Investigo

Universida$_{de}$Vigo

Citation for published version:

Gómez-Rodríguez, C., Fernández-González, D., and Bilbao, V. M. D. (2015), Undirected Dependency Parsing. Computational Intelligence, 31, 348–384. doi: 10.1111/coin.12027.

Peer reviewed version

Link to published version: https://doi.org/10.1111/coin.12027

General rights:

Undirected Dependency Parsing

CARLOS GÓMEZ-RODRÍGUEZ

cgomezr@udc.es

Depto. de Computación, Universidade da Coruña. Facultad de Informática,
Campus de Elviña s/n, 15071 A Coruña, Spain

DANIEL FERNÁNDEZ-GONZÁLEZ
danifg@uvigo.es
Depto. de Informática, Universidade de Vigo. Campus As Lagoas
32004 Ourense, Spain

VÍCTOR MANUEL DARRIBA BILBAO
darriba@uvigo.es
Depto. de Informática, Universidade de Vigo. Campus As Lagoas
32004 Ourense, Spain

Dependency parsers, which are widely used in natural language processing tasks, employ a representation of syntax in which the structure of sentences is expressed in the form of directed links (dependencies) between their words. In this article, we introduce a new approach to transition-based dependency parsing in which the parsing algorithm does not directly construct dependencies, but rather undirected links, which are then assigned a direction in a postprocessing step. We show that this alleviates error propagation, because undirected parsers do not need to observe the single-head constraint, resulting in better accuracy.

Undirected parsers can be obtained by transforming existing directed transition-based parsers as long as they satisfy certain conditions. We apply this approach to obtain undirected variants of three different parsers (the Planar, 2-Planar, and Covington algorithms) and perform experiments on several data sets from the CoNLL-X shared tasks and on the *Wall Street Journal* portion of the Penn Treebank, showing that our approach is successful in reducing error propagation and produces improvements in parsing accuracy in most of the cases and achieving results competitive with state-of-the-art transition-based parsers.

Key words: natural language processing, dependency parsing, parsing, computational linguistics, automata

1. INTRODUCTION

Syntactic parsing is the process of determining the grammatical structure of a sentence: given an input sentence, a parsing algorithm (or parser) will analyze it in order to output a representation of its underlying structure. The format of this representation and the information it contains depends on the particular syntactic theory used by the parser. In constituency
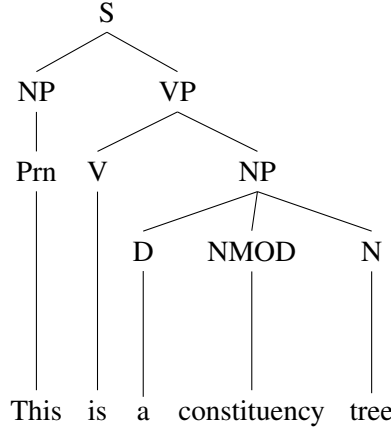
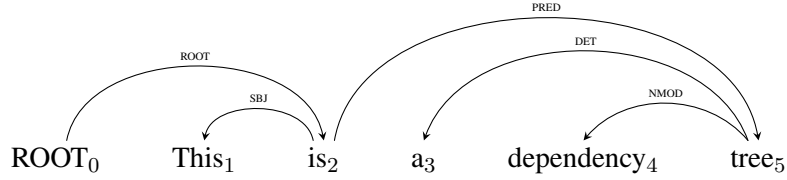FIGURE 1. Constituency tree for an English sentence.

FIGURE 2. Dependency tree for an English sentence.

parsers, or phrase structure parsers, sentences are analyzed by breaking them down into meaningful parts called constituents, which are in turn divided into smaller constituents. The result of such an analysis is represented with a constituency tree, like the one shown in Figure 1. On the other hand, in dependency parsers, the structure of the sentence is represented by a set of directed links (called dependencies) between its words, forming a graph like the one in Figure 2.

Dependency parsing has gained wide popularity in the natural language processing community, and it has recently been applied to a wide range of problems, like machine translation (Ding and Palmer, 2005; Shen et al., 2008; Xu et al., 2009; Katz-Brown et al., 2011), textual entailment recognition (Herrera et al., 2005; Berant et al., 2010), relation extraction (Culotta and Sorensen, 2004; Fundel et al., 2006; Miyao et al., 2009; Katrenko et al., 2010), question answering (Cui et al., 2005; Comas et al., 2010), opinion mining (Joshi and Penstein-Rosé, 2009) or learning for game AI agents (Branavan et al., 2012).

Some important advantages of dependency representations over constituency trees when applied to natural language processing tasks are that they provide a more explicit representation of the semantic information that is useful for applications (for example, the subject and object of a sentence in Figure 2 are explicitly represented in its dependency graph); they do not need intermediate nodes (non-terminals) and hence allow for simpler and more efficient parsing algorithms; and they represent discontinuous linguistic phenomena caused by long-range dependencies or free word order in a natural way by using crossing dependencies.

While there has been research in grammar-driven dependency parsing, where formal grammatical rules are used to define the set of dependency structures that can appear in a language (Tapanainen and Järvinen, 1997; Lombardo and Lesmo, 1996); most current dependency parsers are data-driven, i.e., they use learning techniques to automatically infer
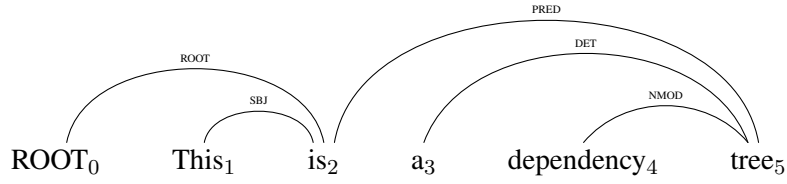
FIGURE 3. Undirected dependency graph for the sentence in Figure 2.

linguistic knowledge from annotated corpora, which can then be used to parse new sentences without the need for an explicit grammar.

In particular, the vast majority of data-driven dependency parsers that have been defined in recent years can be described as being either graph-based or transition-based dependency parsers (Zhang and Clark, 2008; McDonald and Nivre, 2011). Graph-based parsing uses global optimization on models that score dependency graphs (Eisner, 1996; McDonald et al., 2005). In transition-based parsing, which is the focus of this article, dependency graphs are built by sequences of actions by an automaton that transitions between parser states, and each action is scored individually. These scores are used to find a suitable sequence of actions for each given sentence, typically by greedy deterministic search or beam search (Yamada and Matsumoto, 2003; Nivre et al., 2004a; Nivre, 2008). Some key advantages of transition-based parsers are their high efficiency (many of them have linear time complexity, while still providing state-of-the-art accuracy) and the possibility of easily incorporating arbitrarily rich feature models, including non-local features (Zhang and Nivre, 2011).

However, it has been shown by McDonald and Nivre (2007) that transition-based parsers suffer from error propagation: as the generation of a dependency parse for a sentence is modelled as a sequential process, an early erroneous decision may place the parser into an incorrect state, causing more errors later on. In particular, one source of error propagation in transition-based parsers is the need to enforce the single-head constraint, i.e., the common restriction in dependency syntax that forbids each node from having more than one incoming arc. For instance, if we are parsing the sentence in Figure 2 with a transition-based parser that uses greedy deterministic search and we mistakenly make a decision to build an arc from a_3 to $tree_5$ instead of the correct dependency from $tree_5$ to a_3; we will not only miss this dependency but also the one from is_2 to $tree_5$, since we will be in a parser state where the single-head constraint makes it illegal to create it (due to not allowing the node $tree_5$ to have two incoming arcs).

In this article, we introduce a new approach to transition-based parsing that improves accuracy by alleviating this kind of error propagation. To do so, we build novel *undirected* dependency parsers by modifying existing transition-based dependency parsers – namely, the Planar and 2-Planar parsers by Gómez-Rodríguez and Nivre (2010) and the non-projective list-based parser by Nivre (2008), which is a variant of the algorithm by Covington (2001).

The obtained undirected parsers are algorithms that build an undirected graph (like the one on Figure 3) rather than a dependency graph. This means that the single-head constraint need not be observed during the parsing process, since the directed notions of head and dependent (or of incoming and outgoing arcs) are lost in undirected graphs. Therefore, this gives the parser more freedom, and can prevent situations where enforcing the constraint leads to error propagation, like the previous example.

On the other hand, these new algorithms have the obvious disadvantage that their output is an undirected graph, and not a dependency graph. We will need a post-processing step to recover the direction of the dependencies, generating a valid dependency structure. Thus, some complexity is moved from the transition-based parsing process to this post-processing

step; and each undirected parser will outperform the original directed version only if the simplification of the parsing phase is able to avoid more errors than are generated in the post-processing step. Fortunately, as we will see in the experimental sections, this is in fact the case for most of the algorithm-dataset combinations that we tried, showing that the undirected parsing approach is useful to improve parsing accuracy.

The remainder of this article is structured as follows: Section 2 introduces common notation and concepts regarding transition-based dependency parsing, which will be used throughout the article. Section 3 describes a technique to transform transition-based dependency parsers satisfying certain conditions into undirected parsers, and applies it to the Planar, 2-Planar and Covington algorithms. In Section 4, we discuss two different post-processing techniques that can be used to recover dependency trees from undirected graphs. Section 5 puts the ideas in previous sections into practice and presents an empirical study of the accuracy of undirected dependency parsers compared to their directed counterparts, including an evaluation of undirected accuracy measures and direction errors to compare both post-processing approaches. Section 6 presents error analysis to see the effect of undirected parsing on error propagation, and Section 7 discusses related work.. Finally, Section 8 concludes the article with a discussion of the results.

2. TRANSITION-BASED DEPENDENCY PARSING

We now introduce some definitions and notation concerning transition-based dependency parsing, which will serve as a basis to present our undirected parsing approach.

2.1. Dependency parsing and dependency graphs

A *dependency parser* is a system that analyzes natural language sentences and outputs a representation of their syntactic structure in the form of a dependency graph, like the one shown in Figure 2. More formally, a dependency graph can be defined as follows:

Definition 1:   Let $w = w_1 \ldots w_n$ be an input sentence. Given a set L of labels, a *dependency graph* for $w_1 \ldots w_n$ is a labelled directed graph $G = (V_w, A)$, where $V_w = \{0, \ldots, n\}$ is the set of nodes, and $A \subseteq V_w \times L \times V_w$ is the set of labelled directed arcs.

Each of the nodes in V_w encodes the position of a token in w, except for the node 0 which is a dummy node used to mark the root of the sentence and cannot have incoming arcs. Each of the arcs in A encodes a dependency relation between two tokens. We will write the arc $(i, l, j) \in A$ as $i \overset{l}{\to} j$, which will also be called a *dependency link* labelled l from i to j. We then say that i is the *head* of j and, conversely, that j is a syntactic *dependent* of i. The labels on dependency links are typically used to represent their associated syntactic functions, such as SBJ for subject in Figure 2.

Given a dependency graph $G = (V_w, A)$, we will write $i \to j \in A$ if there is a dependency link from i to j, regardless of its label. We will write $i \to^\star j \in A$ if there is a (possibly empty) path from i to j; and $i \leftrightarrow^\star j \in A$ if there is a (possibly empty) path connecting i and j in the undirected graph underlying G. When using these notational conventions, we will omit the references to A when the relevant set of edges is clear from the context.

Most dependency-based syntactic formalisms do not allow arbitrary dependency graphs as syntactic representations. Instead, they are typically restricted to acyclic graphs where each node has at most one head. Such dependency graphs are called *dependency forests*.

Definition 2: A dependency graph G is said to be a *dependency forest* if it satisfies the following:

(1) *Single-head constraint*: if $i \to j$, then there is no $k \neq i$ such that $k \to j$.
(2) *Acyclicity constraint*: if $i \to^\star j$, then there is no arc $j \to i$.

Nodes that have no head in a dependency forest are called *roots*. Apart from the previous two constraints, some dependency formalisms add the additional constraint that a dependency forest can have only one root (or, equivalently, that it must be connected). A forest of this form is called a *dependency tree*.

In this article, we will work with dependency parsers that output dependency forests, i.e., that enforce the single-head and acyclicity constraints. While we do not require the parsers to explicitly enforce connectedness, we note that any dependency forest can be easily converted to a dependency tree (and thus made connected) by linking the dummy root node 0 as the parent of all its other root nodes. Therefore, from now on, we will refer to the outputs produced and the training data used by our parsers as dependency trees, keeping in mind that some of these trees may be a representation of dependency analyses that are forests.

2.2. Transition systems

In a transition-based dependency parser, the dependency analysis for an input sentence is built by a non-deterministic state machine that reads the input sentence and builds dependency arcs. Following the framework of Nivre (2008), this non-deterministic automaton is called a *transition system*, and is defined as follows:

Definition 3: A *transition system* for dependency parsing is a tuple $S = (C, T, c_s, C_t)$, where

(1) C is a set of possible parser *configurations*,
(2) T is a finite set of *transitions*, which are partial functions $t : C \to C$,
(3) c_s is a total initialization function that maps each input string w to a unique *initial configuration* $c_s(w)$, and
(4) $C_t \subseteq C$ is a set of *terminal configurations*.

Although the specific nature of configurations varies among parsers, they are required to contain at least a set A of dependency arcs and a buffer B of unread words, which initially holds all the words in the input sentence. A transition-based parser will be able to read input words by popping them from the buffer, and to create dependency arcs by adding them to the set A.

Given an input string $w = w_1 \ldots w_n$, each of the sequences of configurations that the transition system S can traverse by sequentially applying transitions starting from the initial configuration $c_s(w)$ and ending at some terminal configuration $c_t \in C_t$ is called a *transition sequence* for w in S. The *parse* assigned to w by a transition sequence ending at the terminal configuration c_t is the dependency graph $G = (\{0, \ldots, n\}, A_{c_t})$, where A_{c_t} is the set of arcs stored in the configuration c_t.

Note that, from a theoretical standpoint, it is possible to define a transition system such that no transition sequences exist for a given sentence w. However, this is usually avoided in practice, since robustness (the ability to terminate producing a parse for every possible input) is seen as a desirable quality in data-driven natural language parsers. Therefore, all the transition systems that we will use and define in this article have the property that there is at least one (and typically more than one) transition sequence for every sentence w.

In order to use a transition system to obtain the best dependency analysis for a given sentence, we need to have a mechanism that will select the most suitable among all the

transition sequences that the system allows for that sentence. A standard method to achieve this is by using a classifier to select the best transition to execute at each configuration.

To do so, we define an *oracle* for the transition system $S = (C, T, c_s, C_t)$ as a function $o : C \rightarrow T$, i.e., an oracle is a function that selects a single transition to take at each configuration, and thus can be used to determinize the parsing process. Given a training treebank containing manually annotated dependency trees for sentences, we train a classifier to approximate an oracle by building a canonical transition sequence for each tree in the treebank, and using each of the configurations in the sequence and the corresponding chosen transition as a training instance.

Then, to parse a sentence w, we only need to initialize the parser to the initial configuration $c_s(w)$ and iteratively apply the transitions suggested by the classifier until a terminal configuration is reached. This results in a parser that performs a greedy deterministic search for the best transition sequence, one of the most widely used approaches in transition-based parsing (Yamada and Matsumoto, 2003; Nivre et al., 2004b; Attardi, 2006; Nivre, 2008; Gómez-Rodríguez and Nivre, 2010; Goldberg and Elhadad, 2010; Tratz and Hovy, 2011); although other optimization and search strategies, such as beam search, can also be used (Johansson and Nugues, 2006; Titov and Henderson, 2007; Zhang and Clark, 2008; Huang et al., 2009; Huang and Sagae, 2010; Zhang and Nivre, 2011; Hayashi et al., 2012).

2.3. An example transition system: The Planar transition system

A simple example of a practical transition system for dependency parsing is the *Planar* transition system, introduced by Gómez-Rodríguez and Nivre (2010). The Planar parser is an extension of the well-known arc-eager projective parser by Nivre (2003), which can handle all dependency trees that are *planar*, i.e., those whose arcs can be drawn above the words (as in Figure 2) in such a way that no two arcs cross. In contrast, the arc-eager parser by Nivre (2003) can only build so-called *projective* trees, which is a slightly more restricted set of syntactic structures (Gómez-Rodríguez and Nivre, 2010).

The Planar transition system is a transition system $S = (C, T, c_s, C_t)$ such that

- C is the set of all configurations of the form $c = \langle \sigma, B, A \rangle$, where σ and B are disjoint lists of nodes from V_w (for some input w), and A is a set of dependency arcs over V_w. The list B, called the *buffer*, is used to hold nodes corresponding to input words that have not yet been read. The list σ, called the *stack*, contains nodes for words that have already been read, but still have dependency links pending to be created. For perspicuity, we will use the notation $\sigma|i$ to denote a stack with top i and tail σ, and the notation $j|B$ to denote a buffer with top j and tail B. The set A of dependency arcs contains the part of the output parse that the system has constructed at each given point.
- The initial configuration is $c_s(w_1 \ldots w_n) = \langle [], [1 \ldots n], \emptyset \rangle$, i.e., the buffer initially holds the whole input string while the stack is empty.
- The set of terminal configurations is $C_t = \{ \langle \sigma, [], A \rangle \in C \}$, i.e., final configurations are those where the buffer is empty, regardless of the contents of the stack.
- The set T has the following transitions:

SHIFT $\quad\quad\quad \langle \sigma, i|B, A \rangle \Rightarrow \langle \sigma|i, B, A \rangle$

REDUCE $\quad\quad\quad \langle \sigma|i, B, A \rangle \Rightarrow \langle \sigma, B, A \rangle$

LEFT-ARC$_l$ $\quad \langle \sigma|i, j|B, A \rangle \Rightarrow \langle \sigma|i, j|B, A \cup \{j \xrightarrow{l} i\} \rangle$
$\quad\quad\quad\quad\quad\quad$ only if $\nexists k \mid k \rightarrow i \in A$ (single-head) and $i \leftrightarrow^* j \notin A$ (acyclicity).

RIGHT-ARC$_l$ $\quad \langle \sigma|i, j|B, A \rangle \Rightarrow \langle \sigma|i, j|B, A \cup \{i \xrightarrow{l} j\} \rangle$

The SHIFT transition is used to read words from the input string, by moving the next

Transition	Stack (σ)	Buffer (B)	Added Arc
	$[\text{ROOT}_0]$	$[\text{This}_1,...,\text{tree}_5]$	
SHIFT	$[\text{ROOT}_0, \text{This}_1]$	$[\text{is}_2,..., \text{tree}_5]$	
LA_{SBJ}	$[\text{ROOT}_0, \text{This}_1]$	$[\text{is}_2,..., \text{tree}_5]$	(2, SBJ, 1)
REDUCE	$[\text{ROOT}_0]$	$[\text{is}_2,..., \text{tree}_5]$	
RA_{ROOT}	$[\text{ROOT}_0]$	$[\text{is}_2,..., \text{tree}_5]$	(0, ROOT, 2)
SHIFT	$[\text{ROOT}_0, \text{is}_2]$	$[\text{a}_3,..., \text{tree}_5]$	
SHIFT	$[\text{ROOT}_0, \text{is}_2, \text{a}_3]$	$[\text{dependency}_4, \text{tree}_5]$	
SHIFT	$[\text{ROOT}_0, \text{is}_2, \text{a}_3, \text{dependency}_4]$	$[\text{tree}_5]$	
LA_{NMOD}	$[\text{ROOT}_0, \text{is}_2, \text{a}_3, \text{dependency}_4]$	$[\text{tree}_5]$	(5, NMOD, 4)
REDUCE	$[\text{ROOT}_0, \text{is}_2, \text{a}_3]$	$[\text{tree}_5]$	
LA_{DET}	$[\text{ROOT}_0, \text{is}_2, \text{a}_3]$	$[\text{tree}_5]$	(5, DET, 3)
REDUCE	$[\text{ROOT}_0, \text{is}_2]$	$[\text{tree}_5]$	
RA_{PRED}	$[\text{ROOT}_0, \text{is}_2]$	$[\text{tree}_5]$	(2, PRED, 5)
SHIFT	$[\text{ROOT}_0, \text{is}_2, \text{tree}_5]$	$[\ ]$	
REDUCE	$[\text{ROOT}_0, \text{is}_2]$	$[\ ]$	
REDUCE	$[\text{ROOT}_0]$	$[\ ]$	

FIGURE 4. Transition sequence for parsing the sentence in Figure 2 using the Planar parser (LA=LEFT-ARC, RA=RIGHT-ARC).

node in the buffer to the top of the stack. The LEFT-ARC and RIGHT-ARC transitions build leftward and rightward dependency arcs, respectively, connecting the first node on the buffer and the topmost node on the stack. Finally, the REDUCE transition is used to pop the topmost node from the stack when we have finished building arcs to or from it.

Figure 4 shows a transition sequence in the Planar transition system which derives the labelled dependency graph in Figure 2.

Note that the Planar parser is a linear-time parser, since each word in the input can be shifted and reduced at most once, and the number of arcs that can be built by LEFT-ARC and RIGHT-ARC transitions is strictly bounded by the number of words by the single-head constraint.

2.4. The 2-Planar and Covington transition systems

The undirected dependency parsers defined and tested in this article are based on the Planar transition system described above, the 2-Planar transition system by Gómez-Rodríguez and Nivre (2010) and the version of the Covington (2001) non-projective parser defined by Nivre (2008). We now outline the two latter parsers briefly, a more comprehensive description can be found in the above references.

The 2-Planar transition system is an extension of the Planar system that can recognize a larger set of dependency trees, called 2-planar dependency trees. A dependency tree is said to be 2-planar if it is possible to draw it assigning one out of two colors to each of its dependency arcs, in such a way that arcs sharing the same color do not cross. Gómez-Rodríguez and Nivre (2010) have shown that well over 99% of the dependency trees in natural language treebanks fall into this set, making this parser practical for languages that contain a significant proportion of crossing links, so that planar and projective parsers fall short in coverage.

To handle 2-planar structures, the 2-Planar transition system uses two stacks instead of one, with each stack corresponding to one of the colors that can be assigned to arcs.

At each given configuration, one of the stacks is said to be *active* (meaning that we are building arcs of that color), while the other is *inactive*. Configurations are thus of the form $c = \langle \sigma_0, \sigma_1, B, A \rangle$, where σ_0 is the active stack and σ_1 the inactive stack. The initial configuration is $c_s(w_1 \dots w_n) = \langle [], [], [1 \dots n], \emptyset \rangle$ and the set of terminal configurations is $C_t = \{\langle \sigma_0, \sigma_1, [], A \rangle \in C\}$, analogously to the Planar transition system. The system has the following transitions:

SHIFT $\langle \sigma_0, \sigma_1, i|B, A \rangle \Rightarrow \langle \sigma_0|i, \sigma_1|i, B, A \rangle$

REDUCE $\langle \sigma_0|i, \sigma_1, B, A \rangle \Rightarrow \langle \sigma_0, \sigma_1, B, A \rangle$

LEFT-ARC$_l$ $\langle \sigma_0|i, \sigma_1, j|B, A \rangle \Rightarrow \langle \sigma_0|i, \sigma_1, j|B, A \cup \{j \xrightarrow{l} i\} \rangle$
                  only if $\nexists k \mid k \to i \in A$(single-head) and $i \leftrightarrow^* j \notin A$ (acyclicity).

RIGHT-ARC$_l$ $\langle \sigma_0|i, \sigma_1, j|B, A \rangle \Rightarrow \langle \sigma_0|i, \sigma_1, j|B, A \cup \{i \xrightarrow{l} j\} \rangle$
                  only if $\nexists k \mid k \to j \in A$ (single-head) and $i \leftrightarrow^* j \notin A$ (acyclicity).

SWITCH $\langle \sigma_0, \sigma_1, B, A \rangle \Rightarrow \langle \sigma_1, \sigma_0, B, A \rangle$

The SHIFT transition reads words from the input string exactly as in the Planar transition system, but in this case their corresponding nodes are placed into both stacks. REDUCE, LEFT-ARC and RIGHT-ARC work similarly as in the Planar parser, but they only take into account the active stack, ignoring the inactive one. Finally, a SWITCH transition is added that makes the active stack inactive and vice versa, allowing us to alternate between the two possible arc colors. Despite this added functionality, the 2-Planar parser still runs in linear time.

On the other hand, the Covington algorithm is a transition system that runs in quadratic time, but it has the advantage of being able to parse every possible dependency tree, without restrictions like planarity or 2-planarity. The basic algorithm was first described by Covington (1990, 2001). Nivre (2008) implements a variant of this strategy as a transition system, which is the version we use here.

Configurations in this system are of the form $c = \langle \lambda_1, \lambda_2, B, A \rangle$, where λ_1 and λ_2 are *lists* containing nodes associated with partially processed words, and B is the *buffer* of unprocessed words. The idea of the algorithm is that, after reading each given word, we can do a right-to-left traversal of *all* the nodes for already-read words in the input and create links between them and the first node in the buffer. This traversal is implemented by moving nodes from λ_1 (untraversed nodes) to λ_2 (already traversed nodes). After reading each new input word, all the nodes in both lists are moved back to λ_1 for a new right-to-left traversal to start, hence the quadratic complexity.

The algorithm starts with an initial configuration $c_s(w_1 \dots w_n) = \langle [], [], [1 \dots n], \emptyset \rangle$, and will terminate in final configurations of the set $C_f = \{\langle \lambda_1, \lambda_2, [], A \rangle \in C\}$. The system has the following transitions:

SHIFT $\langle \lambda_1, \lambda_2, i|B, A \rangle \Rightarrow \langle \lambda_1 \cdot \lambda_2|i, [], B, A \rangle$

NO-ARC $\langle \lambda_1|i, \lambda_2, B, A \rangle \Rightarrow \langle \lambda_1, i|\lambda_2, B, A \rangle$

LEFT-ARC$_l$ $\langle \lambda_1|i, \lambda_2, j|B, A \rangle \Rightarrow \langle \lambda_1, i|\lambda_2, j|B, A \cup \{j \xrightarrow{l} i\} \rangle$
                  only if $\nexists k \mid k \to i \in A$ (single-head) and $i \leftrightarrow^* j \notin A$ (acyclicity).

RIGHT-ARC$_l$ $\langle \lambda_1|i, \lambda_2, j|B, A \rangle \Rightarrow \langle \lambda_1, i|\lambda_2, j|B, A \cup \{i \xrightarrow{l} j\} \rangle$
                  only if $\nexists k \mid k \to j \in A$ (single-head) and $i \leftrightarrow^* j \notin A$ (acyclicity).

The SHIFT transition advances the parsing process by reading the first node in the buffer B and inserting it at the head of a list obtained by concatenating λ_1 and λ_2, thus starting a new right-to-left traversal to find candidate nodes to be linked to the one now heading the buffer. This traversal is implemented by the other three transitions: NO-ARC is used when there is no dependency relation between the first node in the buffer and the head of the list α_1, and it moves the head of the list α_1 to α_2 without creating any arcs; while LEFT-ARC and RIGHT-ARC create a leftward (rightward) arc connecting the first node in the buffer and the head of the list α_1, and then move the head of α_1 to α_2. The traversal will end when a new SHIFT transition is executed, signifying that no more arcs will be created involving the first node in the buffer and the nodes in α_1.

3. TRANSFORMING DIRECTED PARSERS INTO UNDIRECTED PARSERS

As mentioned in Section 2.2, practical implementations of transition systems use greedy search or beam search to find the best transition sequence (and thus obtain a dependency tree) for each given input. Since these strategies build transition sequences and dependency arcs in a sequential way from the beginning of the sentence to the end, early parsing decisions may condition and restrict later decisions, causing error propagation. McDonald and Nivre (2007) present an empirical study whose results highlight this phenomenon, showing that a transition-based parser tends to be more accurate than a graph-based parser on arcs that are built early in the transition sequence, but less acurate on arcs built later on.

In particular, one possible source of error propagation is the single-head constraint described in Definition 2. In order to return a valid dependency tree, a transition system must obey this constraint during the whole parsing process. This means that a transition that creates a dependency arc is permissible only if its application does not violate the single-head constraint, i.e., if it does not result in assigning more than one head to the same node. For instance, Figure 4 shows a transition sequence for the Planar parser that correctly parses a sample sentence, assigning it the dependency tree in Figure 2. However, in an alternative scenario where the classifier made a mistake in the eighth transition choosing to apply RA_{NMOD} instead of the correct choice LA_{NMOD}, this would result into building a dependency link from $dependency_4$ to $tree_5$ instead of the correct link from $tree_5$ to $dependency_4$. In turn, this would lead to a situation where creating the (correct) link from is_2 to $tree_5$ would be forbidden by the single-head constraint, as node $tree_5$ would already have an incoming arc. Therefore, in this example, a single erroneous choice of transition initially affecting a single dependency arc propagates to other arcs, due to the single-head constraint, causing at least two attachment errors in the output tree.

In order to remove this source of error propagation, we transform the Planar, 2-Planar and Covington transition systems into variants that build *undirected* graphs instead of directed dependency trees. The goal of this transformation is to allow transition-based parsers to work without needing to obey the single-head constraint. This will make these parsers less sensitive to error propagation, since they will be able to create arcs freely at any point in the parsing sequence, regardless of the existing arcs that have been created before.

The mentioned transformation consists in redesigning the transition systems so that they create dependency links without a direction. In this way, it is not necessary to observe the single-head contraint during the parsing process, since the directed concepts of head and dependent do not apply to undirected links. As a result, the output of these new variants is an undirected graph instead of a tree. This transformation has previously been described in Gómez-Rodríguez and Fernández-González (2012).

3.1. The Undirected Planar, 2-Planar and Covington transition systems

With the goal of obtaining undirected transition systems from the directed ones described in Sections 2.3 and 2.4, we replace the LEFT-ARC and RIGHT-ARC transitions in those systems (which create directed arcs in each direction) with a new transition (ARC) that builds an undirected link. This can be done because, in these three transition systems, the effect of the two directed transitions is the same except for the direction of the created link, so that their behaviour can be collapsed into one common transition: the undirected ARC transition.

In addition to this, the configurations of the undirected transition systems must be changed so that the arc set A is a set of undirected edges, instead of directed arcs. Analogously to our notation for directed arcs, we will use the notation $i \overset{l}{-} j$ as shorthand for an undirected edge labelled l connecting the nodes i and j.

Furthermore, since the direction of the arcs is lost in undirected graphs, the preconditions of transitions that guarantee the single-head constraint are simply removed from the systems.

If we apply these transformations and leave the Planar, 2-Planar and Covington transition systems otherwise unchanged, we will obtain the respective undirected variants: the *undirected Planar*, the *undirected 2-Planar* and the *undirected Covington* transition systems. The transition set of each undirected transition system is as follows:

Undirected Planar

$$\text{SHIFT} \quad \langle \sigma, i|B, A\rangle \Rightarrow \langle \sigma|i, B, A\rangle$$

$$\text{REDUCE} \quad \langle \sigma|i, B, A\rangle \Rightarrow \langle \sigma, B, A\rangle$$

$$\text{ARC}_l \quad \langle \sigma|i, j|B, A\rangle \Rightarrow \langle \sigma|i, j|B, A \cup \{j \overset{l}{-} i\}\rangle$$
$$\text{only if } i \leftrightarrow^* j \notin A \text{ (acyclicity).}$$

Undirected 2-Planar

$$\text{SHIFT} \quad \langle \sigma_0, \sigma_1, i|B, A\rangle \Rightarrow \langle \sigma_0|i, \sigma_1|i, B, A\rangle$$

$$\text{REDUCE} \quad \langle \sigma_0|i, \sigma_1, B, A\rangle \Rightarrow \langle \sigma_0, \sigma_1, B, A\rangle$$

$$\text{ARC}_l \quad \langle \sigma_0|i, \sigma_1, j|B, A\rangle \Rightarrow \langle \sigma_0|i, \sigma_1, j|B, A \cup \{j \overset{l}{-} i\}\rangle$$
$$\text{only if } i \leftrightarrow^* j \notin A \text{ (acyclicity).}$$

$$\text{SWITCH} \quad \langle \sigma_0, \sigma_1, B, A\rangle \Rightarrow \langle \sigma_1, \sigma_0, B, A\rangle$$

Undirected Covington

$$\text{SHIFT} \quad \langle \lambda_1, \lambda_2, i|B, A\rangle \Rightarrow \langle \lambda_1 \cdot \lambda_2|i, [], B, A\rangle$$

$$\text{NO-ARC} \quad \langle \lambda_1|i, \lambda_2, B, A\rangle \Rightarrow \langle \lambda_1, i|\lambda_2, B, A\rangle$$

$$\text{ARC}_l \quad \langle \lambda_1|i, \lambda_2, j|B, A\rangle \Rightarrow \langle \lambda_1, i|\lambda_2, j|B, A \cup \{j \overset{l}{-} i\}\rangle$$
$$\text{only if } i \leftrightarrow^* j \notin A \text{ (acyclicity).}$$

We show in Figure 5 how the undirected Planar parser analyses a sentence using its own set of transitions. Note that the output of this parsing process is the undirected graph presented in Figure 3 instead of the expected dependency tree in Figure 2. In order to obtain a dependency tree as the final output of the analysis, we will need to apply a post-processing step, which will be described in Section 4.

It is worth remarking that, in order to apply this transformation to obtain an undirected

Transition	Stack (σ)	Buffer (B)	Added Arc
	[ROOT$_0$]	[This$_1$,...,tree$_5$]	
SHIFT	[ROOT$_0$, This$_1$]	[is$_2$,..., tree$_5$]	
ARC$_{SBJ}$	[ROOT$_0$, This$_1$]	[is$_2$,..., tree$_5$]	(1, SBJ, 2)
REDUCE	[ROOT$_0$]	[is$_2$,..., tree$_5$]	
ARC$_{ROOT}$	[ROOT$_0$]	[is$_2$,..., tree$_5$]	(0, ROOT, 2)
SHIFT	[ROOT$_0$, is$_2$]	[a$_3$,..., tree$_5$]	
SHIFT	[ROOT$_0$, is$_2$, a$_3$]	[dependency$_4$, tree$_5$]	
SHIFT	[ROOT$_0$, is$_2$, a$_3$, dependency$_4$]	[tree$_5$]	
ARC$_{NMOD}$	[ROOT$_0$, is$_2$, a$_3$, dependency$_4$]	[tree$_5$]	(4, NMOD, 5)
REDUCE	[ROOT$_0$, is$_2$, a$_3$]	[tree$_5$]	
ARC$_{DET}$	[ROOT$_0$, is$_2$, a$_3$]	[tree$_5$]	(3, DET, 5)
REDUCE	[ROOT$_0$, is$_2$]	[tree$_5$]	
ARC$_{PRED}$	[ROOT$_0$, is$_2$]	[tree$_5$]	(2, PRED, 5)
SHIFT	[ROOT$_0$, is$_2$, tree$_5$]	[ ]	
REDUCE	[ROOT$_0$, is$_2$]	[ ]	
REDUCE	[ROOT$_0$]	[ ]	

FIGURE 5. Transition sequence for parsing the sentence in Figure 3 using the Undirected Planar parser.

parser from a directed one, the original transition system must satisfy the condition that their arc-building transitions (conventionally called LEFT-ARC and RIGHT-ARC) be identical except for the directions of the links that they create. This condition holds in some transition systems in the literature – such as the three systems described above, or the arc-eager DAG parser for enriched dependency representations described by Sagae and Tsujii (2008) –, but not in others. For example, in the well-known arc-eager parser by Nivre (2003), LEFT-ARC transitions pop a node from the stack in addition to creating an arc, while RIGHT-ARC transitions instead remove the topmost buffer node and then push the top stack node back to the buffer. One could still try to transform the arc-eager parser into an undirected variant by converting each of its arc transitions into an undirected transition, without necessarily collapsing them into one, but this would result into a parser that violates the acyclicity constraint; because the original system is designed in such a way that both constraints are enforced jointly and acyclicity is only guaranteed if the single-head constraint is also kept. It is easy to check that this problem cannot happen in parsers where LEFT-ARC and RIGHT-ARC transitions are symmetrical in the manner described above: in these systems, if a given graph is not parsable in the original system, then its underlying undirected graph will not be parsable in the transformed system.

3.2. Undirected Feature Models

To implement a practical parser on top of a transition system, we need a feature model to extract relevant information from configurations that will serve to train a classifier. Therefore, apart from modifying the transition system, creating a practical undirected parser necessarily implies to adapt its feature model to work with undirected graphs.

Some features usually employed in transition-based parsers depend on the direction of the arcs that have already been created. Examples of such features are the part-of-speech tag associated with the head of the topmost stack node, or the label of the arc going from

the first node in the buffer to its leftmost dependent.[1] However, since we cannot tell heads from dependents in an undirected graph, these features cannot be used to train an undirected parser.

Therefore, we convert these features into their closest undirected versions: in the previous examples, those would be the part-of-speech tag associated with the closer node linked to the topmost stack node, and the label of the arc that connects the first node in the buffer to the leftmost node linked to it. Notice that now a node (topmost stack or first node in the buffer) has neither head nor dependents, it only has some other nodes linked to it.

More formally, these are the undirected features obtained from the directed ones:

- Information (e.g. part-of-speech, label, lemma, etc.) about the ith node linked to a given node (topmost stack node, topmost buffer node, etc.) on the left or on the right, and about the associated undirected arc, typically for $i = 1, 2, 3$,
- Information (e.g. part-of-speech, label, lemma, etc.) about the closest left and right "undirected siblings" of a given node, i.e., the closest node q located to the left of the given node p such that p and q are linked to some common node r located to the right of both, and vice versa. Note that this notion of undirected siblings does not necessarily correspond to siblings in the directed graph: it can also capture other second-order interactions, such as grandparents.

In addition, we create new features based on undirected relations between nodes that provide further context information for the parser. In particular, we found that the following features worked well in practice:

- A boolean feature representing whether two given nodes are linked or not in the undirected graph, and a feature representing the label of the arc between them.

4. RECOVERING ARC DIRECTIONS

The transformed transition systems described in Section 3 have the drawback that the output they produce is an undirected graph, like the one in Figure 3, rather than a proper dependency tree. In order to use these systems and still obtain a directed dependency tree as the final output of the parsing process, we will apply a post-processing step to assign an orientation to the undirected graph (i.e., choose a direction for each of its edges), in such a way that the single-head constraint is obeyed and the result is a valid dependency tree.

For this purpose, we have developed two different reconstruction techniques to recover arc directions from the undirected graph, previously described in less detail in Gómez-Rodríguez and Fernández-González (2012). The first one, called *naive reconstruction*, is based on using the dummy root node to decide the direction that should be assigned to edges, by choosing the unique orientation of the undirected graph obtained by traversing it from the dummy root. The second technique, *label-based reconstruction*, consists of using the edge labels generated by the transition system to assign a preferred direction to each undirected edge, and then choosing the orientation that conforms to as many preferred directions as possible (note that it will not always be possible to conform to the preferred directions of *all* the arcs, as that may generate a graph violating the single-head constraint).

To describe these reconstruction techniques more formally and view them under a common framework, we can formulate the problem of recovering arc directions as an optimum branching (i.e., directed minimum spanning tree) problem on a weighted graph. Given the undirected graph U produced by an undirected parser, we consider its isomorphic symmetric

---

[1] These example features are taken from the default model for the Planar parser in version 1.5 of MaltParser (Nivre et al., 2006).

directed graph, i.e., the directed graph which has an arc for each of both possible directions of an undirected edge in U. Each directed spanning tree of that graph corresponds to an orientation of U. Then, reconstruction techniques can be implemented by assigning weights to each of the arcs in the symmetric graph, so that they encode a criterion to prefer certain orientations of arcs over others, and then using an optimum branching algorithm to find the minimum spanning tree. Different criteria for assigning weights to arcs will produce different reconstruction techniques.

More formally, let $U = (V_w, E)$ be the undirected graph produced by some undirected parser[2] for an input string w (we omit labels for simplicity and readability).

We define the following sets of arcs:

$$A_1(U) = \{(i,j) \mid j \neq 0 \wedge \{i,j\} \in E\},$$
$$A_2(U) = \{(0,i) \mid i \in V_w\}.$$

The set $A_1(U)$ contains the two possible orientations of each edge in U (i.e., the arcs in the symmetric directed graph isomorphic to U) except for those arcs that would have the node 0 as a dependent, which we disallow because we are using that node as a dummy root, and therefore it cannot be assigned a head. On the other hand, the set $A_2(U)$ contains all the possible arcs that link the nodes in V_w as dependents of the dummy root node, regardless of whether their underlying undirected edges were present in U or not. This is so that the reconstruction techniques defined under this framework are allowed to link unattached tokens to the dummy root.

With these arc sets, we define a graph $D(U)$ containing all the candidate arcs that we will consider when reconstructing a dependency structure from U:

$$D(U) = \{V_w, A(U) = A_1(U) \cup A_2(U)\}.$$

The reconstruction process for an undirected graph U consists of finding an optimum branching (i.e., a directed minimum spanning tree) for a weighted directed graph obtained from assigning a cost $c(i,j)$ to each arc (i,j) of the graph $D(U)$, that is, we are looking for a dependency tree $T = (V_w, A_T \subseteq A(U))$ that minimizes $\sum_{(i,j) \in A_t} c(i,j)$.

Such an tree can be calculated using well-known algorithms for the optimum branching problem, like the Chu-Liu-Edmonds algorithm (Chu and Liu, 1965; Edmonds, 1967). In this particular case, we can take advantage of the fact that the graph $D(U)$ has $O(n)$ nodes and $O(n)$ arcs for a string of length n, and use the implementation by Tarjan (1977) to achieve a time complexity of $O(n \log n)$.

The different reconstruction techniques can be defined by establishing different criteria to assign the costs $c(i,j)$ to the arcs in $A(U)$.

Note that, in practice, the impact of the reconstruction process in overall parsing time is negligible. This is because the optimum branching algorithm does not perform any classification and hence does not need to extract features from configurations, which is the bottleneck of transition-based parsers (Bohnet, 2010; Volokh, 2013). As a result, when applying undirected parsing to a linear-time parsing algorithm such as the Planar and 2-Planar parsers, the total observed runtimes grow linearly with sentence length, in spite of the theoretical $O(n \log n)$ complexity of the reconstruction algorithm.

4.1. Naive Reconstruction

A first, very simple reconstruction technique is based on the fact that a rooted tree is implicitly directed, so if we train the undirected parser to obtain trees and consider the

---

[2]Note that, while the approach taken in this article is to obtain undirected parsers by transforming directed parsers, it would also be possible in theory to design an undirected parser from scratch and apply the same reconstruction techniques to it.

dummy node 0 as their root, we will directly obtain a directed dependency structure. In terms of the generic framework described above, this reconstruction can be defined by assigning costs to the arcs of $D(U)$ as follows:

$$c(i,j) \begin{cases} 1 & \text{if } (i,j) \in A_1(U), \\ 2 & \text{if } (i,j) \in A_2(U) \wedge (i,j) \notin A_1(U). \end{cases}$$

This criterion assigns the same cost to both the orientations of each undirected edge in U, and a higher cost to attaching any node to the dummy root that was not directly linked to it in U. Note that the particular values that we have assigned to the costs (1 and 2) are arbitrary: we can choose any other pair of values for the cost of the arcs as long as the second value is larger than the first one. The resulting optimum branching does not change: it comes from maximizing the number of arcs that come from the "cheaper" set ($A_1(U)$).

To obtain satisfactory results with this approach, we must train the undirected parser to explicitly build undirected arcs from the dummy root node to the root word(s) of each sentence using arc transitions. This means that, if our training treebank contains forests, we need to transform them into trees by explicitly linking each of their roots as dependents of the node 0, as explained at the end of Section 2.1.

Under this assumption, if no classification errors are made, the undirected graph U output by the undirected parser will always be an undirected tree, and the minimum spanning tree will correspond to the unique orientation of U making its edges point away from the dummy root.[3] It is easy to see that this orientation must be the correct parse, since any other orientation violates the assumption that node 0 is a root.

This naive reconstruction technique has the advantage of being very simple, while guaranteeing that the correct parse will be recovered if the undirected parser is able to correctly generate its underlying undirected tree. However, this approach lacks robustness, because it decides the direction of all the arcs in the final output based on which node(s) are chosen as sentence heads and linked to the dummy root. This means that a parsing error affecting the undirected edges that involve the root may propagate and result in many dependency links being erroneous. For this reason, this approach for recovering arc directions will not produce good empirical results, as will be seen in Section 5. Fortunately, we can define a more robust criterion where the orientation of arcs is defined in a more distributed manner, without being so sensible to the edges involving the root.

4.2. Label-based Reconstruction

To obtain a more robust and effective reconstruction technique, we first apply a simple transformation to the training corpus so that arcs will have their direction encoded as a part of their label. To do so, if a leftward arc in the training set is labelled X, we relabel it X_l, meaning "a leftward arc labelled X". If a rightward arc in the training set is labelled X, we relabel it X_r, meaning "a rightward arc labelled X".

After training the undirected parser with this modified treebank, its output for a new sentence will be an undirected graph where each edge's label includes an annotation indicating whether the reconstruction process should prefer to link the corresponding pair of nodes with a leftward or with a rightward arc. Note that those annotations represent *preferred directions* — not hard constraints — because, although in the absence of errors it would be possible to simply use the annotations to decode the correct parse for the sentence, in

---

[3]Note that, while we previously suggested using optimum branching algorithms to find the spanning tree for the sake of generality, in this particular case it is not necessary to use such a generic algorithm: the spanning tree can simply be built in $O(n)$ by starting a traversal from the root and orienting each arc in the sense of the traversal. However, this is only valid for this particular reconstruction technique.

practice parsing errors can create situations where it is not possible to conform to all the annotations without violating the single-head constraint in the directed graph resulting from the reconstruction. In these cases, the reconstruction technique will have to decide which annotations to follow and which have to be ignored. For this purpose, we will assign the costs for our minimum branching algorithm so that it will return a tree agreeing with as many annotations as possible.

To achieve this, we denote by $A_{1+}(U) \subseteq A_1(U)$ the set of arcs in $A_1(U)$ that agree with the annotations, that is, arcs $(i, j) \in A_1(U)$ where either $i < j$ and $\{i, j\}$ is labelled X_r in U, or $i > j$ and $\{i, j\}$ is labelled X_l in U. Conversely, we call $A_{1-}(U)$ the set of arcs in $A_1(U)$ that disagree with the annotations, that is, $A_{1-}(U) = A_1(U) \setminus A_{1+}(U)$. Then, we assign costs to the directed arcs in $A(U)$ as follows:

$$c(i, j) \begin{cases} 1 & \text{if } (i, j) \in A_{1+}(U), \\ 2 & \text{if } (i, j) \in A_{1-}(U), \\ 2n & \text{if } (i, j) \in A_2(U) \land (i, j) \notin A_1(U). \end{cases}$$

where n is the length of the string.

With these costs, the optimum branching algorithm will find a spanning tree that agrees with as many annotations as possible, since assigning the direction that agrees with an edge's annotation has a lower cost than assigning the opposite direction. Additional arcs from the root not appearing in the parsing output (i.e. arcs in $A_2(U) \setminus A_1(U)$) can be added, but only if this is strictly necessary to guarantee connectedness (i.e., if the graph U was disconnected), since the cost of such an arc ($2n$) is greater than the sum of costs of any combination of arcs originating from edges in U.

Analogously to the case of the naive reconstruction, the three cost values need not be exactly 1, 2 and $2n$ as chosen above: it is sufficient that the second value is larger than the first, and the third value is larger than n times the second value. The latter property trivially ensures that a spanning tree with k arcs from the third set always costs less than one with $k + 1$ arcs from the third set; which guarantees that this kind of arcs are only added when this is strictly necessary to guarantee connectedness (i.e., only one such arc will be used for each connected component that does not include the artificial root). The former property ensures that, among the trees with a minimal number of arcs from the third set, we choose those that minimize the number of arcs from the second set.

While this may be the simplest cost assignment to implement label-based reconstruction, we have found experimentally that better practical results are obtained if we give the algorithm more freedom to create new arcs from the root, as follows:

$$c(i, j) \begin{cases} 1 & \text{if } (i, j) \in A_{1+}(U) \land (i, j) \notin A_2(U), \\ 2 & \text{if } (i, j) \in A_{1-}(U) \land (i, j) \notin A_2(U), \\ 2n & \text{if } (i, j) \in A_2(U). \end{cases}$$

The difference with the previous variant is that arcs originating from the root now have a cost of $2n$ *even if* their underlying undirected arcs were present in the output of the undirected parser. Informally, this means that the postprocessor will not "trust" the links from the dummy root created by the parser, and may choose to change them (at no extra cost) if this is convenient to obtain a better agreement with the label annotations of the rest of the arcs (see Figure 6 for an example of the difference between both cost assignments). We believe that the higher empirical accuracy obtained with this criterion is probably due to the fact that it is biased towards changing links from the root, which tend to be more problematic for transition-based parsers, while respecting the parser output as much as possible for links located deeper in the dependency structure, for which transition-based parsers have been shown to be more accurate (McDonald and Nivre, 2007).

Note that both variants of label-based reconstruction share the property that, if the undirected parser produces the correct labelled undirected graph for a given sentence, then the
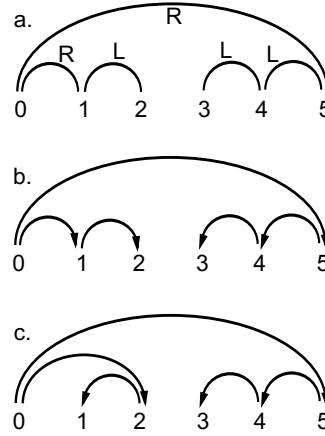
FIGURE 6. a) An undirected graph obtained by the undirected parser trained with a transformed corpus where arcs have been relabelled to specify their direction, b) and c) The dependency graph obtained by each of the variants of the label-based reconstruction (note how the second variant moves an arc from the root).

post-processing will transform it into the correct parse, which is simply the one obtained by following all the annotations in the undirected arcs.

4.3. Example

We can see how the reconstruction techniques work by going back to our running example sentence (Figure 2). In Section 2.3, we saw that the directed Planar parser could parse this sentence with the transition sequence shown in Figure 4. Then, at the beginning of Section 3, we illustrated how a wrong choice by the parser could cause error propagation: if an erroneous link from $dependency_4$ to $tree_5$ were created instead of the correct link from $tree_5$ to $dependency_4$, the single-head constraint would then disallow creating the correct link from is_2 to $tree_5$, causing another attachment error.

If we parsed this sentence with the undirected Planar parser and the naive reconstruction, this error would never happen, since the undirected parser would simply not need to make the choice between linking $dependency_4 \rightarrow tree_5$ or $dependency_4 \leftarrow tree_5$. It would simply execute an ARC transition, as in Figure 5, and produce the undirected arc between $dependency_4$ and $tree_5$ that can be seen in Figure 3. The naive reconstruction technique would then extract from this graph the correct orientation (Figure 2), which is the one where every arc points away from the root.

Note that the naive reconstruction always guarantees that the correct directed parse will be obtained as long as the undirected parse generates the underlying undirected arcs correctly *including those originating from the dummy root*, which point to the root word(s) of the input sentence. This is because rooted trees are implicitly directed. However, if the parser incorrectly identifies the root word of the sentence by linking the dummy root to the wrong node, the naive reconstruction will assign the wrong direction to some of the arcs. An example of this situation can be seen in Figure 7.

On the other hand, if instead of the naive reconstruction we used the label-based reconstruction, the direction error in the example would translate into a labelling error in the undirected parser: instead of creating an undirected edge labelled $NMOD_L$ between the nodes $dependency_4$ and $tree_5$, the edge would be labelled $NMOD_R$, indicating a preference for right attachment. However, this preference would not be followed by the reconstruction
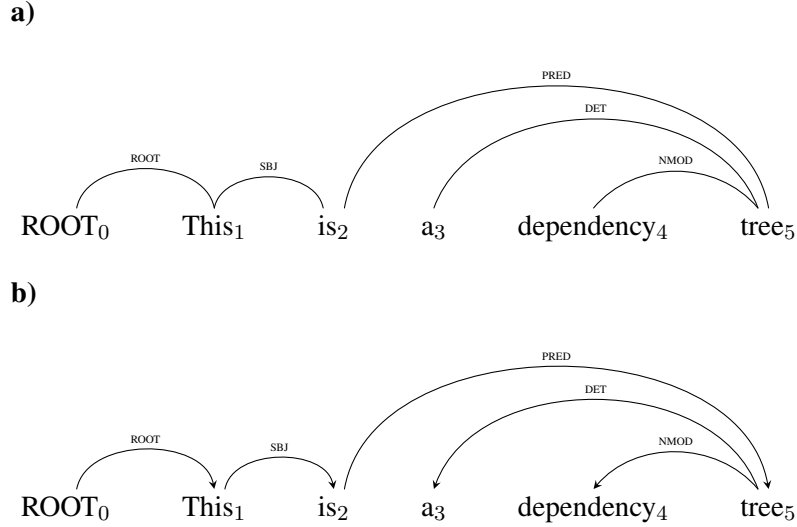
a)

b)

FIGURE 7. Example of an incorrect directed graph obtained by applying the naive reconstruction when the dummy root node was mistakenly attached. **a)** shows the undirected graph obtained after the parsing process where the ROOT was attached to the node $This_1$ instead of the node is_2 as appears in Figure 3; and **b)** shows the resulting dependency tree after applying the naive reconstruction technique, obtaining the arc from $This_1$ to is_2 rather than the correct arc from is_2 to $This_1$ as is described in Figure 2.

technique, since there would be no possible way to conform to all the preferences at the same time without the node $tree_5$ getting two heads, and the only way of disobeying only one annotation (corresponding to the minimum spanning tree, with cost 15 in the second variant of label-based reconstruction) would be to disregard precisely that annotation and output, again, the parse of Figure 2.

Therefore, in this particular toy example, the combination of undirected parsing and any of the reconstruction techniques not only avoids the error propagation due to erroneously linking from $dependency_4$ to $\rightarrow tree_5$, but even eliminates the original error itself during post-processing. Of course, not all the cases will be so favorable when applying these techniques in practice (cf. Figure 7), hence the need to evaluate them empirically to see whether undirected parsing can improve accuracy in real-life settings.

5. EXPERIMENTS

In this section, we evaluate the performance of the undirected Planar, 2-Planar and Covington parsers. For each transition system, we compare the accuracy of the undirected versions with naive and label-based reconstruction to that of the original, directed version. In addition, we provide a comparison to well-known state-of-the-art projective and non-projective parsers.

To evaluate the performance of the parsers in different languages, we use the following eight datasets from the CoNLL-X shared task: Arabic (Hajič et al., 2004), Chinese (Chen et al., 2003), Czech (Hajič et al., 2006), Danish (Kromann, 2003), German (Brants et al., 2002), Portuguese (Afonso et al., 2002), Swedish (Nilsson et al., 2005) and Turkish (Oflazer et al., 2003; Atalay et al., 2003). These datasets have been chosen for their representativity, as they cover a wide range of language families (Germanic, Romance, Semitic, Sino-

Tibetan, Slavic, and Turkic); annotation types (for instance pure dependency annotation in the case of Danish, dependencies extracted from constituents in the case of Chinese, or from discontinuous constituents in German); and degrees of non-projectivity (ranging from the fully projective Chinese treebank to the highly non-projective Czech and German datasets). Buchholz and Marsi (2006) and Havelka (2007) provide detailed information about these and other characteristics of the CoNLL-X treebanks.

In addition to the CoNLL-X datasets, we also perform experiments on English using the WSJ corpus from the well-known Penn Treebank (Marcus et al., 1993). We obtain dependency structures from the constituency trees in the treebank by using the Penn2Malt converter[4], with the head-finding rules of Yamada and Matsumoto (2003), and we follow the standard approach of using sections 2-21 as the training set, 22 as the development set and 23 as the test set.

For our tests, all the algorithms were implemented in MaltParser (Nivre et al., 2006), and trained with classifiers from the LIBSVM (Chang and Lin, 2001) and LIBLINEAR (Fan et al., 2008) packages. In particular, in order to reduce the training time for larger datasets, we employed the LIBLINEAR package for Chinese, Czech, English and German; and we used SVM classifiers from the LIBSVM package for the rest of the languages.

The arc-eager projective and pseudo-projective parsers were trained with the LIBSVM feature models presented in the CoNLL 2006 shared task, where the pseudo-projective version of MaltParser was one of the two top performing systems (Buchholz and Marsi, 2006). The feature models for the 2-Planar parser were taken from Gómez-Rodríguez and Nivre (2010) for the languages included in that paper. In the cases where no dataset was available from these sources (e.g. for English), we tuned our own feature models.

Regarding the new undirected parsers, their feature models for each algorithm and dataset were created from those of the directed parsers as described in Section 3.2.

5.1. Parsing accuracy
To evaluate the accuracy of the undirected parsers with respect to their directed counterparts, we score each parser on the following standard evaluation metrics:

- **Labelled Attachment Score (LAS):** The proportion of tokens (nodes) that are assigned both the correct head and the correct dependency relation label.
- **Unlabelled Attachment Score (UAS):** The proportion of tokens (nodes) that are assigned the correct head (regardless of the dependency relation label).

In our results, we show both LAS and UAS considering every token in the input sentences, including punctuation, as a scoring token.

In particular, Table 1 shows the results obtained by the undirected Planar parser with respect to the original Planar parser by Gómez-Rodríguez and Nivre (2010). Table 2 compares the results of the undirected 2-Planar parser to the 2-Planar parser by Gómez-Rodríguez and Nivre (2010). Finally, Table 3 shows the results obtained by the undirected Covington non-projective parser in comparison to the directed implementation by Nivre (2008).

The results show that the use of undirected parsing with label-based reconstruction (UPlanarL) improves the scores of the Planar parser on all of the nine datasets tested. In most cases, it even attains higher scores than the 2-Planar baseline parser considered, which is remarkable if we take into account that the 2-Planar parser has more theoretical coverage due to its support of crossing links. In the case of 2-planar parsing, applying this technique (U2PlanarL) outperforms the LAS of the directed 2-Planar parser on all the datasets except for Arabic and Portuguese. Finally, the undirected Covington non-projective parser with

---

[4]http://w3.msi.vxu.se/~nivre/research/Penn2Malt.html

TABLE 1. Parsing accuracy of the undirected Planar parser with naive (UPlanarN) and label-based (UPlanarL) postprocessing in comparison to the directed Planar parser (Planar). Results where there is a statistically significant difference at the .05 level between the LAS of Planar and UPlanarL are marked with a star. The last line shows macro averages across all datasets.

| | Planar | | UPlanarN | | UPlanarL | |
Language	LAS	UAS	LAS	UAS	LAS	UAS
Arabic	67.34	77.22	66.33	76.75	**67.50**	**77.57**
Chinese	84.20	88.33	83.10	86.95	**84.50**	**88.35**
Czech	77.70	83.24	75.60	81.14	**77.93**	**83.41**
Danish	82.60	86.64	81.94	86.33	**83.83***	**88.17**
German	83.60	85.67	82.77	84.93	**85.67***	**87.69**
Portug.	83.82	86.88	83.21	86.48	**84.83***	**88.03**
Swedish	82.44	87.36	81.10	85.86	**82.66**	**87.45**
Turkish	71.27	78.57	68.31	75.17	**71.68**	**78.99**
English (WSJ)	85.47	86.55	85.47	86.58	**86.02***	**87.12**
Average	79.83	84.50	78.65	83.35	**80.51***	**85.20**

TABLE 2. Parsing accuracy of the undirected 2-Planar parser with naive (U2PlanarN) and label-based (U2PlanarL) postprocessing in comparison to the directed 2-Planar parser (2Planar). Results where there is a statistically significant difference at the .05 level between the LAS of 2Planar and U2PlanarL are marked with a star. The last line shows macro averages across all datasets.

| | 2Planar | | U2PlanarN | | U2PlanarL | |
Language	LAS	UAS	LAS	UAS	LAS	UAS
Arabic	**67.19**	**77.11**	66.93	77.09	66.52	76.70
Chinese	84.32	**88.27**	82.98	86.81	**84.42**	88.25
Czech	77.91	83.32	75.19	80.80	**78.59**	**84.21**
Danish	83.61	87.63	81.63	85.80	**83.65**	**87.82**
German	85.76	87.86	82.53	84.81	**85.99**	**87.92**
Portug.	**84.92**	**88.14**	83.19	86.33	83.74	87.11
Swedish	82.18	**87.43**	80.71	85.68	**82.25**	87.29
Turkish	70.09	77.39	67.44	74.06	**70.64**	**77.46**
English (WSJ)	85.56	86.66	85.36	86.53	**85.96***	87.04
Average	80.17	**84.87**	78.44	83.10	**80.20**	**84.87**

label-based reconstruction (UCovingtonL) outperforms the results obtained by the baseline parser (Covington) on all the treebanks except for Portuguese and Swedish.

The improvements achieved in LAS by undirected parsers with label-based reconstruction over the directed versions are statistically significant at the .05 level[5] for Danish, English, German and Portuguese for the Planar parser; English in the case of the 2-Planar parser; and Czech, Danish, English and Turkish in the case of the Covington parser. Fur-

---

[5]Statistical significance was assessed using Dan Bikel's randomized comparator: http://www.cis.upenn.edu/~dbikel/software.html

TABLE 3.   Parsing accuracy of the undirected Covington non-projective parser with naive (UCovingtonN) and label-based (UCovingtonL) postprocessing in comparison to the directed algorithm (Covington). Results where there is a statistically significant difference at the .05 level between the LAS of Covington and UCovingtonL are marked with a star. The last line shows macro averages across all datasets.

	Covington		UCovingtonN		UCovingtonL	
Language	LAS	UAS	LAS	UAS	LAS	UAS
Arabic	65.49	**75.69**	62.89	72.83	**65.81**	75.66
Chinese	85.61	89.62	83.48	87.13	**86.17**	**90.04**
Czech	77.43	83.15	71.50	77.96	**78.69***	**84.16**
Danish	82.89	87.06	78.90	83.32	**83.85***	**87.75**
German	85.69	87.78	80.01	82.28	**85.90**	**87.95**
Portug.	**85.21**	**88.21**	81.71	85.17	84.20	87.11
Swedish	**82.76**	**87.61**	81.47	85.96	82.73	87.23
Turkish	72.70	79.75	72.07	79.09	**73.38***	**80.40**
English (WSJ)	84.33	85.44	84.24	85.33	**85.73***	**86.77**
Average	80.23	84.92	77.36	82.12	**80.72***	**85.23**

thermore, no statistically significant *decrease* in accuracy was observed in any of the algorithm/dataset combinations.

As expected, the undirected parsers with naive reconstruction (UPlanarN, U2PlanarN, UCovingtonN) performed worse than those with label-based reconstruction in all the experiments.

To further put these results into context, we provide a comparison of the novel undirected parsers, configured with label-based reconstruction, to well-known projective and non-projective parsers. Table 4 compares the undirected Planar parser to the arc-eager projective parser by Nivre (2003), a well-known algorithm that is also restricted to planar dependency structures.[6] The arc-eager parser is the default parsing algorithm in MaltParser (Nivre et al., 2006), and is also the dependency parser used in other current systems like ZPar (Zhang and Clark, 2011). In addition, Table 5 shows the results of the undirected 2-Planar and the undirected Covington algorithms compared to those of the arc-eager parser with the pseudo-projective transformation of Nivre and Nilsson (2005), which is able to handle non-planar dependencies, and is probably the most widely used method for non-projective transition-based parsing.

As we can see in these experiments, the undirected Planar parser obtains a better score than the arc-eager algorithm (MaltP) on seven out of nine tests; whilst the accuracy of the pseudo-projective arc-eager parser (MaltPP) is outperformed on four out of nine languages by the undirected Covington parser, and on the Arabic, Czech and Danish datasets by the undirected 2-Planar algorithm. Therefore, undirected parsing with label-based reconstruction can be used to improve the accuracy of parsing algorithms, and produces results that are competitive with state-of-the-art transition-based parsers.

5.2. Undirected accuracy

---

[6]The arc-eager parser covers the set of *projective* dependency structures. Planar structures are a very mild relaxaction of projective structures, and in fact both sets become equivalent when sentences are required to have a single root node at position 0.

TABLE 4. Parsing accuracy of the undirected Planar (UPlanarL) with label-based postprocessing in comparison to the MaltParser arc-eager projective (MaltP) algorithm. The last line shows macro averages across all datasets.

	UPlanarL		MaltP	
Language	LAS	UAS	LAS	UAS
Arabic	**67.50**	**77.57**	66.74	76.83
Chinese	84.50	88.35	**86.39**	**90.02**
Czech	**77.93**	**83.41**	77.57	83.19
Danish	**83.83**	**88.17**	82.64	86.91
German	**85.67**	**87.69**	85.48	87.58
Portug.	**84.83**	**88.03**	84.66	87.73
Swedish	**82.66**	87.45	82.44	**87.55**
Turkish	**71.68**	**78.99**	70.96	77.95
English (WSJ)	86.02	87.12	**86.77**	**87.82**
Average	**80.51**	**85.20**	80.41	85.06

TABLE 5. Parsing accuracy of the undirected 2-Planar parser (U2PlanarL) and the undirected Covington non-projective parser (UCovingtonL) with label-based postprocessing in comparison to the MaltParser arc-eager pseudo-projective (MaltPP) algorithm. The last line shows macro averages across all datasets.

	U2PlanarL		UCovingtonL		MaltPP	
Language	LAS	UAS	LAS	UAS	LAS	UAS
Arabic	**66.52**	**76.70**	65.81	75.66	66.02	76.14
Chinese	84.42	88.25	86.17	**90.04**	**86.39**	90.02
Czech	78.59	**84.21**	**78.69**	84.16	78.47	83.89
Danish	83.65	**87.82**	**83.85**	87.75	83.54	87.70
German	85.99	87.92	85.90	87.95	**86.62**	**88.69**
Portug.	83.74	87.11	84.20	87.11	**84.90**	**87.95**
Swedish	82.25	87.29	**82.73**	87.23	82.67	**87.38**
Turkish	70.64	77.46	**73.38**	**80.40**	71.33	78.44
English (WSJ)	85.96	87.04	85.73	86.77	**86.77**	**87.82**
Average	80.20	84.87	80.72	85.23	**80.75**	**85.34**

The results obtained in Section 5.1 show that undirected parsing with the label-based reconstruction is clearly beneficial for parsing accuracy, while the naive reconstruction is not. To obtain a clearer picture of how and why this happens, we can perform an evaluation of the parsers on *undirected evaluation metrics*, i.e., metrics that only take into account the undirected graph underlying the output of each parser, ignoring arc directions. While the standard metrics used in Section 5.1 are the most relevant for practical applications, undirected metrics can provide insights into the influence of each step of undirected parsing on the final accuracy result (i.e., how the accuracy of structural parsing is affected by modifications of feature models described in Section 3.2, and what proportion of errors are introduced by the naive or label-based transformation that is applied afterwards).

Therefore, we will now analyze the results of each parser on the following evaluation metrics:

TABLE 6. Unlabelled parsing accuracy metrics for the undirected Planar parser with naive (UPlanarN) and label-based (UPlanarL) postprocessing in comparison to the directed Planar parser (Planar). The last line shows macro averages across all datasets.

	Planar				UPlanarN				UPlanarL			
Language	ULAS	UUAS	LDE	UDE	ULAS	UUAS	LDE	UDE	ULAS	UUAS	LDE	UDE
Arabic	67.56	78.89	0.22	1.67	67.04	78.67	0.71	1.92	**67.71**	**79.21**	0.21	1.64
Chinese	84.22	**90.34**	0.02	2.01	84.02	89.68	0.92	2.73	**84.58**	90.14	0.08	1.79
Czech	78.06	85.03	0.36	1.79	76.97	84.04	1.37	2.90	**78.35**	**85.34**	0.42	1.93
Danish	82.66	87.66	0.06	1.02	82.79	88.17	0.85	1.84	**83.90**	**88.91**	0.07	0.74
German	83.81	86.34	0.21	0.67	84.26	86.88	1.49	1.95	**85.88**	**88.39**	0.21	0.70
Portug.	83.84	87.35	0.02	0.47	83.62	87.47	0.41	0.99	**84.85**	**88.58**	0.02	0.55
Swedish	82.50	88.33	0.06	0.97	82.02	87.98	0.92	2.12	**82.74**	**88.35**	0.08	0.90
Turkish	71.35	**79.01**	0.08	0.44	70.17	77.57	1.86	2.40	**71.68**	78.99	0.00	0.00
English (WSJ)	85.91	87.75	0.44	1.20	86.55	88.36	1.08	1.78	**86.66**	**88.55**	0.64	1.43
Average	79.99	85.63	0.16	1.14	79.72	85.42	1.07	2.07	**80.71**	**86.27**	0.19	1.08

- **Undirected Labelled Attachment Score (ULAS):** The proportion of labelled arcs that are built correctly by the parser if we disregard direction, i.e., the proportion of arcs $i \xrightarrow{l} j$ in the gold standard annotation such that the parser produces either $i \xrightarrow{l} j$ or $j \xrightarrow{l} i$.
- **Undirected Unlabelled Attachment Score (UUAS):** The proportion of arcs that are built correctly by the parser if we disregard both label and direction, i.e., the proportion of arcs $i \xrightarrow{l} j$ in the gold standard annotation such that the parser produces either $i \xrightarrow{l'} j$ or $j \xrightarrow{l'} i$ for some label l' not necessarily equal to l.
- **Labelled Direction Errors (LDE):** The proportion of labelled arcs that are built correctly by the parser except for their direction, i.e., the proportion of arcs $i \xrightarrow{l} j$ in the gold standard annotation such that the parser outputs the arc $j \xrightarrow{l} i$. This is trivially equivalent to the difference between ULAS and LAS.
- **Unlabelled Direction Errors (UDE):** The proportion of arcs that are built correctly by the parser except for their direction, disregarding labels. This is the proportion of arcs $i \xrightarrow{l} j$ in the gold standard annotation such that the parser outputs an arc of the form $j \xrightarrow{l'} i$ for some label l', and is trivially equivalent to the difference between UUAS and UAS.

Tables 6, 7 and 8 show the undirected scores obtained by the directed and undirected variants of the Planar, 2-Planar and Covington parsers, respectively.

A first interesting insight from the tables is that they tell us the prevalence of direction errors, like the one in the example of Sections 3 and 4.3, in the output of directed parsers. As we can see in the UDE column, more than 1% of the arcs in each parser's output is a direction error (i.e. an arc of the form $a \to b$ that should have been $b \to a$); this represents around 7.5% of the errors made by the parsers. In the majority of the cases where the parsers make a direction error, they miss the dependency label as well, probably because there is a strong correspondence between syntactic functions and dependency directions in most treebanks.

Although undirected parsing can be useful in sentences where the directed parser makes a direction error, because the undirected parser does not need to commit to a direction until the preprocessing step; it is worth reminding that the ultimate goal of undirected parsing is *not* to prevent or correct direction errors, but to prevent error propagation caused by the single-head constraint. For instance, in the example of Section 4.3, the directed parser made

TABLE 7.   Unlabelled parsing accuracy metrics for the undirected 2-Planar parser with naive (U2PlanarN) and label-based (U2PlanarL) postprocessing in comparison to the directed 2-Planar parser (2Planar). The last line shows macro averages across all datasets.

	2Planar				U2PlanarN				U2PlanarL			
Language	ULAS	UUAS	LDE	UDE	ULAS	UUAS	LDE	UDE	ULAS	UUAS	LDE	UDE
Arabic	**67.41**	**78.84**	0.22	1.73	67.41	78.89	0.48	1.80	66.74	78.50	0.22	1.80
Chinese	84.34	**90.24**	0.02	1.97	83.84	89.53	0.86	2.72	**84.48**	90.04	0.06	1.79
Czech	78.32	85.24	0.41	1.92	76.70	84.03	1.51	3.23	**79.07**	**86.06**	0.48	1.85
Danish	83.68	88.70	0.07	1.07	82.45	87.68	0.82	1.88	**83.71**	**88.74**	0.06	0.92
German	85.95	88.39	0.19	0.53	83.79	86.53	1.26	1.72	**86.27**	**88.67**	0.28	0.75
Portug.	**84.92**	**88.68**	0.00	0.54	83.65	87.35	0.46	1.02	83.77	87.73	0.03	0.62
Swedish	82.21	**88.42**	0.03	0.99	81.51	87.43	0.80	1.75	**82.32**	88.19	0.07	0.90
Turkish	70.13	77.83	0.04	0.44	69.74	77.59	2.30	3.53	**70.80**	**77.91**	0.16	0.45
English (WSJ)	86.03	87.91	0.47	1.25	86.52	88.32	1.16	1.79	**86.62**	**88.50**	0.66	1.46
Average	80.33	86.03	0.16	1.16	79.51	85.26	1.07	2.16	**80.42**	**86.04**	0.22	1.17

TABLE 8.   Unlabelled parsing accuracy metrics for the undirected Covington parser with naive (UCovN) and label-based (UCovL) postprocessing in comparison to the directed Covington parser (Cov). The last line shows macro averages across all datasets.

	Cov				UCovN				UCovL			
Language	ULAS	UUAS	LDE	UDE	ULAS	UUAS	LDE	UDE	ULAS	UUAS	LDE	UDE
Arabic	65.77	77.98	0.28	2.29	65.64	77.29	2.75	4.46	**66.72**	**78.41**	0.91	2.75
Chinese	85.65	91.06	0.04	1.44	84.94	90.22	1.46	3.09	**86.27**	**91.50**	0.10	1.46
Czech	77.79	85.08	0.36	1.93	75.11	83.17	3.61	5.21	**79.09**	**86.04**	0.40	1.88
Danish	82.98	88.19	0.09	1.13	79.99	85.59	1.09	2.27	**84.11**	**88.72**	0.26	0.97
German	85.88	88.43	0.19	0.65	82.31	85.16	2.30	2.88	**86.21**	**88.67**	0.31	0.72
Portug.	**85.21**	**88.75**	0.00	0.54	82.75	86.76	1.04	1.59	84.32	87.85	0.12	0.74
Swedish	82.83	**88.42**	0.07	0.81	82.62	87.78	1.15	1.82	**82.89**	88.15	0.16	0.92
Turkish	72.76	80.06	0.06	0.31	73.05	80.38	0.98	1.29	**73.50**	**80.83**	0.12	0.43
English (WSJ)	84.83	86.71	0.50	1.27	85.99	87.69	1.75	2.36	**86.32**	**88.07**	0.59	1.30
Average	80.41	86.08	0.18	1.15	79.16	84.89	1.79	2.77	**81.05**	**86.47**	0.33	1.24

a direction error, but this then caused a non-direction error that was avoided by the undirected parsing technique. Whether the direction error is also finally avoided or not does not depend on the undirected parsing phase, but on the reconstruction technique: thus, in undirected parsers, we can see the ULAS and UUAS metrics as an evaluation of the parsing phase (as its output is an undirected graph), and the LDE and UDE metrics as an evaluation of the reconstruction phase.

In particular, if we compare the direction error metrics for the naive and label-based reconstruction on any of the three algorithms, we can see that the naive reconstruction gets many arc directions wrong. As explained in Section 4.1, this is because the naive approach lacks robustness, as a wrong choice of root can propagate and cause many arcs to be assigned the wrong direction. In contrast, the label-based reconstruction gives very good results, to the point that it does not produce more direction errors than the original directed parsers.

While the differences in direction errors introduced by the transformations explain part of the differences in LAS and UAS scores between the naive and label-based reconstructions

reported in Section 5.1, the tables show that there is also a difference in ULAS and UUAS between both approaches, i.e., the label-based reconstruction already produces a better undirected graph than the naive one even before the reconstruction process. Since the parsing algorithm is the same in both cases, this can only be due to the feature models: in the label-based reconstruction, labels are augmented with preferred directions, and this information is included (in the same way as standard labels) in the feature models, while in the naive reconstruction no directionality information is encoded in labels or used as features.

Therefore, we can see that having no directionality information during parsing has a negative effect on the accuracy of structural prediction. However, adding the information about preferred direction that is available when using the label-based approach compensates for this, making the feature models in Section 3.2 obtain good results that improve over the directed approach.

6. ERROR ANALYSIS

The results in the previous section show that undirected parsing with label-based reconstruction can improve the accuracy of several parsing algorithms, and obtain results that are competitive with the state of the art in transition-based parsing. These results seem to support our hypothesis that the undirected parsing approach can successfully alleviate error propagation, formulated in Section 3.

To further test whether this is in fact the reason for the improvements in accuracy, we conduct a more in-depth analysis of the outputs produced by the parsers in Section 5, going beyond the global directed and undirected accuracy metrics to see the circumstances under which parsing errors are being made.

In particular, as observed by McDonald and Nivre (2007), a good indicator of error propagation in transition-based parsers is the loss of dependency precision for longer dependency arcs (with the length of an arc $i \xrightarrow{l} j$ being defined as $|j-i|$). In most transition-based parsers, shorter arcs tend to be created earlier in transition sequences than longer ones: for example, all the algorithms considered here (like the one used in McDonald and Nivre (2007)) have the property that given nodes $i < j < k < l$, an arc connecting j and k will always be created before an arc connecting i and l. This means that longer arcs, being created later, will be more likely to be affected by the propagation of errors made in previous parsing decisions.

Figure 8 shows the labelled dependency arc precision obtained by the directed and undirected planar parsers for different predicted dependency lengths, i.e., for each length l, it shows the percentage of correct arcs among the arcs of length l in the parser output. Figure 9 shows the labelled dependency arc recall for those same parsers as a function of gold dependency length, i.e., the percentage of gold standard arcs of each length l in the test set that were correctly predicted by the parser.

Instead of doing a language by language analysis, we measure these values across all datasets by aggregating the parser outputs for every language into a single file, and evaluating it with respect to a corresponding joined gold standard file, thus following the same method as McDonald and Nivre (2007). This ensures that we gain a better insight into the distribution of parsing errors, since the datasets for each individual language are too small to have a meaningful sample of arcs for each given distance. Note that, although this form of data aggregation means that the obtained precision and recall values are micro averages, the CoNLL-X test sets purposely have roughly the same size (5000 tokens). In the case of the Penn Treebank we randomly chose a sample of that size as well, so that no language is overrepresented and the result would be roughly the same if we computed a macro average instead.

Starting from length 10, we group the data into bins of size 5 — $[10, 15)$, etc. — for a
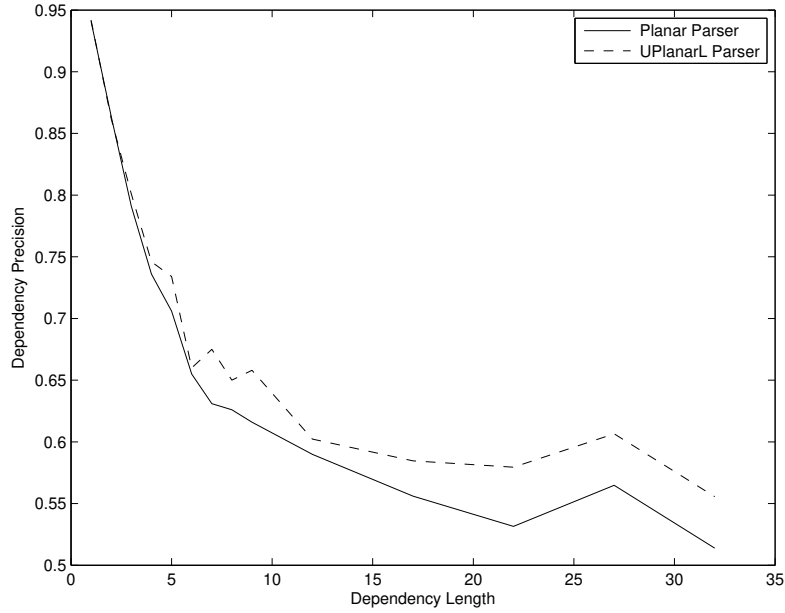
FIGURE 8. Dependency arc precision relative to predicted dependency length for the Planar parser (Planar) and the undirected Planar parser with label-based reconstruction (UPlanarL) on the nine datasets.

more meaningful visualization because, even with the mentioned aggregation, data become sparser from that point and there may be no arcs at all for some particular lengths.

As we can see in Figure 8, both the directed and the undirected planar parsers experiment a drop in precision for longer-distance dependencies, which is expected, among other reasons, because both are transition-based parsers and will exhibit error propagation.[7] However, this phenomenon is significantly more pronounced in the directed parser, and while the accuracies of both algorithms are practically indistinguishable for arcs of length smaller than 5, the undirected Planar parser obtains a clearly better precision on longer arcs, with the difference in precision typically being between 3 and 5 points. Note that the difference in global LAS was not so huge (cf. Section 5) because shorter dependencies are more frequent in treebanks than longer ones, and thus they have a higher weight in the overall LAS metric.

The recall measurements, shown in Figure 9, exhibit a very similar trend. Note that, in principle, precision is a more useful metric than recall for the purpose of estimating the impact of error propagation, because the order in which dependency arcs are built is directly

---

[7]While error propagation is probably the most important reason why a transition-based parser's performance drops for longer dependencies, it is not the only cause for this phenomenon: McDonald and Nivre (2007) observe that even in the graph-based parser by McDonald et al. (2006), which does not analyze sentences sequentially and thus cannot exhibit error propagation, there is a slight drop in accuracy for longer dependencies. A reason for this is that there is less training data for longer dependencies than for shorter ones, since most dependencies in natural language treebanks connect words that are close to each other. Additionally, longer dependencies tend to occur more frequently in ambiguous constructions and in complex, relatively infrequent linguistic phenomena that are difficult to parse (Nivre et al., 2010), while short dependencies include many trivial instances like the attachment of determiners to nouns, where there is little ambiguity and many examples in training sets allowing for a high parsing accuracy.
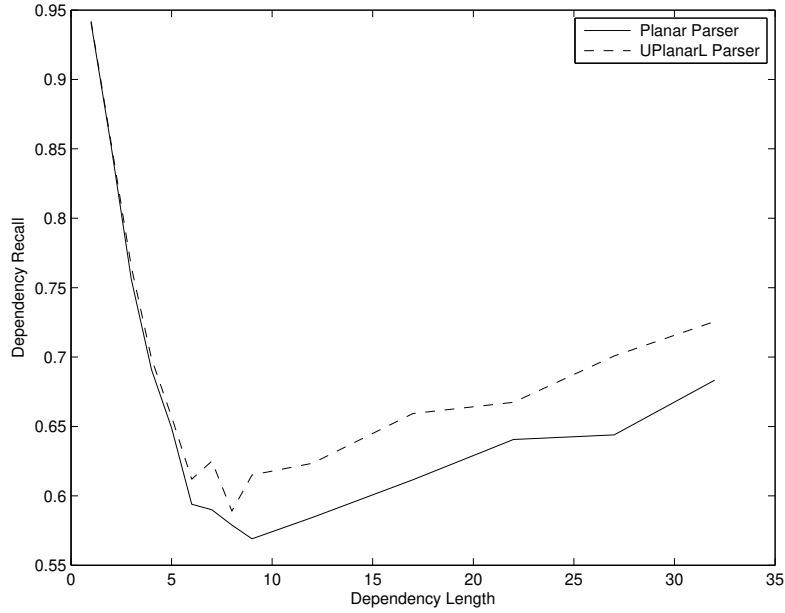
FIGURE 9. Dependency arc recall relative to gold dependency length for the Planar parser (Planar) and the undirected Planar parser with label-based reconstruction (UPlanarL) on the nine datasets.

related to their relative length and position in the output trees, with the relation to the gold standard tree being more indirect. However, we include recall for completeness.

These results for the directed and undirected planar parsers suggest that, as we hypothesized, the undirected variant of the parser is less affected by error propagation than the original, directed version. This is the cause of the higher precision and recall for longer-distance dependencies reflected in Figures 8 and 9, and of the improvement in overall accuracy that we observed in Section 5.

Figures 10 and 11 provide precision and recall measurements analogous to those in Figures 8 and 9, but this time for the directed and undirected 2-planar parsers. In this case, the results are less conclusive, with the undirected parser slightly outperforming the directed one in precision for arcs with lengths between 5 and 10, but behaving worse for those with lengths greater than 10. It should be noted that the overall precision of the undirected parser for arcs of length at least 5 is better than that of the directed version, even though the area under the curve is slightly larger for the latter, because arcs of lengths 5 to 10 are more frequent in treebanks than those of length larger than 10. However, the results do not seem marked enough to suggest that there is a consistent mitigation of error propagation across languages in this case. This is coherent with the results of Section 5, where the 2-Planar algorithm was the one that benefitted the least from undirected parsing in terms of LAS, obtaining improvements in several datasets but only a very slight improvement on the cross-language average.

Finally, Figures 12 and 13 provide the same comparison for the two variants of the Covington algorithm. In this case, it is again clear that the undirected variant exhibits less error propagation than the directed one: it obtains better precision for dependencies of length between 5 and 20 (which are the vast majority of long-distance dependencies), and much
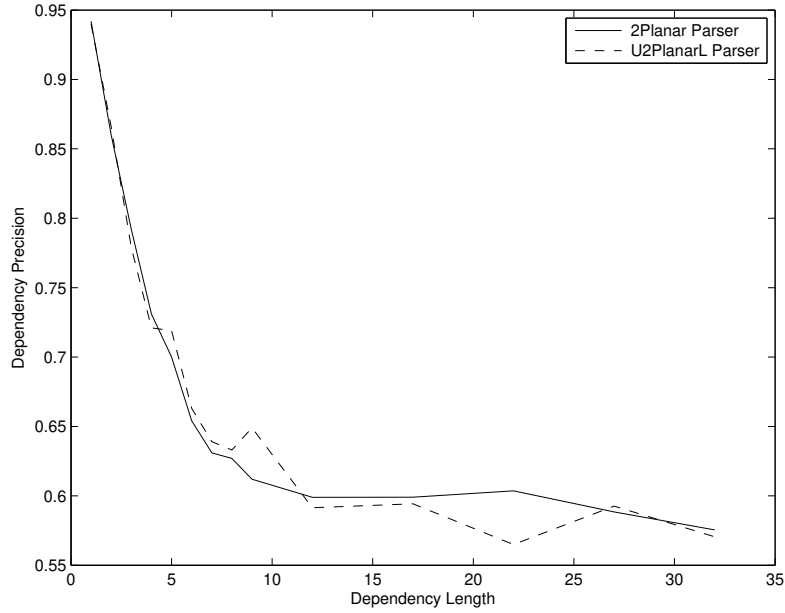
FIGURE 10. Dependency arc precision relative to predicted dependency length for the 2-Planar parser (2Planar) and the undirected 2-Planar parser with label-based reconstruction (U2PlanarL) on the nine datasets.

better recall for those with length between 5 and 30, with differences up to 5 percentage points for some of these lengths. The fact that the improvement is larger in recall than in precision indicates that the undirected Covington parser tends to generate more long arcs than the directed one.

Summing up the results of the error analysis, we can see that undirected parsing clearly increases the precision for longer-distance dependency arcs, at least in the Planar and Covington cases; providing evidence that this technique successfully alleviates error propagation. In the case of the 2-Planar algorithm, the results are less conclusive, since there no clear increase (or decrease) in precision has been observed for longer dependencies. This is in line with the LAS results in Section 5, where applying the undirected technique achieved much larger LAS improvements on the Planar and Covington algorithms than on the 2-Planar parser.

One possible reason why the undirected technique works better for Planar and Covington than for 2-Planar is that the latter parser suffers from less error propagation than the others in the first place. This can be seen by comparing the results for the directed parsers in Figures 8, 10 and 12: in the Planar and Covington parsers the precision soon drops below 0.6 when we increase dependency length, while the 2-Planar parser keeps precisions around 0.6 for longer, suggesting that it is less affected by error propagation.

7. RELATED WORK

In this section we review existing work in the literature related to using undirected relations between words for parsing, as well as other proposals to address the problem of error propagation in transition-based parsers.
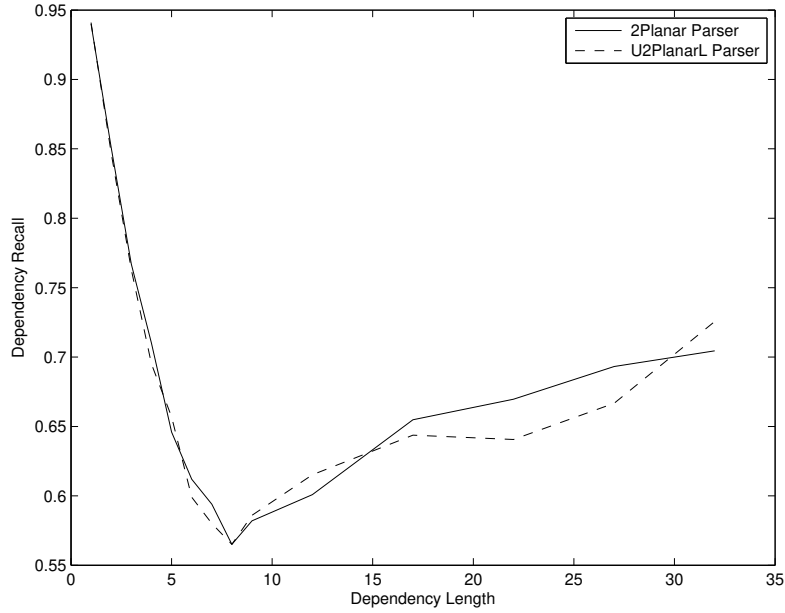
FIGURE 11. Dependency arc recall relative to gold dependency length for the 2-Planar parser (2Planar) and the undirected 2-Planar parser with label-based reconstruction (U2PlanarL) on the nine datasets.

7.1. Undirected parsing

The idea of parsing with undirected relations between words has been applied before in the work on Link Grammar (Sleator and Temperley, 1991). However, in that case undirected arcs are the desired final result, and not an intermediate result, since the Link Grammar formalism itself represents the syntactic structure of sentences by means of undirected links.

While most languages exhibit both left and right dependency relations, some languages that are strongly head-initial or head-final can be annotated assuming that all the dependency arcs have the same orientation. In such languages, dependency parsers need not worry about arc directions and therefore directed and undirected parsing are equivalent. For example, the Japanese parser by Kudo and Matsumoto (2002) simplifies its processing by assuming that all arcs must point to the left, since Japanese is strongly head-final. This is equivalent to producing an undirected dependency structure as the final output. However, such an approach would not be useful for most other languages, as they contain both left and right dependencies.

To the best of our knowledge, the idea of obtaining an undirected graph as an intermediate step for parsing directed dependency structures has not been explored before in the literature.

7.2. Error propagation mitigation

On the other hand, the problem of error propagation in transition-based dependency parsers has received considerable attention in the last few years.

One way to reduce error propagation is to expand the parser's search space by using beam search instead of greedy deterministic search, in such a way that the parser is not
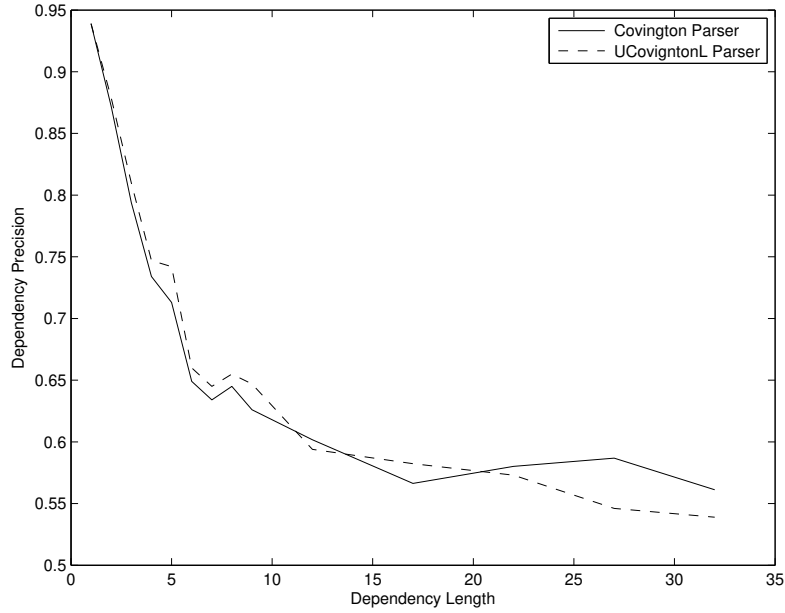
FIGURE 12. Dependency arc precision relative to predicted dependency length for the Covington parser (Covington) and the undirected Covington parser with label-based reconstruction (UCovingtonL) on the nine datasets.

limited to exploring a single transition sequence (Zhang and Clark, 2008; Zhang and Nivre, 2011; Bohnet and Kuhn, 2012). This approach has been shown to produce considerable improvements in accuracy when it is combined with global structural learning, instead of local learning in each parser configuration (Zhang and Nivre, 2012). While this technique produces larger accuracy improvements than undirected parsing, it also has a much larger computational cost. In theory, using a beam increases parsing time by at least a factor proportional to the beam size. In practice, Goldberg et al. (2013) show that most actual implementations additionally add an extra $O(n)$ factor to each transition, making linear transition systems run in quadratic time. In contrast, our undirected parsing approach does not incur any performance penalty during parsing: the only extra processing needed is the post-processing step that assigns directions to arcs, but this has a negligible impact in total runtime because it is a simple graph algorithm that does not need to perform classification or feature extraction, which are the bottlenecks in transition-based parsing (Bohnet, 2010; Volokh, 2013). It is also worth noting that undirected parsing and beam search are not mutually exclusive, since the former is a transformation on transition systems and the latter is a search strategy that can be applied on top of any transition system. Thus, both approaches can be applied simultaneously.

An alternative approach, introduced very recently, tackles the problem in a more efficient way by employing dynamic oracles (Nivre and Goldberg, 2012). In this approach, the parser uses greedy search, but it is trained by an online procedure that considers more than one transition sequence for each given tree (spurious ambiguity) and simulates erroneous transitions to help the parser learn to recover from errors. However, this cannot be applied to any parser because defining a dynamic oracle requires computing a cost function on transitions, and it is only known how to do this efficiently for a very small set of transition-based parsers.
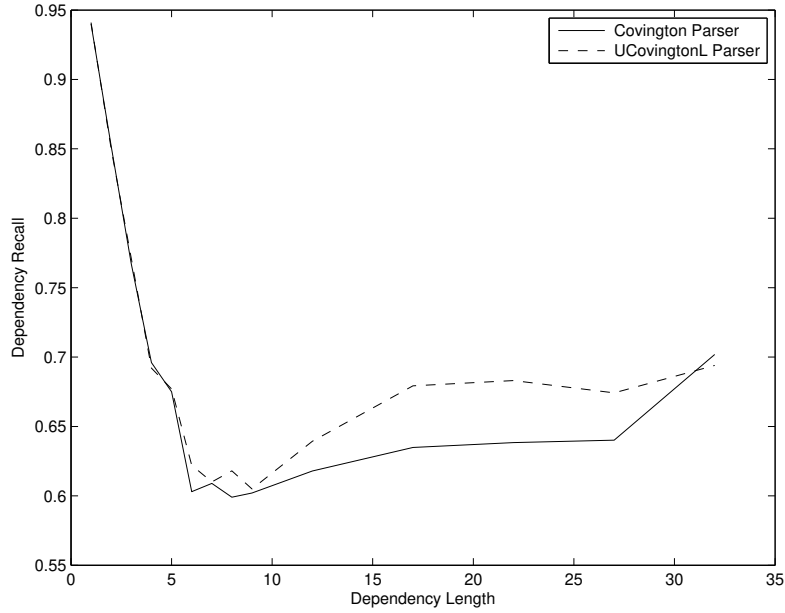
FIGURE 13. Dependency arc recall relative to gold dependency length for the Covington parser (Covington) and the undirected Covington parser with label-based reconstruction (UCovingtonL) on the nine datasets.

In particular, no dynamic oracles have been defined for non-projective parsers, while our undirected approach can be applied to non-projective parsers like 2-Planar and Covington. Note that, if dynamic oracles were defined in the future for algorithms supporting undirected parsing, then both techniques would also be applicable simultaneously.

Finally, the easy-first parser by Goldberg and Elhadad (2010) is a transition-based parser that builds the easier and more reliable dependency arcs first, regardless of their position in the sentence. This limits error propagation because it leaves the decisions more likely to cause errors for the end of the process, and allows the parser to have more contextual information available for those hard decisions, which produces significant improvements in accuracy. However, note that this is an entirely new algorithm running in $O(nlog(n))$, while our approach is a transformation applicable to several different algorithms.

8. CONCLUSION

In this article, we have presented a technique to transform transition-based dependency parsers satisfying certain conditions into undirected dependency parsers, which can then be implemented and trained with feature models that do not depend on the directions of dependency links. The resulting parsers have the drawback that they generate undirected graphs instead of dependency trees, but we have shown how arc directions can be recovered from the undirected output by means of one of two different post-processing techniques, so that the final result is a fully functional dependency parser.

The advantage of the parsers obtained in this way is that they do not need to obey the single-head constraint until the post-processing step. This gives the parser more freedom when choosing transitions to apply, and alleviates error propagation, thus producing im-

provements in accuracy with respect to directed parsers. We have backed this claim with experiments in which we evaluated the directed and undirected version of the Planar, 2-Planar (Gómez-Rodríguez and Nivre, 2010) and Covington (Nivre, 2008; Covington, 2001) parsing algorithms, obtaining improvements in labelled attachment score in 23 out of 27 algorithm-dataset combinations, with statistically significant differences for several of them, outperforming well-known state-of-the-art transition-based parsers. A more in-depth analysis has shown that the undirected parsers tend to perform especially well for longer-distance dependencies, which supports the hypothesis that the increase in accuracy is due to alleviation of error propagation.

Note that this article we have obtained undirected parsers by transforming existing directed parsers, and this provided a good baseline to assess the usefulness of the undirected parsing technique. However, it would also be possible to define undirected parsers from scratch, without necessarily being based on any directed parsers, and apply the same reconstruction techniques to them so as to obtain directed dependency structures as output. This is an interesting avenue for future work.

ACKNOWLEDGMENT

REFERENCES

AFONSO, SUSANA, ECKHARD BICK, RENATO HABER, and DIANA SANTOS. 2002. "Floresta sintá(c)tica": a treebank for Portuguese. *In* Proceedings of the 3rd International Conference on Language Resources and Evaluation (LREC 2002), ELRA, Paris, France, pp. 1968–1703.

ATALAY, NART B., KEMAL OFLAZER, and BILGE SAY. 2003. The annotation process in the Turkish treebank. *In* Proceedings of EACL Workshop on Linguistically Interpreted Corpora (LINC-03), Association for Computational Linguistics, Morristown, NJ, USA, pp. 243–246.

ATTARDI, GIUSEPPE. 2006. Experiments with a multilanguage non-projective dependency parser. *In* Proceedings of the 10th Conference on Computational Natural Language Learning (CoNLL), pp. 166–170.

BERANT, JONATHAN, IDO DAGAN, and JACOB GOLDBERGER. 2010. Global learning of focused entailment graphs. *In* Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, ACL '10, Association for Computational Linguistics, Stroudsburg, PA, USA, pp. 1220–1229. `http://dl.acm.org/citation.cfm?id=1858681.1858805`.

BOHNET, BERND. 2010. Very high accuracy and fast dependency parsing is not a contradiction. *In* Proceedings of the 23rd International Conference on Computational Linguistics, COLING '10, Association for Computational Linguistics, Stroudsburg, PA, USA, pp. 89–97. `http://dl.acm.org/citation.cfm?id=1873781.1873792`.

BOHNET, BERND, and JONAS KUHN. 2012. The best of both worlds: a graph-based completion model for transition-based parsers. *In* Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics, EACL '12, Association for Computational Linguistics, Stroudsburg, PA, USA. ISBN 978-1-937284-19-0. pp. 77–87. `http://dl.acm.org/citation.cfm?id=2380816.2380828`.

BRANAVAN, S. R. K., DAVID SILVER, and REGINA BARZILAY. 2012. Learning to win by reading manuals in a monte-carlo framework. J. Artif. Int. Res., **43**(1):661–704. ISSN 1076-9757. `http://dl.acm.org/citation.cfm?id=2387915.2387932`.

BRANTS, SABINE, STEFANIE DIPPER, SILVIA HANSEN, WOLFGANG LEZIUS, and GEORGE SMITH. 2002.

The tiger treebank. *In* Proceedings of the Workshop on Treebanks and Linguistic Theories, September 20-21, Sozopol, Bulgaria. http://www.coli.uni-sb.de/~sabine/tigertreebank.pdf.

BUCHHOLZ, SABINE, and ERWIN MARSI. 2006. CoNLL-X shared task on multilingual dependency parsing. *In* Proceedings of the 10th Conference on Computational Natural Language Learning (CoNLL), pp. 149–164.

CHANG, CHIH-CHUNG, and CHIH-JEN LIN. 2001. LIBSVM: A Library for Support Vector Machines. Software available at http://www.csie.ntu.edu.tw/~cjlin/libsvm.

CHEN, K., C. LUO, M. CHANG, F. CHEN, C. CHEN, C. HUANG, and Z. GAO. 2003. Sinica treebank: Design criteria, representational issues and implementation. *In* Treebanks: Building and Using Parsed Corpora. *Edited by* A. Abeillé. Kluwer, pp. 231–248.

CHU, Y. J., and T. H. LIU. 1965. On the shortest arborescence of a directed graph. Science Sinica, **14**:1396–1400.

COMAS, PERE R., JORDI TURMO, and LLUÍS MÁRQUEZ. 2010. Using dependency parsing and machine learning for factoid question answering on spoken documents. *In* Proceedings of the 13th International Conference on Spoken Language Processing (INTERSPEECH 2010), Makuhari, Japan.

COVINGTON, MICHAEL A. 1990. A dependency parser for variable-word-order languages. Technical Report AI-1990-01, University of Georgia, Athens, GA.

COVINGTON, MICHAEL A. 2001. A fundamental algorithm for dependency parsing. *In* Proceedings of the 39th Annual ACM Southeast Conference, pp. 95–102.

CUI, HANG, RENXU SUN, KEYA LI, MIN-YEN KAN, and TAT-SENG CHUA. 2005. Question answering passage retrieval using dependency relations. *In* SIGIR '05: Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, ACM, New York, NY, USA. ISBN 1-59593-034-5. pp. 400–407.

CULOTTA, ARON, and JEFFREY SORENSEN. 2004. Dependency tree kernels for relation extraction. *In* ACL '04: Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Morristown, NJ, USA, pp. 423–429.

DING, YUAN, and MARTHA PALMER. 2005. Machine translation using probabilistic synchronous dependency insertion grammars. *In* ACL '05: Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Morristown, NJ, USA, pp. 541–548.

EDMONDS, JACK. 1967. Optimum branchings. Journal of Research of the National Bureau of Standards, **71B**:233–240.

EISNER, JASON M. 1996. Three new probabilistic models for dependency parsing: An exploration. *In* Proceedings of the 16th International Conference on Computational Linguistics (COLING), pp. 340–345.

FAN, R.-E., K.-W. CHANG, C.-J. HSIEH, X.-R. WANG, and C.-J. LIN. 2008. LIBLINEAR: A library for large linear classification. Journal of Machine Learning Research, **9**:1871–1874.

FUNDEL, KATRIN, ROBERT KÜFFNER, and RALF ZIMMER. 2006. RelEx—Relation extraction using dependency parse trees. Bioinformatics, **23**(3):365–371. ISSN 1367-4803.

GOLDBERG, YOAV, and MICHAEL ELHADAD. 2010. An efficient algorithm for easy-first non-directional dependency parsing. *In* Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL HLT), pp. 742–750.

GOLDBERG, YOAV, KAI ZHAO, and LIANG HUANG. 2013. Efficient implementation of beam-search incremental parsers. *In* Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (ACL 2013), Short Papers.

GÓMEZ-RODRÍGUEZ, CARLOS, and DANIEL FERNÁNDEZ-GONZÁLEZ. 2012. Dependency parsing with undirected graphs. *In* Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics, Association for Computational Linguistics, Stroudsburg, PA, USA. ISBN 978-1-937284-19-0. pp. 66–76.

GÓMEZ-RODRÍGUEZ, CARLOS, and JOAKIM NIVRE. 2010. A transition-based parser for 2-planar dependency structures. *In* Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, ACL '10, Association for Computational Linguistics, Stroudsburg, PA, USA, pp. 1492–1501. http://portal.acm.org/citation.cfm?id=1858681.1858832.

HAJIČ, JAN, JARMILA PANEVOVÁ, EVA HAJIČOVÁ, JARMILA PANEVOVÁ, PETR SGALL, PETR PAJAS, JAN ŠTĚPÁNEK, JIŘÍ HAVELKA, and MARIE MIKULOVÁ. 2006. Prague Dependency Treebank 2.0. CDROM CAT: LDC2006T01, ISBN 1-58563-370-4. Linguistic Data Consortium.

HAJIČ, JAN, OTAKAR SMRŽ, PETR ZEMÁNEK, JAN ŠNAIDAUF, and EMANUEL BEŠKA. 2004. Prague Arabic Dependency Treebank: Development in data and tools. *In* Proceedings of the NEMLAR International Conference on Arabic Language Resources and Tools.

HAVELKA, JIŘÍ. 2007. Beyond projectivity: Multilingual evaluation of constraints and measures on non-projective structures. *In* ACL 2007: Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Morristown, NJ, USA, pp. 608–615.

HAYASHI, KATSUHIKO, TARO WATANABE, MASAYUKI ASAHARA, and YUJI MATSUMOTO. 2012. Head-driven transition-based parsing with top-down prediction. *In* Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Jeju Island, Korea, pp. 657–665. http://www.aclweb.org/anthology/P12-1069.

HERRERA, JESÚS, ANSELMO PEÑAS, and FELISA VERDEJO. 2005. Textual entailment recognition based on dependency analysis and WordNet. *In* Machine Learning Challenges, *volume 3944 of* Lecture Notes in Computer Science, Springer-Verlag, Berlin-Heidelberg-New York, pp. 231–239.

HUANG, LIANG, WENBIN JIANG, and QUN LIU. 2009. Bilingually-constrained (monolingual) shift-reduce parsing. *In* Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 1222–1231.

HUANG, LIANG, and KENJI SAGAE. 2010. Dynamic programming for linear-time incremental parsing. *In* Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics (ACL), pp. 1077–1086.

JOHANSSON, RICHARD, and PIERRE NUGUES. 2006. Investigating multilingual dependency parsing. *In* Proceedings of the 10th Conference on Computational Natural Language Learning (CoNLL), pp. 206–210.

JOSHI, MAHESH, and CAROLYN PENSTEIN-ROSÉ. 2009. Generalizing dependency features for opinion mining. *In* Proceedings of the ACL-IJCNLP 2009 Conference Short Papers, ACLShort '09, Association for Computational Linguistics, Stroudsburg, PA, USA, pp. 313–316. http://dl.acm.org/citation.cfm?id=1667583.1667680.

KATRENKO, SOPHIA, PIETER ADRIAANS, and MAARTEN VAN SOMEREN. 2010. Using local alignments for relation recognition. J. Artif. Int. Res., **38**(1):1–48. ISSN 1076-9757. http://dl.acm.org/citation.cfm?id=1892211.1892212.

KATZ-BROWN, JASON, SLAV PETROV, RYAN MCDONALD, FRANZ OCH, DAVID TALBOT, HIROSHI ICHIKAWA, and MASAKAZU SENO. 2011. Training a parser for machine translation reordering. *In* Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing (EMNLP '11). http://petrovi.de/data/emnlp11b.pdf.

KROMANN, MATTHIAS T. 2003. The Danish dependency treebank and the underlying linguistic theory. *In* Proceedings of the 2nd Workshop on Treebanks and Linguistic Theories (TLT), Växjö University Press, Växjö, Sweden, pp. 217–220.

KUDO, TAKU, and YUJI MATSUMOTO. 2002. Japanese dependency analysis using cascaded chunking. *In* Proceedings of the 6th Workshop on Computational Language Learning (CoNLL), pp. 63–69.

LOMBARDO, VINCENZO, and LEONARDO LESMO. 1996. An Earley-type recognizer for dependency grammar. *In* Proceedings of the 16th International Conference on Computational Linguistics (COLING 96), ACL / Morgan Kaufmann, San Francisco, CA, USA, pp. 723–728.

MARCUS, MITCHELL P., BEATRICE SANTORINI, and MARY ANN MARCINKIEWICZ. 1993. Building a large annotated corpus of English: The Penn Treebank. Computational Linguistics, **19**:313–330.

MCDONALD, RYAN, KEVIN LERMAN, and FERNANDO PEREIRA. 2006. Multilingual dependency analysis with a two-stage discriminative parser. *In* Proceedings of the 10th Conference on Computational Natural Language Learning (CoNLL), pp. 216–220.

MCDONALD, RYAN, and JOAKIM NIVRE. 2007. Characterizing the errors of data-driven dependency parsing models. *In* Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL), pp. 122–131.

MCDONALD, RYAN, and JOAKIM NIVRE. 2011. Analyzing and integrating dependency parsers. Comput. Linguist., **37**:197–230. ISSN 0891-2017.

MCDONALD, RYAN, FERNANDO PEREIRA, KIRIL RIBAROV, and JAN HAJIČ. 2005. Non-projective dependency parsing using spanning tree algorithms. *In* Proceedings of the Human Language Technology Conference and the Conference on Empirical Methods in Natural Language Processing (HLT/EMNLP), pp. 523–530.

MIYAO, Y., K. SAGAE, R. SÆTRE, T. MATSUZAKI, and J. TSUJII. 2009. Evaluating contributions of natural language parsers to protein-protein interaction extraction. Bioinformatics, **25**(3):394–400. http://bioinformatics.oxfordjournals.org/cgi/content/abstract/25/3/394.

NILSSON, JENS, JOHAN HALL, and JOAKIM NIVRE. 2005. MAMBA meets TIGER: Reconstructing a Swedish treebank from Antiquity. *In* Proceedings of the NODALIDA Special Session on Treebanks. *Edited by* P. J. Henrichsen.

NIVRE, JOAKIM. 2003. An efficient algorithm for projective dependency parsing. *In* Proceedings of the 8th International Workshop on Parsing Technologies (IWPT), pp. 149–160.

NIVRE, JOAKIM. 2008. Algorithms for Deterministic Incremental Dependency Parsing. Computational Linguistics, **34**(4):513–553. ISSN 0891-2017. . http://www.mitpressjournals.org/doi/abs/10.1162/coli.07-056-R1-07-027.

NIVRE, JOAKIM, and YOAV GOLDBERG. 2012. A dynamic oracle for arc-eager dependency parsing. *In* Proceedings of the 24th International Conference on Computational Linguistics (COLING), pp. 959–976.

NIVRE, JOAKIM, JOHAN HALL, and JENS NILSSON. 2004a. Memory-based dependency parsing. *In* Proceedings of the 8th Conference on Computational Natural Language Learning (CoNLL-2004), Association for Computational Linguistics, Morristown, NJ, USA, pp. 49–56.

NIVRE, JOAKIM, JOHAN HALL, and JENS NILSSON. 2004b. Memory-based dependency parsing. *In* Proceedings of the 8th Conference on Computational Natural Language Learning (CoNLL), pp. 49–56.

NIVRE, JOAKIM, JOHAN HALL, and JENS NILSSON. 2006. Maltparser: A data-driven parser-generator for dependency parsing. *In* Proceedings of the 5th International Conference on Language Resources and Evaluation (LREC), pp. 2216–2219.

NIVRE, JOAKIM, and JENS NILSSON. 2005. Pseudo-projective dependency parsing. *In* Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL), pp. 99–106.

NIVRE, JOAKIM, LAURA RIMELL, RYAN MCDONALD, and CARLOS GÓMEZ RODRÍGUEZ. 2010. Evaluation of dependency parsers on unbounded dependencies. *In* Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010), pp. 833–841.

OFLAZER, KEMAL, BILGE SAY, DILEK ZEYNEP HAKKANI-TÜR, and GÖKHAN TÜR. 2003. Building a Turkish treebank. *In* Treebanks: Building and Using Parsed Corpora. *Edited by* A. Abeillé. Kluwer, pp. 261–277.

SAGAE, KENJI, and JUN'ICHI TSUJII. 2008. Shift-reduce dependency DAG parsing. *In* COLING '08: Proceedings of the 22nd International Conference on Computational Linguistics, Association for Computational Linguistics, Morristown, NJ, USA. ISBN 978-1-905593-44-6. pp. 753–760.

SHEN, LIBIN, JINXI XU, and RALPH WEISCHEDEL. 2008. A new string-to-dependency machine translation algorithm with a target dependency language model. *In* Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (ACL-08: HLT), Association for Computational Linguistics, Morristown, NJ, USA, pp. 577–585.

SLEATOR, DANIEL, and DAVY TEMPERLEY. 1991. Parsing English with a link grammar. Technical Report CMU-CS-91-196, Carnegie Mellon University, Computer Science.

TAPANAINEN, PASI, and TIMO JÄRVINEN. 1997. A non-projective dependency parser. *In* Proceedings of the fifth conference on Applied natural language processing, ANLC '97, Association for Computational Linguistics, Stroudsburg, PA, USA, pp. 64–71. . http://dx.doi.org/10.3115/974557.974568.

TARJAN, R. E. 1977. Finding optimum branchings. Networks, **7**:25–35.

TITOV, IVAN, and JAMES HENDERSON. 2007. A latent variable model for generative dependency parsing. *In* Proceedings of the 10th International Conference on Parsing Technologies (IWPT), pp. 144–155.

TRATZ, STEPHEN, and EDUARD HOVY. 2011. A fast, accurate, non-projective, semantically-enriched parser. *In* Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Edinburgh, Scotland, UK., pp. 1257–1268. http://www.aclweb.org/anthology/D11-1116.

VOLOKH, ALEXANDER. 2013. Performance-Oriented Dependency Parsing. Doctoral dissertation, Saarland University, Saarbrücken, Germany.

XU, PENG, JAEHO KANG, MICHAEL RINGGAARD, and FRANZ OCH. 2009. Using a dependency parser to improve smt for subject-object-verb languages. *In* Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, NAACL '09, Association for Computational Linguistics, Stroudsburg, PA, USA. ISBN 978-1-932432-41-1. pp. 245–253. http://dl.acm.org/citation.cfm?id=1620754.1620790.

YAMADA, HIROYASU, and YUJI MATSUMOTO. 2003. Statistical dependency analysis with support vector machines. *In* Proceedings of the 8th International Workshop on Parsing Technologies (IWPT), pp. 195–206.

ZHANG, YUE, and STEPHEN CLARK. 2008. A tale of two parsers: Investigating and combining graph-based and transition-based dependency parsing. *In* Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 562–571.

ZHANG, YUE, and STEPHEN CLARK. 2011. Syntactic processing using the generalized perceptron and beam search. Computational Linguistics, **37**(1):105–151.

ZHANG, YUE, and JOAKIM NIVRE. 2011. Transition-based dependency parsing with rich non-local features. *In* Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers - Volume 2, HLT '11, Association for Computational Linguistics, Stroudsburg, PA, USA. ISBN 978-1-932432-88-6. pp. 188–193.

ZHANG, YUE, and JOAKIM NIVRE. 2012. Analyzing the effect of global learning and beam-search on transition-based dependency parsing. *In* COLING (Posters). *Edited by* M. Kay and C. Boitet. Indian Institute of Technology Bombay, pp. 1391–1400.