



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Elimination of Spurious Ambiguity in Transition-Based Dependency Parsing

Citation for published version:

Cohen, SB, Gómez-Rodríguez, C & Satta, G 2012 'Elimination of Spurious Ambiguity in Transition-Based Dependency Parsing' ArXiv. <<http://arxiv.org/pdf/1206.6735v1>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Elimination of Spurious Ambiguity in Transition-Based Dependency Parsing

Shay B. Cohen
Department of Computer Science
Columbia University, USA
scohen@cs.columbia.edu

Carlos Gómez-Rodríguez
Departamento de Computación
Universidade da Coruña, Spain
cgomezr@udc.es

Giorgio Satta
Department of Information Engineering
University of Padua, Italy
satta@dei.unipd.it

Abstract

We present a novel technique to remove spurious ambiguity from transition systems for dependency parsing. Our technique chooses a canonical sequence of transition operations (computation) for a given dependency tree. Our technique can be applied to a large class of bottom-up transition systems, including for instance Nivre [2004] and Attardi [2006].

1 Introduction

In parsing, spurious ambiguity refers to ambiguity in a grammar that occurs because several derivations exist for an identical syntactic analysis. When the grammar is enriched with probabilities, the existence of spurious ambiguity implies that the statistical model is defined over *derivations*, a more fine-grained version of the actual syntactic structures of interest. The probability of a syntactic structure then becomes the marginalized probability over all derivations that map to that syntactic structure.

Spurious ambiguity can exist in various grammatical models such as combinatory categorial grammars [Steedman, 2001], tree adjoining grammars [Joshi et al., 1975], data-oriented parsing [Bod, 1992] and transition-based dependency parsing [Nivre, 2005].

While models with spurious ambiguity are statistically more expressive than models without spurious ambiguity,¹ an obstacle exists in the need to marginalize out derivations in order to compute the total probability of a syntactic structure, which is necessary for training and decoding with such models. For many models with spurious ambiguity, it is in fact provably NP-hard to do such marginalization [Sima'an, 1996].

Various heuristics exist to sidestep the need for marginalization. For example, during decoding, one can find the highest-scoring derivation instead of the highest-scoring structure. Under the assumption that most of the probability mass of a given syntactic structure is concentrated on a single derivation, this alternative decoding can be successful. However, this assumption often fails when the probability mass is evenly divided for one syntactic structure but concentrated on a single derivation for another. Even when marginalization can be done efficiently, the likelihood of observed data often becomes non-convex, which is undesirable for training the model because of the local optima problem. For these reasons, it is preferable in most cases to eliminate spurious ambiguity.

In this paper, we focus on eliminating spurious ambiguity that exists in transition-based dependency parsing. Ambiguity arises because several sequences of shift and reduce operations (which assemble a derivation) could yield

¹By this we mean that there are distributions over syntactic structures which can be obtained using models with spurious ambiguity but can not be obtained using models without spurious ambiguity.

identical dependency trees. The transition-based parsing literature has implicitly tackled the issue of spurious ambiguity by defining an **oracle** which, after receiving a dependency tree as input, outputs a unique derivation for that tree based on a canonical ordering of the transition operations. This oracle is then used on the training data (pairs of sentences and dependency trees), yielding new training data (pairs of sentences and shift-reduce derivations) to train multi-class classifiers that decide at each transition step which operation to take [Nivre et al., 2004].

Rather than eliminating spurious ambiguity from the model, this heuristic creates a bias through training to prefer certain derivations for a given dependency tree when doing decoding. In addition, as we discuss in §5, some of the existing oracles for supervised dependency parsing are based on incomplete heuristics (which are often undocumented).

We present a more principled approach to eliminate spurious ambiguity in transition-based dependency parsing. We first define a wide class of bottom-up transition systems, which includes the arc-standard transition system [Nivre, 2008] as well as the transition system from Attardi [2006]. One could also define a transition-based parser using a strategy which is a hybrid between the arc-standard strategy and the easy-first strategy from Goldberg and Elhadad [2010], in which a set of shift actions would need to be taken before a reduction decision is made affecting elements at some deeper position on the stack: this decision can depend on the “easiness” of the reduction. Such a parser can be easily encapsulated into our framework.

We then provide a general technique to enrich the transitions of these systems in order to remove spurious ambiguity while maintaining the completeness of the enriched system with respect to the original. Each tree is associated with a single derivation, which is a sequence of shift and reduce operations such that reduce operations are performed as soon as possible, and conflicts between several reductions are resolved by first attaching dependents that are closer to the current focus point of the parser (top of the stack). This is coherent with psycholinguistic models postulating that humans tend to process local attachments first [Gibson, 2000].

Our approach eliminates ambiguity from a declarative transition system. However, it is extensible to a decoding algorithm as well. The transition systems we introduce can be made probabilistic in a manner similar to the one that appears in Cohen et al. [2011]. Then, a dynamic programming algorithm for these probabilistic systems can be derived so that one can identify the highest scoring derivation and compute the expectations of features in the model [Kuhlmann et al., 2011, Cohen et al., 2011]. Our removal of spurious ambiguity is efficient: the dynamic programming algorithm which is based on the transformed transition system has the same asymptotic complexity as a dynamic programming algorithm for the original transition system.

Our original motivation was to construct a probabilistic model for transition-based dependency parsing, such that a unique (canonical) derivation exists for each dependency tree. This avoids the computational complexity involved in marginalizing derivations. Removal of spurious ambiguity in such a case has to be done at the level of the transition system and not at the level of a tabular method simulating the system or at the level of the resulting parse forest: removing undesired derivations from the chart does not tell us how to set transition probabilities in the original system in such a way that the probability mass of each dependency tree is allocated to a single canonical derivation.

The rest of this paper is organized as follows. We provide an overview of transition-based dependency parsing in §2. We then describe the main details of the spurious ambiguity removal technique in §3. We provide proofs and formal analysis in §4. We apply our technique to the parser from Attardi [2006] and run some experiments in §5. We describe other applications of our technique in §6, and we conclude with an open problem in §7.

2 Transition-Based Dependency Parsing

In this section we briefly introduce the basic definitions for transition-based dependency parsing; we refer the reader to Nivre [2008] for a more detailed presentation. We also define the class of transition-based parsers which is investigated in this paper.

2.1 General Transition Systems

Let Σ be an input alphabet and let $w = a_1 \cdots a_n$, $n \geq 1$, be the input string with $a_i \in \Sigma$ for each i with $1 \leq i \leq n$. A **dependency tree** for w is a directed tree $G = (V_w, A)$ where $V_w = \{0, 1, \dots, n\}$ is the set of nodes and $A \subseteq V_w \times V_w$ is a set of arcs. Each node encodes the position of a token in w , with 0 being a dummy node used as an artificial root,

and each arc encodes a dependency relation between two tokens. We write $i \rightarrow j$ to denote a directed arc $(i, j) \in A$, where node i is the head and node j is the dependent.

A **transition system** for dependency parsing is a tuple $S = (C, T, I, C_t)$, where C is a set of configurations, defined below, T is a finite set of **transitions**, which are partial functions $t: C \rightarrow C$, I is a total initialization function mapping each input string to a unique initial configuration, and $C_t \subseteq C$ is a set of terminal configurations.

A **configuration** is defined relative to input string w , and is a triple (σ, β, A) . Symbols σ and β are disjoint lists of nodes from V_w , called **stack** and input **buffer**, respectively, and $A \subseteq V_w \times V_w$ is a set of arcs. If t is a transition and c_1, c_2 are configurations such that $t(c_1) = c_2$, we write $c_1 \vdash_t c_2$, or simply $c_1 \vdash c_2$ if t is understood from the context.

We denote the stack with its topmost element to the right and the buffer with its first element to the left. We indicate concatenation in the stack and buffer by a vertical bar. For example, for $i \in V_w$, $\sigma|i$ denotes some stack with topmost element i and $i|\beta$ denotes some buffer with first element i . For $1 \leq i \leq n$, β_i denotes the buffer $[i, i + 1, \dots, n]$; for $i > n$, β_i denotes the empty buffer $[]$.

A **computation** of S is a sequence $\gamma = c_0, \dots, c_m, m \geq 1$, of configurations such that, for every i with $1 \leq i \leq m$, $c_{i-1} \vdash_{t_i} c_i$ for some $t_i \in T$. In other words, each configuration in a computation is obtained as the value of the preceding configuration under some transition. A computation can be uniquely specified by its initial configuration c_0 and the sequence t_1, \dots, t_m of its transitions. Thus we will later denote γ in the form $(c_0; t_1, \dots, t_m)$.

2.2 Spurious Ambiguity

A computation $\gamma = c_0, \dots, c_m$ is called **complete** whenever $c_0 = I(w)$ for some input string w , and $c_m \in C_t$. For a complete computation γ we denote as $D(\gamma)$ the unique dependency tree consisting of nodes V_w and all arcs in the final configuration c_m . We say that a transition system has **spurious ambiguity** if, for some pair of complete computations γ and γ' with $\gamma \neq \gamma'$, we have $D(\gamma) = D(\gamma')$.

Informally, the existence of spurious ambiguity implies that there are at least two computations that derive the same dependency tree. Spurious ambiguity exists in various transition systems, such as those in Nivre [2004] and Attardi [2006].

Example 1. *The well-known arc-standard transition system by Nivre [2004] can be defined as follows: its initialization function is $I(a_1 \dots a_n) = ([0], [1 \dots n], \emptyset)$, its set of terminal configurations is $C_t = ([0], [], A)$, and it has the following transitions:*

$$\begin{aligned} \text{shift} &: (\sigma, i|\beta, A) \vdash (\sigma|i, \beta, A) \\ \text{la} &: (\sigma|i|j, \beta, A) \vdash (\sigma|j, \beta, A \cup \{j \rightarrow i\}) \\ \text{ra} &: (\sigma|i|j, \beta, A) \vdash (\sigma|i, \beta, A \cup \{i \rightarrow j\}) \end{aligned}$$

The two following complete computations for a string $w = a_1 a_2 a_3$ produce the same tree with arcs $\{0 \rightarrow 2, 2 \rightarrow 1, 2 \rightarrow 3\}$:

- (i) $(I(w); \text{shift}, \text{shift}, \text{la}, \text{shift}, \text{ra}, \text{ra})$;
- (ii) $(I(w); \text{shift}, \text{shift}, \text{shift}, \text{ra}, \text{la}, \text{ra})$.

Therefore, this transition system has spurious ambiguity, caused by the fact that it allows words (in the example, a_2) to choose whether to collect a left or a right dependent first.

We remark that while in the case of the arc-standard model spurious ambiguity is restricted to a certain set of *permutations* over sequences of operations, i.e., all derivations of a given syntactic tree consist of the same transitions in some permutation, this does not hold in the case of non-projective models.

2.3 Bottom-Up Shift-Reduce Transition Systems

Many of the transition systems for dependency parsing that have been proposed in the literature adopt a bottom-up strategy, meaning that they construct dependency trees starting from the leaves and finishing with the root, by always collecting all the dependents of a given node before assigning it as a dependent of another node. This includes for

instance the already mentioned arc-standard parser, and the non-projective parser of Attardi [2006]. These parsers tend to present spurious ambiguity because, as in Example 1, the left and right dependents of a given node can be collected in different orders. This is in contrast with parsers derived from the arc-eager model [Nivre, 2003] which are not bottom-up and instead impose a unique left-to-right order in which arcs must be constructed.

Some bottom-up transition systems use reduce transitions that affect the buffer, but they can be cast in an alternative form in which all reductions involve only elements from the stack. This is done by considering the first element of the buffer as the topmost stack symbol, as discussed by Cohen et al. [2011]; in this way reductions might take place between stack elements placed at positions deeper than the topmost one. The following definition captures the general form of such models.

Definition A transition system is **bottom-up shift-reduce** if its initialization function is $I(a_1 \cdots a_n) = ([0], [1 \cdots n], \emptyset)$, its set of terminal configurations is $C_t = ([0], [], A)$, and its set of transitions consists of the following:

(i) a shift transition sh of the form $(\sigma, i|\beta, A) \vdash (\sigma|i, \beta, A)$;

(ii) a set of left arc transitions $\text{la}_{p \leftarrow q}$ with $p > q \geq 1$, each of the form

$$(\sigma|i_p|i_{p-1}|\cdots|i_1, \beta, A) \vdash (\sigma|i_{p-1}|\cdots|i_1, \beta, A \cup \{i_q \rightarrow i_p\});$$

(iii) a set of right arc transitions $\text{ra}_{p \rightarrow q}$ with $p > q \geq 1$, each of the form

$$(\sigma|i_p|i_{p-1}|\cdots|i_1, \beta, A) \vdash (\sigma|i_p|\cdots|i_{q+1}|i_{q-1}|\cdots|i_1, \beta, A \cup \{i_p \rightarrow i_q\}).$$

Transitions in (ii) and (iii) above are called **reductions**. The **degree** of reductions $\text{la}_{p \leftarrow q}$ and $\text{ra}_{p \rightarrow q}$ is defined as $p - q$ and is always positive. The **depth** of reductions $\text{la}_{p \leftarrow q}$ and $\text{ra}_{p \rightarrow q}$ corresponds to the index p . The degree of a transition system S , written $\text{deg}(S)$, is the maximum degree among all its reductions. Analogously, the depth of a transition system S , written $\text{depth}(S)$, is the maximum depth among all its reductions.

The next definition introduces a condition that allows us to remove spurious ambiguity from bottom-up shift-reduce parsers. Informally, the condition requires that the existence in the system of a reduction of some type involving stack positions p and q , $p > q$, always implies the existence in the system of reductions of the same type involving stack positions p' and q' with $p' < p$ and $q' \leq q$. We need some additional notation. Let $\mu(\text{la}_{p \leftarrow q})$ be a set of transitions including $\text{la}_{p-1 \leftarrow q}$ if $p > q + 1$, $\text{la}_{p-1 \leftarrow q-1}$ if $q > 1$, and no other transition. Similarly, $\mu(\text{ra}_{p \rightarrow q})$ includes $\text{ra}_{p-1 \rightarrow q}$ if $p > q + 1$, $\text{ra}_{p-1 \rightarrow q-1}$ if $q > 1$, and no other transition.

Definition Let S be a bottom-up shift-reduce transition system with set of transitions T . S is **monotonic** if for each $t \in T$ we have $\mu(t) \subseteq T$.

Example 2. *The transition-based parser of Attardi [2006] can be written as the bottom-up shift-reduce system with transitions sh, $\text{la}_{p \leftarrow 1}$ and $\text{ra}_{p \rightarrow 1}$ for every p with $2 \leq p \leq d$, $d = \text{depth}(S)$. The system with depth 3, as used by Kuhlmann and Nivre [2010], Cohen et al. [2011], has transitions sh, $\text{la}_{2 \leftarrow 1}$, $\text{ra}_{2 \rightarrow 1}$, $\text{la}_{3 \leftarrow 1}$ and $\text{ra}_{3 \rightarrow 1}$.*

These systems are monotonic for every value of d , since for a transition $\text{la}_{p \leftarrow 1}$, we have that $\mu(\text{la}_{p \leftarrow 1}) = \{\text{la}_{p-1 \leftarrow 1}\}$ (if $p > 2$) or \emptyset (otherwise), and therefore $\mu(\text{la}_{p \leftarrow 1})$ is included in T . The same also holds for $\mu(\text{ra}_{p \rightarrow 1})$.

The monotonicity property is crucial for the main result of this paper: if a bottom-up shift-reduce transition system is monotonic, we can systematically obtain an equivalent system without spurious ambiguity, as described in the next section.

3 Removal of Spurious Ambiguity

Let S be a bottom-up shift-reduce transition system that is monotonic. We show how we can systematically obtain a new transition system S' without spurious ambiguity that is equivalent to S , that is, S' parses the same set of trees as S . In essence, this is the main result of this paper, which can be formally stated as follows:

Theorem 3. Any transition system S which is bottom-up shift-reduce and monotonic, can always be converted into an equivalent transition system S' that does not have spurious ambiguity, such that:

- (i) for each complete computation γ' of S' on w there is a complete computation γ of S such that $D(\gamma) = D(\gamma')$; and
- (ii) for each complete computation γ of S on w there is a complete computation γ' of S' such that $D(\gamma) = D(\gamma')$.

Next, we describe how S' is created, and give full formal proofs of this theorem in §4.

3.1 Stack Symbols

Recall that in S each stack symbol is an integer i representing the word occurrence a_i in the input string. Each stack symbol in S' is obtained by annotating i with the following Boolean features:

- a feature $i.\text{stop}$ indicating whether, in the current analysis, the word a_i has collected all of its dependents (T) or it is still seeking some of them (F);
- for each k with $1 \leq k \leq \text{deg}(S)$, a feature $i.\text{left}_k$ indicating that a left reduction is allowed (T) or forbidden (F) between symbol i and the symbol k positions below i in the stack;
- for each k with $1 \leq k \leq \text{deg}(S)$, a feature $i.\text{right}_k$ indicating that a right reduction is allowed (T) or forbidden (F) between symbol i and the symbol k positions below i in the stack.

We now introduce some predicates that will be used later to define the new transition system S' . Let i and j be stack symbols of S' . The predicate $\text{bu}(i, j) \equiv \neg i.\text{stop} \wedge j.\text{stop}$ indicates whether a bottom-up link from node i to node j is admissible in the current configuration, i.e., whether node i can accept a dependent and node j has already collected all of its dependents. Assume that i and j are located at stack positions p and q , respectively, with $p > q$. Then the predicates²

$$\begin{aligned} \text{left}(i, j; p, q) &\equiv j.\text{left}_{(p-q)} \wedge \text{bu}(j, i) \wedge \text{la}_{p \leftarrow q} \in T, \\ \text{right}(i, j; p, q) &\equiv j.\text{right}_{(p-q)} \wedge \text{bu}(i, j) \wedge \text{ra}_{p \rightarrow q} \in T \end{aligned}$$

indicate that reductions $\text{la}_{p \leftarrow q}$ and $\text{ra}_{p \rightarrow q}$, respectively, are **available** in the current configuration, i.e., these reductions can be performed by the parser. As we will see later, the notion of available reduction plays a crucial role in the construction of S' .

3.2 Transitions

The basic idea underlying the construction of S' is to perform reductions as early as they become available in a computation, according to the notion of available reduction that we have just introduced. This is implemented as follows.

We define a **priority** relation among transitions in T such that, in choosing between several reductions that are compatible with some dependency tree, we give highest priority to the reduction with its dependent closest to the top of the stack. This reduction is necessarily unique, given that in a dependency tree each dependent has a unique head. The shift transitions are always assigned the lowest priority.

Note that the priority relation can be seen as a partial order between reductions, but the set of reductions that are compatible with a given tree is totally ordered, due to the restriction that a node cannot have more than one head.

In the new transition system S' we simulate S as follows. Given a configuration c'_1 of S' representing a configuration c_1 of S , we consider the set $T_{c'_1}$ of all transitions from S that are available at c'_1 . We nondeterministically choose a transition $t \in T_{c'_1}$ and simulate it on c'_1 under S' , moving into a new configuration c'_2 . Most important, in c'_2 we set the

²Here we are overloading symbols left and right , with related meanings: it will always be clear from the context whether these symbols refer to features or else to predicates.

feature of the stack symbols in such a way that all transitions in T_{c_1} that had higher priority than t are now blocked, meaning that no computation spanning from c'_2 will ever be able to apply such transitions. We can now specify our construction.

For a stack symbol i of S , we write $i[\mathbb{T}]$ to denote the stack symbol of S' such that $i.\varphi = \mathbb{T}$ for every feature φ . For a feature φ and a value v , we write $j = i[\varphi \leftarrow v]$ if $j.\varphi = v$ and $j.\varphi' = i.\varphi'$ for every other feature φ' . We generalize this notation to a set of features \mathcal{F} , and write $j = i[\varphi \leftarrow v \mid \varphi \in \mathcal{F}]$ if $j.\varphi = v$ for each $\varphi \in \mathcal{F}$ and $j.\varphi = i.\varphi$ for each $\varphi \notin \mathcal{F}$. Finally, as a shorthand, we write $i[\varphi \leftarrow v \mid \varphi \in \mathcal{F}; \varphi' \leftarrow v' \mid \varphi' \in \mathcal{F}']$ in place of $(i[\varphi \leftarrow v \mid \varphi \in \mathcal{F}])[\varphi' \leftarrow v' \mid \varphi' \in \mathcal{F}']$.

The system S' obtained by removing spurious ambiguity from S has a set of transitions T' including all and only the transitions reported below, where δ is $\text{depth}(S)$:

$$\begin{aligned} \text{sh}^s : (\sigma | i_\delta | i_{\delta-1} | \dots | i_1, i | \beta, A) \vdash \\ (\sigma | i'_\delta | i'_{\delta-1} | \dots | i'_1 | \beta, A) \end{aligned}$$

where we let $i' = i[\mathbb{T}]$, and for every u with $1 \leq u \leq \delta$ we let

$$\begin{aligned} i'_u = i_u[\text{left}_k \leftarrow \mathbb{F} \mid \text{left}(i_{u+k}, i_u; u+k, u); \\ \text{right}_k \leftarrow \mathbb{F} \mid \text{right}(i_{u+k}, i_u; u+k, u)]. \end{aligned}$$

Transition sh^s simulates a shift of S . The superscript s means that the new symbol i' added to the stack has the feature stop set to \mathbb{T} , that is, we (nondeterministically) guess that i' is now ready for bottom-up reduction. Since the shift transition has always the lowest priority in S , sh^s blocks any reduction that was available in the antecedent configuration, by setting the features of each i'_q , as indicated above.

We also add to T' a transition $\text{sh}^{\bar{s}}$ defined exactly as sh^s but with the only difference that we let $i' = (i[\mathbb{T}])[\text{stop} \leftarrow \mathbb{F}]$, that is, we guess that node i' is still seeking dependents in the current analysis.

For each $\text{ra}_{p \rightarrow q}$ in T , we add to T'

$$\begin{aligned} \text{ra}_{p \rightarrow q}^s : (\sigma | i_p | i_{p-1} | \dots | i_1, \beta, A) \vdash \\ (\sigma | i'_p | \dots | i'_{q+1} | i'_{q-1} | \dots | i'_1, \beta, A \cup \{i_p \rightarrow i_q\}) \end{aligned}$$

which can only be applied under the precondition $\text{right}(i_p, i_q; p, q)$. Here we let $i'_p = i_p[\text{stop} \leftarrow \mathbb{T}]$, and for every u with $1 \leq u \leq d$ we let

$$\begin{aligned} i'_u = i_u[\text{left}_k \leftarrow \mathbb{F} \mid u+k < q \wedge \text{left}(i_{u+k}, i_u; u+k, u); \\ \text{right}_k \leftarrow \mathbb{F} \mid u < q \wedge \text{right}(i_{u+k}, i_u; u+k, u)]. \end{aligned}$$

As for the shift transition, we also add to T' a transition $\text{ra}_{p \rightarrow q}^{\bar{s}}$ defined exactly as $\text{ra}_{p \rightarrow q}^s$ but with $i'_p = i_p[\text{stop} \leftarrow \mathbb{F}]$. Reductions $\text{ra}_{p \rightarrow q}^s$ and $\text{ra}_{p \rightarrow q}^{\bar{s}}$ block every reduction t allowable in the antecedent configuration that has priority higher than the reduction $p \rightarrow q$, that is, with a dependent at a position closer to the top of the stack than q .

Similarly to the above, for each $\text{la}_{p \leftarrow q}$ in T we add to T'

$$\begin{aligned} \text{la}_{p \leftarrow q}^s : (\sigma | i_p | i_{p-1} | \dots | i_1, \beta, A) \vdash \\ (\sigma | i'_{p-1} | \dots | i'_1, \beta, A \cup \{i_q \rightarrow i_p\}) \end{aligned}$$

which can only be applied under the precondition $\text{left}(i_p, i_q; p, q)$. Here we let $i'_q = i_q[\text{stop} \leftarrow \mathbb{T}]$, and for every u with $1 \leq u \leq d$ we let

$$\begin{aligned} i'_u = i_u[\text{left}_k \leftarrow \mathbb{F} \mid u+k < p \wedge \text{left}(i_{u+k}, i_u; u+k, u); \\ \text{right}_k \leftarrow \mathbb{F} \mid u < p \wedge \text{right}(i_{u+k}, i_u; u+k, u)]. \end{aligned}$$

We also add to T' a transition $\text{la}_{p \leftarrow q}^{\bar{s}}$ defined exactly as $\text{la}_{p \leftarrow q}^s$ but with $i'_q = i_q[\text{stop} \leftarrow \mathbb{F}]$.

The initialization function and final configuration set of S' are like those of S , but we have to specify feature values for the stack symbol corresponding to the dummy root node 0: all its features will be \mathbb{F} in the initial configuration, and in final configurations it must have the left_k and right_k features set to \mathbb{F} but stop set to \mathbb{T} .

Example 4. If we apply the transformation defined in this section to remove spurious ambiguity from the arc-standard transition system of Example 1, we obtain a system S' where the only valid computation for the tree with arcs $\{0 \rightarrow 2, 2 \rightarrow 1, 2 \rightarrow 3\}$ is

$$(I(w); \text{sh}^s, \text{sh}^{\bar{s}}, \text{la}_{2 \leftarrow 1}^{\bar{s}}, \text{sh}^s, \text{ra}_{2 \rightarrow 1}^s, \text{ra}_{2 \rightarrow 1}^{\bar{s}}),$$

which builds the arcs in the same order as the computation (i) of Example 1.

It is easy to check that an alternate computation building the arcs in the order of the computation (ii) does not exist in S' . Such a computation would have to start with the transitions $\text{sh}^s, \text{sh}^{\bar{s}}, \text{sh}^s, \text{ra}_{2 \rightarrow 1}^{\bar{s}}$ (the need to use the s or \bar{s} variant of each configuration is uniquely determined by whether nodes have pending dependents or not).

However, after applying these transitions the parser will be in a configuration $([0, 1, 2], [], \emptyset)$ with:

$$\begin{aligned} 0.\text{stop} &= F, 0.\text{left}_1 = F, 0.\text{right}_1 = F, \\ 1.\text{stop} &= T, 1.\text{left}_1 = T, 1.\text{right}_1 = F, \\ 2.\text{stop} &= F, 2.\text{left}_1 = F, 2.\text{right}_1 = T. \end{aligned}$$

At such configuration, the feature value $2.\text{left}_1 = F$ blocks the left reduction creating the arc $2 \rightarrow 1$. This is so because the sh^s transition that moved the node 3 to the stack set this value to F , blocking this left reduction since it could have been executed at that point with higher priority than sh^s .

4 Formal Properties and Proofs

We now proceed to prove that the described transformation for the removal of spurious ambiguity is correct (i.e. prove Theorem 3). To do so, we first show that transition systems S and S' defined as in §3 are equivalent, i.e., they assign the same set of trees to any input string. Afterward, we show that S' has no spurious ambiguity, i.e., different complete computations of S' will always produce different dependency trees.

4.1 Equivalence of Unambiguous System to Original System

Let S and S' be defined as in Section 3, with associated transition sets T and T' , respectively. To show that S and S' are equivalent, we need to prove that for every input string w

- (i) for each complete computation γ' of S' on w there is a complete computation γ of S such that $D(\gamma) = D(\gamma')$; and
- (ii) for each complete computation γ of S on w there is a complete computation γ' of S' such that $D(\gamma) = D(\gamma')$.

The proof of (i) is rather straightforward. We show a mapping from the complete computations of S' to the complete computations of S that preserves the associated trees. We define a homomorphism τ from T' to T by letting

$$\begin{aligned} \tau(\text{la}_{p \leftarrow q}^s) &= \tau(\text{la}_{p \leftarrow q}^{\bar{s}}) = \text{la}_{p \leftarrow q}, \\ \tau(\text{ra}_{p \rightarrow q}^s) &= \tau(\text{ra}_{p \rightarrow q}^{\bar{s}}) = \text{ra}_{p \rightarrow q}, \\ \tau(\text{sh}^s) &= \tau(\text{sh}^{\bar{s}}) = \text{sh}, \end{aligned}$$

and extend it to (complete) computations (recall that we represent a computation by its initial configuration and its sequence of transitions) by letting $\tau((c_0; t_1, \dots, t_m)) = (c_0; \tau(t_1), \dots, \tau(t_m))$.

It is not difficult to see that if γ is complete, then $\tau(\gamma)$ is also complete. Furthermore, this mapping preserves trees, i.e., for any computation γ of S' we have $D(\gamma) = D(\tau(\gamma))$, because transitions $t \in T'$ and $\tau(t) \in T$ create the same arc, if any. This concludes the proof of (i).

To prove statement (ii) above, let $\gamma = c_0, \dots, c_m = (c_0; t_1, \dots, t_m)$ be a complete computation of S for an input string w , and let A_γ be the set of arcs in $D(\gamma)$. We show that we can always find a computation γ' of S' such that

$D(\gamma') = D(\gamma)$. To do this, we introduce below the notion of canonical computations of S . Then we proceed in two steps: first we transform γ into a canonical computation γ_f of S equivalent to γ , and then we transform γ_f into an equivalent computation γ' of S' .

Consider a configuration c_k , $0 \leq k \leq m$, appearing in γ . Let $\mathcal{R}_{k,\gamma}$ be the set of reductions of S that can be applied to c_k , and that are compatible with $D(\gamma)$, i.e., these reductions construct an arc $(h \rightarrow d) \in A_\gamma$. Here a_h is the head word and a_d is the dependent word, both from w .

Assume that $\mathcal{R}_{k,\gamma} \neq \emptyset$, and let t_ρ be the reduction in $\mathcal{R}_{k,\gamma}$ with the highest priority. This means that t_ρ is the reduction in $\mathcal{R}_{k,\gamma}$ with dependent node d placed at the position closest to the top in the stack associated with c_k or, equivalently, the reduction with the largest value of index d in w . Note that there cannot be more than one such reduction, due to the single-head constraint in $D(\gamma)$.

We say that c_k is a **troublesome** configuration in γ if $t_{k+1} \neq t_\rho$. This means that t_{k+1} is either a shift transition, or else a reduction in $\mathcal{R}_{k,\gamma}$ that, when applied to c_k , creates a dependency link $h' \rightarrow d'$ with $d' < d$, i.e., a reduction with lower priority than t_ρ , since node d' will be placed at a deeper position than node d in the stack associated with c_k .

We say that a computation of S is in **canonical** form if it does not contain any troublesome configuration. This means that, at each configuration c_k of a canonical computation, the reduction in $\mathcal{R}_{k,\gamma}$ with the highest priority is taken, in case set $\mathcal{R}_{k,\gamma}$ is not empty. We now show that for every computation γ of S there exists an equivalent canonical computation γ_f of S . We show how to eliminate the leftmost troublesome configuration in γ ; iteration of this process will always produce a computation where no configurations are troublesome.

Let c_k be the leftmost troublesome configuration in γ . We show that we can build a computation γ_k of S which is equivalent to γ , and such that its first k configurations are not troublesome. The transition sequence t_{k+1}, \dots, t_m can be written in the form

$$t_{k+1}, t_{k+2}, \dots, t_{j-1}, t'_\rho, t_{j+1}, \dots, t_m$$

where t'_ρ is a reduction creating the same link $h \rightarrow d$ that should have been created by the reduction $t_\rho \in \mathcal{R}_{k,\gamma}$ with the highest priority. Note that reduction t'_ρ must take place at some c_j in γ with $j > k + 1$, because $h \rightarrow d$ is in $D(\gamma)$, and this link cannot be present in the arc set associated with c_k (if it were, the reduction t_ρ could not be available at c_k because d would not be in the stack at that configuration).

The sequence t_{k+1}, \dots, t_m in γ can then be replaced (generating the same tree) with

$$t_\rho, \tau_d(t_{k+1}), \dots, \tau_d(t_{j-1}), t_{j+1}, \dots, t_m$$

where $\tau_d(t)$ represents the transition that creates the same arc in a stack where the node j has been removed as t would create in a stack where the node j is present. Formally, for a transition applied at a configuration c with stack $\sigma|i_p| \dots |i_q| \dots |i_1$, we define $\tau_d(\text{sh}) = \text{sh}$ and

$$\tau_d(\text{ra}_{p \rightarrow q}) = \begin{cases} \text{ra}_{p \rightarrow q} & \text{if } i_p > d \text{ and } i_q > d, \\ \text{ra}_{p-1 \rightarrow q} & \text{if } i_p < d \text{ and } i_q > d, \\ \text{ra}_{p-1 \rightarrow q-1} & \text{if } i_p < d \text{ and } i_q < d. \end{cases}$$

$$\tau_d(\text{la}_{p \leftarrow q}) = \begin{cases} \text{la}_{p \leftarrow q} & \text{if } i_p > d \text{ and } i_q > d, \\ \text{la}_{p-1 \leftarrow q} & \text{if } i_p < d \text{ and } i_q > d, \\ \text{la}_{p-1 \leftarrow q-1} & \text{if } i_p < d \text{ and } i_q < d. \end{cases}$$

Note that, since S is monotonic, the existence of a transition t implies the existence of $\tau_d(t)$.

The computations γ_k and γ produce the same tree. Also, in γ_k the first k configurations are not troublesome, since applying the reduction t_ρ at c_k makes c_k not troublesome, and by construction the configurations to the left of c_k in γ_k are not troublesome.

By iteratively applying the above process, we eventually obtain a computation γ_f of S such that $D(\gamma_f) = D(\gamma)$. It then remains to show that we can obtain a computation γ' of S' with the same associated dependency tree as γ_f .

Let $\gamma_f = (c_0; t_1, \dots, t_m)$ and assume that for each j , $1 \leq j \leq m$, transition t_j in γ_f applies to configuration $c_{j-1} = (\sigma|i_p| \dots |i_q| \dots |i_1, i_0|\beta, A)$. The computation γ' is obtained as $\gamma' = (c_0; t'_1, \dots, t'_m)$, where for each j , t'_j is specified as follows.

- If $t_j = \text{ra}_{p \rightarrow q}$, then t'_j is $\text{ra}_{p \rightarrow q}^{\bar{s}}$ if $A_\gamma \setminus (A \cup \{(i_p, i_q)\})$ contains a dependency link of the form (i_p, u) for some u , and t'_j is $\text{ra}_{p \rightarrow q}^s$ otherwise.
- If $t_j = \text{la}_{p \leftarrow q}$, then t'_j is $\text{la}_{p \leftarrow q}^{\bar{s}}$ if $A_\gamma \setminus (A \cup \{(i_q, i_p)\})$ contains a dependency link of the form (i_q, u) for some u , and t'_j is $\text{la}_{p \leftarrow q}^s$ otherwise.
- If $t_j = \text{sh}$, then t'_j is $\text{sh}^{\bar{s}}$ if $A_\gamma \setminus A$ contains a dependency link of the form (i_0, u) for some u , and t'_j is sh^s otherwise.

It is not difficult to see that γ' is a valid computation of S' for w . This follows from the fact that the transitions t'_j above satisfy the $\text{bu}(i, j)$ predicates in S' , and the fact that in γ_f reductions are applied in accordance to the priority relation. We also observe that if γ_f is complete then γ' is complete as well. Finally, the fact that $D(\gamma') = D(\gamma_f)$ follows immediately from the above mapping from transitions t_j to transitions t'_j . This concludes the proof of (ii) and thus the proof of the equivalence of S and S' .

4.2 Non-ambiguity of the Transition System

To prove that our transformed system S' has no spurious ambiguity, we need to show that different complete computations of S' for w always produce different trees, i.e., if $\gamma_1 \neq \gamma_2$ are complete computations of S' for input string w , then $D(\gamma_1) \neq D(\gamma_2)$.

To do so we write γ_1 as $\alpha c_1 \beta_1$ and γ_2 as $\alpha c_2 \beta_2$, with α the common prefix among both computations, and c_1, c_2 configurations such that $c_1 \neq c_2$. Note that α cannot be empty, since both computations must at least have the initial configuration $I(w)$ in common. We call c_0 the last configuration in α , and t_1, t_2 the transitions that produce c_1, c_2 (respectively) from c_0 . We distinguish four cases below.

Case 1: t_1 and t_2 are transitions that differ only in the stop feature of some new node u in the configuration they produce. As an example, we have $t_1 = \text{la}_{p \leftarrow q}^s$ and $t_2 = \text{la}_{p \leftarrow q}^{\bar{s}}$, which differ in the stop feature of node $u = q$. Without loss of generality, we assume $u.\text{stop} = \text{T}$ in c_1 , and $u.\text{stop} = \text{F}$ in c_2 . Let $c_0 = (\sigma, \beta, A)$. Then $D(\gamma_2)$ must contain at least one arc originating from u that is not present in A , while $D(\gamma_1)$ cannot contain any arc originating from u that is not already in A , because $u.\text{stop} = \text{T}$ prevents the addition of dependents of u after t_1 is executed. Therefore, $D(\gamma_1) \neq D(\gamma_2)$.

Case 2: t_1 and t_2 are reduce transitions with different head nodes but the same dependent node u . In this $D(\gamma_1) \neq D(\gamma_2)$ follows from the single-head constraint, since the node u will be assigned different heads in γ_1 and γ_2 , respectively.

Case 3: t_1 and t_2 are reduce transitions involving different dependent nodes. Suppose that t_1 creates the arc $h_1 \rightarrow d_1$ and t_2 creates the arc $h_2 \rightarrow d_2$. Without loss of generality, we assume that $d_1 > d_2$, i.e., t_1 has higher priority than t_2 . Then $D(\gamma_1)$ contains the arc $h_1 \rightarrow d_1$, but $D(\gamma_2)$ cannot contain this arc, since the system's features block its construction after the application of the transition t_2 at configuration c_0 .

Case 4: t_1 is a reduce transition and t_2 is a shift transition. The same reasoning of Case 3 applies: the arc $h_1 \rightarrow d_1$ created by t_1 cannot appear in $D(\gamma_2)$, because the system's features block its construction after the shift transition t_2 is applied. This concludes the proof that S' does not have spurious ambiguity.

4.3 Complexity

Let S be a bottom-up monotonic transition system, and let $\text{deg}(S) = \delta$. The construction in §3 adds $2\delta + 1$ binary features to each stack symbol of S . This results in $2^{2 \cdot \delta + 1}$ new symbols in S' for each stack symbol of S . While for projective dependency parsing we have $\delta = 1$, degree larger than one is needed in non-projective parsing. However, it has been observed by Attardi [2006] that most of the non-projective trees in the CoNLL data can be parsed with $\delta = 2$ or 3. This means that, in practical cases, the blow-up of stack symbols by our construction can be considered a small constant.

To discuss a concrete application, consider the non-projective system S of [Attardi, 2006], also shown in Example 2, restricted to $\delta = 2$, which is still heavily affected by spurious ambiguity. We have applied the construction in §3 to S with some ad-hoc optimization of the features for that system, resulting in a new system S' with a blow-up of

Language	Size	Attardi	This paper
Arabic	1,460	27	2
Bulgarian	12,823	47	36
Czech	72,703	1,334	602
Danish	5,190	179	159
Dutch	13,349	1,448	1,018
German	39,216	2,140	1,538
Japanese	17,044	121	45
Portuguese	9,071	295	203
Slovene	1,534	48	27
Spanish	3,306	11	10
Swedish	11,042	197	105
Turkish	4,997	208	102

Table 1: Coverage of Attardi’s oracle versus the coverage of our oracle for various treebanks from the CoNLL 2006 data sets [Buchholz and Marsi, 2006]. “Size” denotes the number of sentences in the treebank (we used the training portion only), “Attardi” denotes the number of sentences that Attardi’s oracle could not parse and “this paper” denotes the number of parse trees that our oracle could not parse.

stack symbols of $2^{\delta+1} = 8$. This means that we can apply to S' the inside/outside algorithm presented in Cohen et al. [2011], working in time $\mathcal{O}(|w|^7)$ for an input string w , with an extra hidden constant of 8.

5 Experiments

As mentioned earlier, transition-based dependency parsing uses an oracle to convert training data which consists of pairs of sentences and dependency trees to pairs of sentences with shift-reduce sequences, in order to sidestep the issue of spurious ambiguity. The new training data is then used to train multi-class classifiers. In several cases, oracles are based on heuristics and are incomplete. The oracle that is provided in the DeSR dependency parsing package,³ which is based on the parser from Attardi [2006], is an example for such incomplete heuristics.

We compared the coverage of Attardi’s oracle, restricted to transitions of degree at most 2, to the oracle of an equivalent transition system without spurious ambiguity.⁴ Our findings are given in Table 1. As theoretically guaranteed, there were no cases where Attardi’s parser recognized a tree using transitions of degree 2, and our oracle did not recognize it. The reverse, however, holds quite often.

6 Discussion

We note that monotonic bottom-up shift-reduce transition systems can be made probabilistic and generative, in a manner similar to Cohen et al. [2011]. The issue with spurious ambiguity is especially crucial with generative models in the unsupervised setting, when using algorithms such as the expectation-maximization (EM) algorithm. Cohen et al. [2011] describe an EM algorithm for the system from Attardi [2006], which can be extended to any monotonic bottom-up transition system. The EM algorithm they describe can be further extended to monotonic bottom-up transition systems after removal of spurious ambiguity (as we describe in this paper), making these systems readily available for transition-based unsupervised learning for dependency parsing.

³<http://desr.sourceforge.net/>.

⁴Note that the algorithm implemented in the latest version of DeSR, which we used for these experiments, differs from the description provided in Attardi [2006] and Example 2 in that $la_{3\leftarrow 1}$ and $ra_{3\rightarrow 1}$ transitions push a node from the stack back to the buffer after reducing. This does not affect our method to remove spurious ambiguity, which is correct both for the version described in Attardi [2006] and for the latest implementation of Attardi’s parser.

7 Conclusion

We provided a principled treatment to the issue of spurious ambiguity in transition-based dependency parsing. We defined a large class of transition systems, which we call monotonic bottom-up shift-reduce transition systems, that cover existing systems such as the arc-standard parser of Nivre [2008] and the non-projective parser of Attardi [2006], as well as systems in which reductions affect elements at positions in the stack deeper than the topmost element [Goldberg and Elhadad, 2010]. We then showed how to eliminate spurious ambiguity from these systems. Our technique has applications for unsupervised and supervised dependency parsing. The transition model that we present can be used as a substitute for models such as the dependency model with valence that have long been used for dependency grammar induction [Klein and Manning, 2004, Cohen and Smith, 2010, Spitkovsky et al., 2010].

In this paper we have discovered some sufficient conditions under which spurious ambiguity can be removed from bottom-up dependency transition systems, which we hope are as “tight” as possible. However, our technique does not work for all dependency transition systems, and it remains an open problem to show whether removal of spurious ambiguity can be carried out in the general case. There might as well be dependency parsing strategies for which removal of spurious ambiguity is not only difficult, but simply impossible. A similar scenario is observed, for instance, for structural ambiguity in context-free grammars, where some context-free languages can only be generated using ambiguous context-free grammars; see for instance Hopcroft et al. [2006].

References

- Giuseppe Attardi. Experiments with a multilanguage non-projective dependency parser. In *Proceedings of the Tenth Conference on Computational Natural Language Learning (CoNLL)*, pages 166–170, 2006.
- Rens Bod. A computational model of language performance: Data oriented parsing. In *Proceedings of COLING*, 1992.
- Sabine Buchholz and Erwin Marsi. CoNLL-X shared task on multilingual dependency parsing. In *Proceedings of the Tenth Conference on Computational Natural Language Learning (CoNLL)*, pages 149–164, 2006.
- Shay B. Cohen and Noah A. Smith. Covariance in unsupervised learning of probabilistic grammars. *Journal of Machine Learning Research*, 11:3017–3051, 2010.
- Shay B. Cohen, Carlos Gómez-Rodríguez, and Giorgio Satta. Exact inference for generative probabilistic non-projective dependency parsing. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1234–1245, 2011.
- Edward Gibson. The dependency locality theory: A distance-based theory of linguistic complexity. In *Image, language, brain: Papers from the first mind articulation project symposium*, pages 95–126, 2000.
- Yoav Goldberg and Michael Elhadad. 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*, pages 742–750, Los Angeles, California, June 2010. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/N10-1115>.
- John E. Hopcroft, R. Motwani, and Jeffery D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison Wesley; 3rd edition, 2006.
- Aravind K. Joshi, Leon S. Levy, and Masako Takahashi. Tree Adjunct Grammars. *Journal of Computer and System Sciences*, 10(2):136–163, 1975.
- Dan Klein and Christopher D. Manning. Corpus-based induction of syntactic structure: Models of dependency and constituency. In *Proceedings of the 42th Annual Meeting of the Association of Computational Linguistics (ACL)*, 2004.

- Marco Kuhlmann and Joakim Nivre. Transition-based techniques for non-projective dependency parsing. *Northern European Journal of Language Technology*, 2(1):1–19, 2010.
- Marco Kuhlmann, Carlos Gómez-Rodríguez, and Giorgio Satta. Dynamic programming algorithms for transition-based dependency parsers. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics (ACL)*, Portland, Oregon, USA, 2011.
- Joakim Nivre. An efficient algorithm for projective dependency parsing. In *Proceedings of the Eighth International Workshop on Parsing Technologies (IWPT)*, pages 149–160, Nancy, France, 2003.
- Joakim Nivre. Incrementality in deterministic dependency parsing. In *Workshop on Incremental Parsing: Bringing Engineering and Cognition Together*, pages 50–57, Barcelona, Spain, 2004.
- Joakim Nivre. Dependency grammar and dependency parsing. Technical report, Växjö University, 2005.
- Joakim Nivre. Algorithms for deterministic incremental dependency parsing. *Computational Linguistics*, 34(4):513–553, 2008.
- Joakim Nivre, Johan Hall, and Jens Nilsson. Memory-based dependency parsing. In *Proceedings of the 8th Conference on Computational Natural Language Learning (CoNLL-2004)*, pages 49–56, Morristown, NJ, USA, 2004. Association for Computational Linguistics.
- Khalil Sima'an. Computational complexity of probabilistic disambiguation by means of tree-grammars. In *Proceedings of COLING*, pages 1175–1180, 1996.
- Valentin Spitzkovsky, Hiyam Alshawi, and Daniel Jurafsky. From baby steps to leapfrog: How “less is more” in unsupervised dependency parsing. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, pages 751–759, 2010.
- Mark Steedman. *The Syntactic Process*. MIT Press, 2001.