preserved in the alignment. We assign a score to each alignment based on the labels of the corresponding nodes and the arcs from these nodes, as described below. The algorithm seeks an alignment with maximal score.

### 4.2 The Algorithm

Let $T_{s}$ and $T_{t}$ be the source and the target trees. The algorithm uses dynamic programming to build up, in a bottom-up fashion, the scores for matching each node in $T_{s}$ against each node of $T_{t}$. There are $O\left(n^{2}\right)$ such scores, where $n=\max \left(\left|T_{s}\right|,\left|T_{t}\right|\right)$ is the number of nodes in the trees. Let $d(v)$ be the degree of a node $v$. We denote children of $v$ by $v_{i}, i=1, \ldots, d(v)$, and the $\operatorname{arc}\left(v, v_{i}\right)$ by $\vec{v}_{i}$.

Procedure $S C O R E_{L C A}$ : The dynamic programming builds up a score function $S\left(v, v^{\prime}\right)$ for all $v \in T_{s}$ and $v^{\prime} \in T_{t}$, which is stored in a $\left|T_{s}\right| \times\left|T_{t}\right|$ matrix $S$. The value $S\left(v, v^{\prime}\right)$ is the score assigned to the best match between the two subtrees rooted at $v$ in $T_{s}$ and at $v^{\prime}$ in $T_{t}$. Initially, $S$ is filled with undefined values. When a value for $S\left(v, v^{\prime}\right)$ is required, and the corresponding entry in
the matrix is undefined, it is recursively computed by the following formula:
$S\left(v, v^{\prime}\right)=\max \left\{\begin{array}{l}\operatorname{MATCH}_{l c a}\left(v, v^{\prime}\right) \\ \max _{i=1, \ldots, d(v)} S\left(v_{i}, v^{\prime}\right)-P\left(\vec{v}_{i}\right) \\ \max _{j=1, \ldots, d\left(v^{\prime}\right)} S\left(v, v_{j}^{\prime}\right)-P\left(\vec{v}_{j}^{\prime}\right)\end{array}\right.$
The function $M A T C H_{l c a}\left(v, v^{\prime}\right)$ is a measure of how well the nodes $v$ and $v^{\prime}$ align, and is computed as follows:

$$
\begin{align*}
& \text { MATCH } H_{l c a}\left(v, v^{\prime}\right)=\operatorname{Lex} x_{n o d e}\left(v, v^{\prime}\right)+ \\
& \quad+\max _{p \in \mathcal{P}\left(v, v^{\prime}\right)} \sum_{(i, j) \in p} \operatorname{Lex}_{a r c}\left(\vec{v}_{i}, \vec{v}_{j}^{\prime}\right)+ \\
& \quad+S\left(v_{i}, v_{j}^{\prime}\right) \tag{2}
\end{align*}
$$

where:

- Lex $x_{\text {node }}\left(v, v^{\prime}\right) \geq 0$ is a measure of how closely the label on source node $v$ corresponds to the label on target node $v^{\prime}$ in the bilingual dictionary. Le $\boldsymbol{x}_{\text {arc }}\left(\vec{v}, \vec{v}^{\prime}\right)$ is the corresponding measure for arcs.
- $\mathcal{P}\left(v, v^{\prime}\right)$ is the set of all possible pairings of the children of $v$ against the children of $v^{\prime}$. There are $O(d!)$ such pairings, where $d$ is the smaller of the degrees of $v$ and $v^{\prime}$.
- $P\left(\vec{v}_{i}\right) \geq 0$ is the penalty for collapsing the edge $\vec{v}_{i}$, which may depend on the label of that edge.

The summation in (2) ranges over all pairs, denoted by $(i, j)$, which appear in a given pairing $p \in \mathcal{P}\left(v, v^{\prime}\right)$. The summation is evaluated for all $O(d!)$ possible pairings. The pairing with the maximum score is then selected.

The total running time for computing the scores of all of the $O\left(n^{2}\right)$ node pairs $v$ and $v^{\prime}$, is $O\left(d!n^{2}\right)$, where $d$ is the lesser of the degrees of the source and target trees. Computing the max term in (2) can be mapped into the Maximum-Weight Clique problem (which is NP-complete), cf. (Farach et al., 1995b). However, in the NLP domain, the running time is contained because $d<6$ for most trees encountered in practice. Next we describe a heuristic which achieves a time bound quadratic in the size of the tree.

## 5 A Greedy Heuristic

We can reduce the computation time of the max term in (2), if we do not consider all of the $O(d!)$ pairings of the children of $v$ and $v^{\prime}$. Instead we
use a greedy approach and choose the $d$ highestscoring, mutually disjoint pairs from among the $d^{2}$ possible pairs of children of $v$ and $v^{\prime}$. The justification for this heuristic is that we expect that the high-scoring pairs will dominate, and will be a priori mutually disjoint.

The following is an alternative, greedy procedure for computing $S\left(v, v^{\prime}\right)$ :

Procedure GREEDY LCA :

1. $\forall i, j$ s.t. $1 \leq i \leq d(v), 1 \leq j \leq d\left(v^{\prime}\right)$ compute the corresponding entry in a $d(v) \times d\left(v^{\prime}\right)$ matrix $M$ :

$$
M_{i j}=L e x_{a r c}\left(\vec{v}_{i}, \vec{v}_{j}^{\prime}\right)+S\left(v_{i}, v_{j}^{\prime}\right)
$$

The entry $M_{i j}$ of $M=M\left(v, v^{\prime}\right)$ is the score of matching the $i$ th child of $v$ with $j$ th child of $v^{\prime} .{ }^{6}$
2. Let $T^{\prime} O P \leftarrow\{ \}$ be the set of highest scoring pairs.
3. Find the largest entry $M_{i_{0} j_{0}}$ in the matrix, such that neither its row nor its column is already occupied by some pair in TOP:

$$
T O P \leftarrow T O P \cup\left\{\left(i_{0}, j_{0}\right)\right\}
$$

where the coordinates $\left(i_{0}, j_{0}\right)$ are such that

$$
\begin{align*}
& M_{i_{0}, j_{0}}=\max _{i, j}\left\{M_{i j}\right. \\
& \left.\quad \mid \forall\left(i^{\prime}, j^{\prime}\right) \in T O P, i \neq i^{\prime}, j \neq j^{\prime}\right\} \tag{3}
\end{align*}
$$

4. Repeat the above step $d$ times, where $d=$ $\min \left(d(v), d\left(v^{\prime}\right)\right)$.

5 . Compute the result:

$$
\begin{aligned}
& M A T^{\prime} C H_{l c a}\left(v, v^{\prime}\right)= \\
& \quad=\operatorname{Lex}_{n o d e}\left(v, v^{\prime}\right)+\sum_{(i, j) \in T^{\prime} O P} M_{i j}(4)
\end{aligned}
$$

With sorting, this can be done in $O\left(d \log d+d^{2}\right)$ time.

The validity of this heuristic can be tested by comparing the performance of the procedures using the computation in (2) and in (4).

[^3]
## 6 Strict Lexical Matching Heuristic

(Grishman, 1.994) employed an optimization heuristic which favored lexical matches. For each source node $v$ with label $h(v)$, the procedure using this heuristic would first, attempt to find a target node $v^{\prime}$ with label $L\left(v^{\prime}\right)$ such that $L(v)$ translated as $L\left(v^{\prime}\right)$ in the bilingual dictionary (a perfect lexical match). If such a lexical match was found, the procedure did not attempt to match $v$ with any other target node.

A similar heuristic (Lex-Match) was incorporated into our program as the following preprocessing steps:

1 For each source node $v$, all possible lexical matches are identified in the target tree. ${ }^{7}$ If $v$ has at least one possible lexical match, all of those positions in the score matrix $S$ which do not correspond to a lexical match of $v$ are set to \%ero.

2 For each target node $v^{\prime}$ which has at least onc lexical match, all of those positions in the score matrix which do not correspond to a lexical match of $v^{\prime}$ are set to zero.
By setting to zero those positions in the score matrix which represent unlikely matches, this heuristic prevents these scores from ever being calculated, substantially reducing the running time. Lex-Match, unlike the (Grishman, 1994) heuristic, allows one source node to match lexically with more than one node in the target tree.

## 7 Implementation

We have implemented the greedy LCA-preserving algorithm with the following features:

Penalties: The penalties for collapsing edges were set to $0 .{ }^{8}$

Scores: A Lex node score of 100 and a Lex $x_{a r e}$ score of 21 was awarded for each match using our bilingual dictionary. These functions have the value 0 if there is no lexical match.

[^4]| Text | Baseline | Struc-Share | Lex-Match | Struc-Share and Lex-Match |
| :--- | :---: | :---: | :---: | :---: |
| Fl Camino Real | 11.5 sec | 11.3 sec | 8.3 sec | 7.7 sec |
| Curious George | 98.0 sec | 48.8 sec | 87.4 sec | 44.7 sec |
| Total | 109.5 sec | 60.1 sec | 95.7 sec | 52.4 sec |

Table 1: Time Improvements Due to Optimizations

| Text | Lex-Match Off | Lex-Match On |
| :--- | :--- | :--- |
| EI Camino Real | 47 out of $57(82 \%)$ | 47.5 out of $57(83 \%)$ |
| Curious George | 44.6 out of $55(81 \%)$ | 44.6 out of $55(81 \%)$ |
| Total | 91.6 out of $112(82 \%)$ | 92.1 out of $112(82 \%)$ |

Table 2: Changes in Accuracy due to Lex-Match Heuristic

Optimization Variables: We experimented with variants of the procedures which included Structure Sharing (Struc-Share) and the Lexical Match Optimization (Lex-Match), as well as with those that did not.

Table 1 shows the time consumed by our program to align sentences under different conditions. The baseline refers to our program without any optimizations (which is at least 6 times faster than bofore using this algorithm.) The optimization variables have different effects on the different texts. We believe that structure sharing has a much stronger effect on Curious George than on El Camino Real because the former has longer sentences which produced more parses. The hexMatch optimization has a greater effect on Lb Camino Real than on Curious George because all of the words contained in El Camino Real are included in our bilingual dictionary, but only a small portion of the words in Curious George are included. We expect that as the size of our dictionary increases, the Lex-Match optimization will have a greater effect.
'The precision for cach aligned pair ol' sentences is computed according to the formula:

$$
\frac{\mid \text { ResullSSet } \cap \text { AnswerKcy } \mid}{\mid \text { ResullSet } \mid}
$$

where ResultSel is the sel of source parses to which the alignment procedure assigned the highest score, and Answerkey is the set of best source parses as judged by one of the experimenters. ${ }^{9}$ This precision measure was previously used in (Matsumoto et al., 1993) and (Grishman, 1994). Table 2 compares the precision of the alignment, procedure with and without, the Lex-Match heuristic (structure sharing had no effect on the scores.) The slight increase in precision observed with the

[^5]Lex-Match optimization, may be an indication that we should raise the score for lexical matches of node labels.

## 8 Results and Future Directions

The current implementation aligns trees 63 times faster than our previous program (Grishman, 1994), with a $2.3 \%$ improvement in precision. ${ }^{10}$ We expect finc-tuning of the parameters in our procedures to improve our performance. We expect to gain greater efficiency if all common nodes between forests are shared, rather than just the NPs. Another efficiency improvement will be achicved by factoring all ambiguity into the parse trec, as in (Matsumoto et al., 1993). In our current approach, disjunctions are represented only at the root level.

In order to improve the precision of alignment, we plan to experiment with varying the values of the Lex functions and penalties in our scoring algorithm and expanding our bilingual dictionary. We will also experiment with the non-greedy algorithm discussed above and a dominancepreserving algorithm (a less constrained version of the algorithm which we have omitted due to space limitations). In the dominance-preserviug algorithm we relax the requirement of lea-preservation, and require the preservation of the dominance relationship between nodes:

If, for two nodes $a \in T$ and $b \in T, a$ dominates $b$ (denoted as $a \prec b$ ), then for $f(a) \in l^{\prime \prime}$ and $f(b) \in l^{\prime \prime}, f(a) \prec$ $f(b)$.
The idea which makes it possible to align sentences quickly is that we place restrictions on the ways in which we align the parse trecs. We

[^6]disallow alignments which violate the L ( iA con straint, or the dominance requirement, and permit only one-to-one aligmments; between nodes. Some cases where one might posit a correspondence between a single node $x$ and a group of nodes $G=\left\{y_{1} \ldots y_{n}\right\}$, can be interpreted as an alignment between $x$ and $y_{j}$, for some $j, 1 \leq j \leq n$, where $y_{j}$ dominates the remaining nodes in $(f$. We do not consider other types of one-to-many alignments.

## Acknowledgements

We wish to thank Antonio Moreno Sandoval and Cristina Olmeda Moreno for preparation of the Spanish analyses for Jorge el Curioso. We also thank Catherine Macleod for preparation of the English parses.

This research was supported by the National Science Foundation meder Grant IRI-9303013.

## References

Peter Brown, Stephen A. Della Pietra, Vincent J. Della Pietra and Robert L. Mcrcer. 1993. 'The Mathematics of Statistical Machine 'Iransiation: Parameter bistimation. In "Compulational Linguistics, 19: 26:3 312.
S. Chen. 1993. Aligning Sontences in Bilingual Corpora using lexical information. In Proceedings of the 31 st Annual Meeting of the Association for Computational Linguistics, pages (9) 16, Columbus Ohio. Association for Computational Linguistics, Morristown, New Jersey.
'T'. II. Cormen, C. B. Leiserson and R. L. Rivest. 1990. Introduction to Algorithms, 'The MIT' Presss, Cambridge, Mass.

Martin Sarach, Teresa M. Pryytycka and Mikkel Thorup. 1995. 'The maximum agreement subtree problem for binary trees. Unpublished manuseript, Rutgers University, Odense University, and University of Copenhagen.
Martin Farach, 'İresa M. Pryylycka and Mikkel Thorup. 1995. On the agreement of many trees. Unpublished manuscript, Rutgers University, Odense University, and University of Copenhagen.

Osamu l'uruse and Hitoshi lida. 1994. Con stituent Bomdary Parsing for Example-Based Machine Transtation. hi COLINC OA Proceedings, Volume 1, pages 105 111, Kyoto Japan.

Ralph Grishman. 1994. Iterative Aligmment of Syntactic Structures for a Bilimgual Corpus. In I'roceedings of the Second Ammal Workshop for Very Large Corpora, 'Tokyo, Japan.

Ralph Grishman and Michiko Kosaka. 1992. Combining Rationalist and Empiricist Ap proaches to Machine 'Translation. In Procectings of the lourth International Conference one Theortical and Methodotogical Issues in Machine ITanstation, Montreal, Canada.

Hiroyuki Kaji, Yunko Kida and Yasututsugo Morimoto. 1992. Lcaruing 'Translation 'Pmplates, from Bilingual Text. In COLING 92 Procectings.

Lauri Kartumen. 1985. Structure-Sharing with Binary 'Trees. In I'roceedings of the ©3nd Ammal Mecting of the Assoriation for Computational Linguistics.

Judith Klavans and Evelyne 'Tzoukermann. 1990). The BICORI) System. Tn COLIN( 90 Proceedings, Volume 3, pages 174179.
Y. Matsmmoto, II. Ishimoto, 'T' Utsuro and M Nagao. 1993. Structural Matching of Paralled lexts. In 3ist Anmual Mecting of the Association for Compulational Linguistics: "Procedings of the Conference".

Makao Nagao. 1984. A Framowork of a Mechan ical Translation between Japanese and English by Analogy Principle. In Alick lilithorn and Ranan Banerji, editors, Artificial and Human Intelligence. Blsevier Science Publishers B.V., Amsterdam, 'The Netherlands.
bernando (\%. N. Percira. 1985. A Structure Sharing Representation for Unification-Based Grammatical l'ormalisms. In Proceedings of the \&3vd Annual Meeting of the Association for Compulational Linguistics.
Satoshi Sato and Makoto Nagao, 1990. Toward Memory-based 'Lranslation. In COLING 9O Procectings, Volume 3, pages 247252.

Mike Stred and 'Tandy Warnow. 1993. Kaiokura 'Tree 'Theorems: Computing the Maximum Agrecment Subtree. In Information Processint Letters, 48: 7782.


[^0]:    ${ }^{1}$ Cf. e.g. (Cormen et al., 1990), pp.299-328
    ${ }^{2}$ The expression $T\left(V_{s}, L_{s}^{\prime}\right)$ denotes a trec as a pair of sets: $V_{s}$ is the set of vertices (nodes) in the tree, and $E_{s}$ is the set of edges (arcs).

[^1]:    ${ }^{3}$ By improving the asymptotic speed of alignment on these few sentences, we open the possibility for as ing, much larger corpota in luture work.
    ${ }^{4}$ (Sato and Nagao, 1990 ) and (Matisumoto et al., 1903) atso assume dependency bype structures in their example-based work.

[^2]:    "Tor simplicity of presentation we state the problem in terms of alignment of trees. In practice we are using an optimized variant of the algorithm, which aligns pairs of structure-sharing forests.

[^3]:    ${ }^{6}$ Note: if we disregard the are labels for simplicity, and set $\operatorname{Lex} \operatorname{arc}(\cdot, \cdot)=0$, then we do not need to build $M$, and may simply use $M_{i j}=S\left(v_{i}, v_{j}^{\prime}\right)$.

[^4]:    ${ }^{7} \Lambda$ match $M\left(v, v^{\prime}\right)$ is also a lexical match if cither $M\left(v, w^{\prime}\right)$ or $M\left(w, v^{\prime}\right)$ is a lexical match, where $w$ and $w^{\prime}$ are children of $v$ and $v^{\prime}$, respectively.
    ${ }^{8}$ When penalties are set to zero and an empty bilingual dictionary is used, the alignment algorithm fills the scoring matrix with zeros. When we introduce non-zero penalties, the alignment procedure prefers matches between nodes dominating similar structures, since nodes dominating dissimilar structures reccive negative scores. We expect that non-zero peraltics will improve precision with a nonempty bilingual dictionary, because they will favor similar structures. In preliminary testing, penalty values of 20 and 30 yielded improvements in precision.

[^5]:    ${ }^{9}$ If there was no correct parse, the parses with the fowest errors were used for purposes of aligmment.

[^6]:    ${ }^{10}$ The dynamic programming algorithm accounts for an approximately $600 \%$ increase in speed of aligmment
    a rough estimate since much of the program has been re-implemented.

