

Inherent Dependency Displacement Bias of Transition-Based Algorithms

Mark Anderson Carlos Gómez-Rodríguez

Universidade da Coruña, CITIC

FASTPARSE Lab, LyS Research Group, Departamento de Computación y Tecnologías de la Información

Campus Elviña, s/n, 15071

A Coruña, Spain

{m.anderson, carlos.gomez}@udc.es

Abstract

A wide variety of transition-based algorithms are currently used for dependency parsers. Empirical studies have shown that performance varies across different treebanks in such a way that one algorithm outperforms another on one treebank and the reverse is true for a different treebank. There is often no discernible reason for what causes one algorithm to be more suitable for a certain treebank and less so for another. In this paper we shed some light on this by introducing the concept of an algorithm’s inherent dependency displacement distribution. This characterises the bias of the algorithm in terms of dependency displacement, which quantify both distance and direction of syntactic relations. We show that the similarity of an algorithm’s inherent distribution to a treebank’s displacement distribution is clearly correlated to the algorithm’s parsing performance on that treebank, specifically with highly significant and substantial correlations for the predominant sentence lengths in Universal Dependency treebanks. We also obtain results which show a more discrete analysis of dependency displacement does not result in any meaningful correlations.

Keywords: dependency parsing, transition-based algorithms, dependency displacement

1. Introduction

Dependency parsing, and in particular the transition-based family of parsers, has a large variety of parsing algorithms to choose from. When comparing different algorithms, empirical results on collections of corpora often show differences in accuracy that can heavily vary across different languages or treebanks, so that a given algorithm can be the best choice for one corpus while being outperformed in another.

However, despite these nontrivial patterns in accuracy variations having been observed in many experiments in the last decade, both with non-neural and neural implementations (Nivre, 2008; Ballesteros and Nivre, 2013; Chen and Manning, 2014; Fernández-González et al., 2016), very little is known about what makes an algorithm more fitting for a corpus beyond obvious facts (like non-projective algorithms being better for highly non-projective treebanks). This makes the practical choice of a particular algorithm for a language a matter of trial and error. For example, Malt-Optimizer chooses between projective and non-projective algorithm according to the amount of non-projective dependencies observed in the treebank, but then the choice of a specific algorithm among projective (or non-projective) options is made by running all of the projective (or the non-projective) algorithms and choosing the one that achieves the highest accuracy (Ballesteros and Nivre, 2012).

Contribution We shed some light on this issue by performing two complementary analyses:

1. An analysis of the performance of transition-based algorithms with respect to dependency displacement (signed distance) which highlights that a shallow level of analysis is not sufficient to account for the aforementioned differences.
2. An analysis focusing on the distribution of dependency displacements which shows the performance of a transition-based algorithm on a given corpus can be

partly explained by the similarity of the dependency displacement distribution of the corpus (sampled from the test set) and what we call the algorithm’s *inherent distribution*. This is a random distribution that describes the bias of the algorithm in terms of dependency lengths and directions, and it can be approximated by running the algorithm in random mode. This entails assigning a uniform probability to each of the transitions available for a given configuration.

Experiments with Universal Dependencies show a highly significant correlation between said similarity and the accuracy of algorithms on different treebanks for the most prevalent sentence lengths in the corpora.

2. Related work

2.1. How language-specific performance differs across algorithms

With the proliferation of dependency treebanks for multiple languages, various papers performed comparisons including multiple transition-based parsing algorithms, where language-specific differences in performance across algorithms are apparent.

One of the first papers that included a large number of languages and algorithms is that of Nivre (2008), who found language-specific differences between the performance of Arc-Standard and Arc-Eager. While they hypothesized that the proportion of left arcs in a language could account for these differences, the evidence was not conclusive.

Other analyses have found differences in performances between Arc-Eager and Arc-Standard such as Ballesteros and Nivre (2013) or Fernández-González et al. (2016), but they provided no explanation for this phenomenon. Ballesteros and Nivre (2013) did consider the effect the dummy root placement has on different algorithms. They found that the placement is not trivial and had a noticeable effect on the performance of Arc-Eager.

The differences between algorithms have also been observed with different architectural implementations, namely neural networks. Chen and Manning (2014) found that for some treebanks Arc-Eager performed better whereas for others Arc-Standard performed better.

2.2. Analysis of parsing errors

There have been analyses on the strengths and weaknesses of different parsers by comparing the errors they make on different kinds of dependencies, with dependency distance and direction often being considered.

The most comprehensive analysis on the performance of dependency parsers came from McDonald and Nivre (2011), who compared a graph-based parser with a transition-based parser. They investigated the relative strengths of each of the parsing paradigms with respect to factors like dependency length, distance to root, and sentence length. They found that the transition-based parser performed worse for longer dependencies and those closer to the root, due to error propagation, and better in the converse cases. However, while this analytical methodology gave interesting insights when comparing parsers from very different paradigms, it is not fine-grained enough to draw conclusions when the comparison is between transition-based algorithms, as we will verify in Section 4.

de Lhoneux et al. (2017) investigated the differences in performances when using the same transition-based algorithm with a neural network implementation, namely UD-Pipe (Straka and Straková, 2017), and with a classical implementation, MaltParser (Nivre et al., 2007). They observed a similar trend for both implementations with respect to the change in F1 score for different dependency distance bins.

More recently, Kulmizev et al. (2019) evaluated the difference in performance between graph-based and transition-based parsers and observed that the use of contextualised word embedding off-set the typical error-propagation found when using transition-based parsers.

Rehbein et al. (2017) investigated what made certain corpora harder to parse than others by looking at how the dependency encoding scheme affects dependency length and arc direction entropy. They, however, did not observe a correlation in parsing accuracy for different dependency lengths resulting from different encoding styles, but they did find a correlation with respect to arc direction entropy. This is somewhat corroborated by a previous analysis undertaken by Gulordava and Merlo (2016), who modified treebanks to minimise dependency length and the arc direction entropy and obtained an increase in parser performance using these treebanks.

Additionally, Kirilin and Versley (2015) evaluated error patterns (over- and underproduction of certain structures); Goldberg and Elhadad (2010) made a similar analysis by classifying parser output and test sentences based on underproduced (indicating test) and overproduced (indicating prediction) structures; Dickinson (2010) also showed that algorithms generate idiosyncratic structures not found in the training data; and Schwartz et al. (2012) found that certain structures work better with certain heads. Together these findings suggest that algorithms have some sort of

bias towards generating certain types of structures.

2.3. Dependency distance

There has also been more linguistic-specific work undertaken with respect to dependency distance, which can be relevant for parsing. For example, Gibson (2000) introduced *dependency locality theory* (DLT) which postulates that dependency distances are minimised in order to make parsing linguistic input more cognitively friendly. Liu et al. (2017) tested this theory by evaluating how syntactic patterns are governed by dependency distance. They provided a review of psycholinguistic experiments and treebank analyses that suggest natural languages have a tendency to minimize dependency distance, arguing that this makes sense under working-memory restrictions. So if longer distance dependencies are difficult for humans to parse, it would make sense then that parsing algorithms would struggle in a similar vein.

Many analyses have observed this tendency to minimise dependency distance, further corroborating this hypothesis (Ferrer-i-Cancho, 2004; Liu, 2008; Liu, 2007; Buch-Kromann, 2006; Futrell et al., 2015; Temperley and Gildea, 2018). However, the extent to which this restriction is adhered has been observed to vary significantly across languages (Jiang and Liu, 2015; Gildea and Temperley, 2010). Other works have investigated different syntactic traits of languages associated with dependency length such as highlighting a correlation with an increase in dependency length and free-order languages (Gulordava and Merlo, 2015) and with an increase in non-projective dependencies (Ferrer-i-Cancho and Gómez-Rodríguez, 2016; Gómez-Rodríguez and Ferrer-i-Cancho, 2017).

In response to Liu et al. (2017), Gómez-Rodríguez (2017) hypothesised that the good practical results achieved by transition-based algorithms could be because they are biased towards short dependencies. This is substantiated by Eisner and Smith (2010) who improved parser performance by imposing upper bounds on dependency length and using dependency lengths as a parsing feature. In another response to Liu et al. (2017), Hudson (2017) highlighted mean dependency distance varies significantly between treebanks and that the direction of dependencies could impact parsing difficulty. Liu (2010) and Jiang and Liu (2015) actually found that dependency direction analysis can be used to typologically classify languages. Hence, for our analyses we use dependency displacement which quantifies both length and direction as defined in Equation 1.

3. Data and setup

We report two levels of analysis to explain performance differences between different transition-based algorithms. The first is based on measuring the differences in performance for each dependency displacement. It transpires that this does not offer any explanation how language-specific performance differs between algorithms. For this, we need the second analysis, based on the concept of an inherent displacement distribution of each algorithm.

We evaluate 3 projective algorithms (Arc-Standard, Arc-Eager and Covington) and 2 unrestricted non-projective

algorithms (Covington and Swap-Eager). The results for each algorithm were obtained by running MaltParser v1.91, which has the benefit of having multiple algorithms implemented of which several are projective and several are non-projective. While there are more modern systems that outperform MaltParser in accuracy, none of them provide such a range of algorithms.¹ Furthermore, as we are focused on the algorithms, the architecture of the implementation is considered a controlled variable. Also, de Lhoneux et al. (2017) showed that the change in accuracy with respect to dependency length follows a similar trend for MaltParser and UD-Pipe (a neural network implementation), so it is justifiable to use MaltParser for analyses on algorithms. Furthermore, MaltParser is potentially better suited for this analysis as we want to evaluate the performance of algorithms against one another and using a more robust system that can more readily overcome the inherent biases we show in this paper would actually obscure this. For this reason too, we do not finetune the feature functions for each language as it is obvious that different features will benefit different languages more or less depending on the linguistic features of a given language. So these features are kept constant for each treebank and should be considered controlled variables for the following analyses.

Beyond these experimental considerations, MaltParser is still competitive with respect to parsing accuracy even if it is not state-of-the-art and it is very efficient which makes it much more readily deployable when compared to large and unwieldy neural networks.

Version 2.1 of the Universal Dependency treebanks was used for all of the subsequent analysis (Nivre et al., 2017).

4. Dependency displacement

4.1. Analysis details

Similar to the analysis undertaken by (McDonald and Nivre, 2011), we investigated how parsing performance varies based on the dependency displacement. Whereas they compared a graph-based and a transition-based parser, we have compared different transition-based parsing algorithms. Also, we look at the effect dependency displacement, s_{dep} , has on the performance of the algorithms, rather than dependency distance. Dependency displacement is defined as:

$$s_{\text{dep}} = x_{\text{head}} - x_{\text{dependent}} \quad (1)$$

where x_i refers to the position of word i in a given sentence. So a right arc of length 3 would have a displacement of -3 and conversely a left arc of length 3 would have a displacement of 3. Hudson (2017) highlighted that different languages have a tendency towards being head-final (OSV or SOV), head-medial (OVS or SVO), or head-initial (VOS or VSO) and argue that the direction of dependencies cannot be ignored when analysing dependency distances. In addition, dependency direction was hypothesised by Nivre (2008) to affect how language-specific performance differs between parsers.

¹For example, UDPipe has one projective, one partially non-projective and one unrestricted non-projective parser.

We evaluate the attachment precision and recall in order to have a more fine-grained analysis similar to de Lhoneux et al. (2017).

For this analysis we used all 76 treebanks that contained a training and test set.

4.2. Results

Figure 1 shows that this coarse analysis does not capture any statistically meaningful variation across projective algorithms with regard to attachment precision or recall. It is interesting to note that long-distant displacement to the right and to the left follow different trends for both precision and recall, with precision being higher for right arcs (negative displacement) but recall being higher for left arcs (positive displacement).

Figure 2 shows the corresponding results for the non-projective algorithms. We see in Figure 2a the only statistically meaningful difference between algorithms. Precision for left arc dependencies is higher for Swap-Eager than non-projective Covington. It is also of note that the non-projective precision results are much more symmetric with regard to dependency displacement than those of the projective algorithms. These results show that there is a need for a more fine-grained analysis.

5. Comparing displacement distributions

5.1. Analysis overview

In this analysis we test the hypothesis that the similarity of the dependency displacement distribution generated by the latent biases of an algorithm, as defined as its inherent displacement distribution in Section 5.2, and the actual distribution of a treebank can account for the difference in performance across parsing algorithms. We measure the difference between these distributions by using the Wasserstein distance - also known as the earth mover’s distance (EMD). It can be intuitively thought of as the cost or work required to change one distribution into another by moving distribution mass (Rubner et al., 1998). An example comparison for the German test data for sentences of length 10 to 12 tokens is shown in Figure 3, where Arc-Standard is seen to perform worse than Arc-Eager by 0.58 UAS and correspondingly has a higher EMD. We evaluate dependency displacements according to sentence-length bins, as it has been shown that dependency distance (and thus displacement) distributions change with sentence length, both in real sentences and in random models (Ferrer-i-Cancho and Liu, 2014).

5.2. Inherent displacement distributions

Let $P = (C, T, c_s, C_t)$ be a transition-based algorithm where C is the set of possible configurations, T the set of transitions, c_s an initialization function mapping a sentence length k to an initial configuration $c_s(k)$, and C_t a set of terminal configurations. We assume configurations in C to be of the form (D, A) , where A is a set of dependency arcs built so far, and D is a state of the data structures associated with the algorithm (e.g. a stack and a buffer, for stack-based algorithms).

Then, we define the *inherent distribution* of P for sentences of length k , written $\iota_k(P)$, as the discrete probability dis-

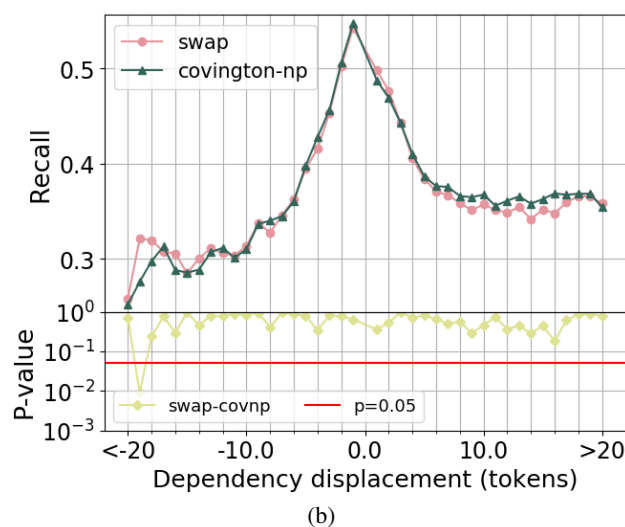





Figure 1: Attachment precision (a) and recall (b) for the three projective algorithms used: Arc-Eager (eager, blue, circle), Arc-Standard (standard, magenta, square), Covington (covington, green, triangle). The corresponding p-values (derived from a t-test respectively using the average precision and recall and the corresponding standard deviation across treebanks for each displacement value for each algorithm) are shown below: Arc-Eager and Arc-Standard (e-s); Arc-Eager and Covington (e-c); and Covington and Arc-Standard (c-s). No statistically significant differences are observed for any comparison with regard to attachment precision or recall.

Figure 2: Attachment precision (a) and recall (b) for the two non-projective algorithms tested: Swap-Eager (swap, magenta, circle) and non-projective Covington (covington-np, green, triangle). The corresponding p-values (derived from a t-test respectively using the average precision and recall and the corresponding standard deviation across treebanks for each displacement value for each algorithm) are shown below: Swap-Eager and non-projective Covington (swap-covnp). Almost no statistically significant differences are observed for any comparison with regard to attachment precision or recall, except for left arcs (positive displacement) with regard to precision.

tribution of the random variable generated by the following process:

- Start at the initial configuration $c_s(k)$.
- At each configuration, let $t_1 \dots t_q$ be the available transitions. Choose one of them randomly with probability $1/q$, and go to the resulting next configuration.
- The process ends when a terminal configuration $c_t = (D_t, A_t) \in C_t$ has been reached. Then, choose a dependency arc uniformly at random from A_t and take its displacement as the value of the random variable.

Note that the inherent distribution of an algorithm does not depend on the contents of the particular sentence being parsed in the stochastic process, but just its length. Therefore, it can be seen as a variable that describes the distribution of displacements that the algorithm is “biased” to produce, in the absence of any training data. The transition is selected using a uniform probability across all possible transitions for a given configuration as there is no underlying reason why an algorithm would select one transition over another *without* using the feature function at a given timestep. Our hypothesis is that a given language or corpus will be parsed more accurately by algorithms whose inher-




Figure 3: Example comparisons between the dependent displacement distribution of the German-GSD treebank (de, magenta) for sentences of length 10 to 12 in version 2.1 of the Universal Dependency treebanks and the inherent displacement distributions of two algorithms: Arc-Standard (std, green) and Arc-Eager (eager, purple). The corresponding UAS and EMD values are displayed.

ent distribution is closer (as measured by the EMD) to the actual observed displacements in that language or corpus.

While the inherent distribution of an algorithm for sentences of length k would be difficult to obtain analytically, especially for the algorithms that support arbitrary non-projectivity where exact inference is intractable (McDonald and Satta, 2007), in practice we will approximate it by running a number of simulations of the above stochastic process.

The above definition can be extended to corpora (or subsets of them, such as the sentence-length bins we use in this paper). Let S be a set of n sentences containing n_k sentences of length k , for a range of values of k . Then, the inherent distribution of P with respect to S is the discrete probability distribution of the random variable generated by taking a random sentence length from S (where each length k is taken with probability n_k/n), and then taking a random displacement using the process above. The inherent distribution of P with respect to a set S , denoted $\iota_S(P)$, can be approximated by running a number of simulations of the above stochastic process on all the sentences of S .

Approximating inherent distributions

In order to approximate the biased distributions for each parsing algorithm P , we implemented a version of each of them so that they randomly select a transition from the set of available transitions for any given configuration. For each tree in a treebank that fell within the range of a sentence-length bin, a random tree was generated this way. This was done so as to ensure the EMD of different distributions was due to differences in the dependency displacement and to minimize other factors affecting the EMD. In other words, we wanted to obtain inherent distributions that echo the output of a parsing algorithm if it had been run normally (i.e. making predictions based on a feature function). For each treebank and each sentence length bin B in its test set, a random displacement distribution was generated 10 times to approximate $\iota_B(P)$, and their average EMD of

the observed distribution of displacements in the trees of B was taken. We opted to run it 10 times individually rather than run it once with more data points as this way we can more easily evaluate the uncertainty of a given inherent distribution. The standard error for the average EMD can be observed in Figures 5 and 6. It is clear that the variation across each generated distribution is quite small and that 10 instances are enough to minimise the uncertainty of this procedure.

We split the data according to sentence length to account for the differences in arc lengths that arise from longer or shorter sentences. Optimally, we would have undertaken our analyses according to sentence length and would not have used bins, however, the statistics were too low for many treebanks in order to this. The sentence lengths bins used and their statistics can be seen in Figure 4.




Figure 4: Corresponding stats for each bin used in the subsequent analysis with the average across the 26 treebanks shown and the first and third quartile limits.

Because of the varying difficulty of the datasets in the universal dependency collection, we have opted to compare the average EMD against δUAS , defined as the difference between the performance of an algorithm on a treebank minus the average score across algorithms for that treebank:

$$\delta\text{UAS}_{L_A} = \text{UAS}_{L_A} - \overline{\text{UAS}}_L \quad (2)$$

where A is the algorithm and L is the treebank.

We then compare the δUAS for an algorithm against its average EMD. We do this for 26 treebanks from version 2.1 of the Universal Dependencies treebanks. These languages were selected based on their size. We removed all languages with less than 1,000 trees in the training set and in the test set. This was necessary because if a treebank was too small then most sentence-length bins would not have enough stats to compute a meaningful EMD.

We split our analysis into projective and non-projective algorithms. The performance between projective and non-projective algorithms on certain datasets would be dominated by the percentage of non-projective arcs in the data and would potentially cloud any effect that the displacement distribution similarities might have. Nivre (2008) found a strong correlation between the percentage of non-projective dependencies and the improvement in accuracy for a non-projective parser ($r=0.815$, $p=7.0 \times 10^{-4}$) using a




Figure 5: δUAS (as defined in Equation 2) for each algorithm against the corresponding average EMD ($k=10$) for projective algorithms in the 10-12 token sentence-length bin.

much more limited dataset (CoNLL-X shared task 2006). Furthermore, the search space affects the random distributions (for example, non-projective trees in general have larger average dependency lengths than projective trees, see for example (Ferrer-i-Cancho and Gómez-Rodríguez, 2016)) and this could also be a confounding factor. So by separating projective and non-projective algorithms, we have made the search space a fixed factor.

Finally, we also compared algorithms directly. We did this by comparing ΔUAS and ΔEMD , defined as:

$$\Delta UAS_L = UAS_{L_{A_1}} - UAS_{L_{A_2}} \quad (3)$$

$$\Delta EMD_L = \overline{EMD}_{L_{A_1}} - \overline{EMD}_{L_{A_2}} \quad (4)$$

where A_1 is the first algorithm, A_2 is the second, and L is the treebank.

5.3. Results

There is a lack of meaningful correlation for both projective and non-projective when looking at the displacement distributions for unbinned treebanks. The correlation and corresponding p-value for the projective algorithms were $r = -0.045$ and $p = 6.97 \times 10^{-1}$. For the non-projective algorithms they were $r = -0.252$ and $p = 7.17 \times 10^{-2}$. Neither result is statistically significant nor does either show any strong correlation despite that. This corroborates the findings of Ferrer-i-Cancho and Liu (2014) and further justifies our binning procedure.

Figure 5 shows an example plot comparing δUAS against \overline{EMD} for the projective algorithms. For this bin (10-12 tokens) there is a strong negative correlation of -0.533 . Hence, r^2 is 0.284, meaning that the \overline{EMD} accounts for 28.4% of the variance seen in δUAS . The correlation is statistically significant ($p = 4.98 \times 10^{-7}$).

The corresponding results for all of the sentence-length bins can be observed in Table 1, where unsurprisingly the shortest sentences (lengths 1 to 3) show no correlation (they are too short for any meaningful difference between dependency displacement distributions) and the correlations start to diminish as the sentences get larger, but are still statistically meaningful until sentence lengths of 25.

Figure 6 shows the correlation for the non-projective algorithms for the same bin as Figure 5. It is clear that the corre-

token bin	r	r ²	p
1-3	-0.060	0.004	6.09×10^{-01}
4-6	-0.401	0.161	2.75×10^{-04}
7-9	-0.503	0.253	2.74×10^{-06}
10-12	-0.533	0.284	4.98×10^{-07}
13-15	-0.526	0.277	7.49×10^{-07}
16-18	-0.514	0.264	1.47×10^{-06}
19-21	-0.402	0.161	2.68×10^{-04}
22-24	-0.304	0.093	6.78×10^{-03}
25-27	-0.202	0.041	7.65×10^{-02}
28-33	-0.072	0.005	5.29×10^{-01}
34-39	-0.034	0.001	7.70×10^{-01}
40-99	0.139	0.019	2.43×10^{-01}

Table 1: Full results for each sentence-length bin for projective algorithms where token bin is the sentence length range, r is the Pearson coefficient of the correlation between δUAS and the EMD of each algorithm (e.g. as shown in Figure 5), r^2 is the squared Pearson coefficient which gives an indication of how much variation in the data this correlation accounts for, and finally p is the p-value for a given correlation.




Figure 6: δUAS (as defined in Equation 2) for each algorithm against the corresponding average EMD ($k=10$) for non-projective algorithms in the 10-12 token sentence-length bin.

lation is not as strong for the non-projective algorithms, but it is still large enough to be meaningful and is statistically significant. Table 2 shows the results for all of the bins used. Interestingly, the correlation does not diminish so severely for the non-projective algorithms as the sentence length increases as is the case for the projective algorithms.

Finally, focusing on direct comparisons between two algorithms, Figure 7 shows the comparison between Arc-Standard and Arc-Eager for the same bin as above. A strong negative correlation can be seen which is statistically significant. The results for the direct comparison between the three projective algorithms and the two non-projective algorithms can be seen in Figure 8. These comparisons yield significant results for Arc-Standard when compared with Arc-Eager and Covington for moderate lengthed sentences (4-25 tokens), but not for the non-projective algorithms and for Arc-Eager and Covington. This suggests the

token bin	r	r ²	p
1-3	0.001	0.000	9.92×10^{-01}
4-6	-0.243	0.059	8.32×10^{-02}
7-9	-0.327	0.107	1.79×10^{-02}
10-12	-0.386	0.149	4.70×10^{-03}
13-15	-0.344	0.118	1.25×10^{-02}
16-18	-0.364	0.133	7.90×10^{-03}
19-21	-0.344	0.118	1.26×10^{-02}
22-24	-0.350	0.122	1.11×10^{-02}
25-27	-0.349	0.122	1.11×10^{-02}
28-33	-0.347	0.121	1.17×10^{-02}
34-39	-0.354	0.125	1.01×10^{-02}
40-99	-0.298	0.089	3.97×10^{-02}

Table 2: Full results for each sentence-length bin for non-projective algorithms where token bin is the sentence length range, r is the Pearson coefficient of the correlation between δ UAS and the EMD of each algorithm (e.g. as shown in Figure 6), r^2 is the squared Pearson coefficient which gives an indication of how much variation in the data this correlation accounts for, and finally p is the p-value for a given correlation.




Figure 7: Δ UAS as defined in Equation 3 against Δ EMD as defined in Equation 4 comparing Arc-Eager and Arc-Standard (10-12 token sentence-length bin).

previous analysis of comparing all projective and the two non-projective algorithms is a statistically more powerful means of analysing the effect of the similarity of dependency displacement distribution on algorithm performance.

6. Discussion

A coarse analysis focusing on dependency displacements individually does not show a significant difference in performance across algorithms except with regards to the attachment precision for non-projective algorithms. However, the difference between an algorithm’s biased latent dependency displacement distribution and the target treebank being parsed is correlated with the performance of the algorithm for that treebank.

The obtained correlations are statistically significant for the sentence lengths that comprise most of the sentences found in actual corpora, both when analyzing projective and non-projective algorithms. In the case of projective algorithms,




Figure 8: Absolute Pearson coefficients and the corresponding p-values from comparisons between each pair of projective algorithms and the two non-projective algorithms: Arc-Eager and Arc-Standard (eager-std, blue); Arc-Eager and Covington (eager-cov, yellow); Arc-Standard and Covington (std-cov, orange); and non-projective Covington and Swap-Eager (covnp-swap, magenta). Statistically relevant comparisons can be seen between Arc-Standard and the other projective algorithms for mid-range tree lengths.

this factor accounts for more than 25% of the variance in UAS across algorithms. This is a remarkable proportion given the complexity in explaining how treebank-specific accuracy differs between algorithms and the variety of factors involved.

In fact, to the best of our knowledge, this is the first study in which these differences are studied quantitatively between transition-based algorithms of the same search space, thus casting light on a question that has been open since the introduction of the first transition-based parsers in the early 2000s.

It is worth noting that the effect of displacement distributions on parsing accuracy is independent of the effect of transition sequence length. It has been hypothesized that short transition sequences reduce error propagation (Fernández-González and Gómez-Rodríguez, 2018), but this effect does not help when comparing the relative performance of algorithms on different treebanks, as we do here: for example, the arc-standard and arc-eager algorithms produce transition sequences of identical lengths, independently on the syntactic structures found (they always need exactly $2n$ transitions for a sentence of length n). However, as we have seen, their inherent displacement distributions are different and can be used to explain their suitability to different treebanks.

The insights provided in this paper could be useful to guide parsing algorithm design: since algorithms tend to be more

accurate on corpora that are closer to their inherent distribution, a potential avenue for designing better transition-based parsing algorithms is to try to make their inherent distribution match that of human languages more closely. To further validate the hypothesis investigated here, it would be interesting to generate artificial treebanks in such a way so as to create a spread of arc distributions so we can control the EMD range.

Beyond the explicit findings of this paper, it is interesting to observe that linguistic considerations can have an impact in natural language processing systems and more analyses like this, such as considering what makes certain languages harder to model than others, should hopefully prove to be useful in the future (Mielke et al., 2019).

7. Conclusion

We have introduced the concept of an algorithm’s inherent displacement distribution, which captures the algorithm’s bias towards implicitly preferring certain dependency lengths and directions to others. We have shown that given a treebank, the similarity between each transition-based algorithm’s inherent dependency displacement distribution and the treebank’s distribution is a strongly correlated to the corresponding algorithm’s performance on that treebank.

Acknowledgments

This work has received funding from the European Research Council (ERC), under the European Union’s Horizon 2020 research and innovation programme (FAST-PARSE, grant agreement No 714150), from the ANSWER-ASAP project (TIN2017-85160-C2-1-R) from MINECO, and from Xunta de Galicia (ED431B 2017/01, ED431G 2019/01).

8. Bibliographical References

- Ballesteros, M. and Nivre, J. (2012). MaltOptimizer: an optimization tool for MaltParser. In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 58–62. Association for Computational Linguistics.
- Ballesteros, M. and Nivre, J. (2013). Going to the roots of dependency parsing. *Computational Linguistics*, 39(1):5–13.
- Buch-Kromann, M. (2006). *Discontinuous grammar: A dependency-based model of human parsing and language learning*. Ph.D. thesis, Copenhagen Business School.
- Chen, D. and Manning, C. (2014). A fast and accurate dependency parser using neural networks. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 740–750.
- de Lhoneux, M., Szymne, S., and Nivre, J. (2017). Old school vs. new school: Comparing transition-based parsers with and without neural network enhancement. In *TLT*, pages 99–110.
- Dickinson, M. (2010). Detecting errors in automatically-parsed dependency relations. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 729–738.
- Eisner, J. and Smith, N. A. (2010). Favor short dependencies: Parsing with soft and hard constraints on dependency length. In *Trends in Parsing Technology*, pages 121–150. Springer.
- Fernández-González, D. and Gómez-Rodríguez, C. (2018). Non-projective dependency parsing with non-local transitions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 693–700. Association for Computational Linguistics.
- Fernández-González, D., Gómez-Rodríguez, C., and Vilares, D. (2016). Improving the arc-eager model with reverse parsing. *Computing and Informatics*, 35(3):555–585.
- Ferrer-i-Cancho, R. and Gómez-Rodríguez, C. (2016). Crossings as a side effect of dependency lengths. *Complexity*, 21(S2):320–328.
- Ferrer-i-Cancho, R. and Liu, H. (2014). The risks of mixing dependency lengths from sequences of different length. *Glottometry*, 5(2):143–155.
- Ferrer-i-Cancho, R. (2004). Euclidean distance between syntactically linked words. *Physical Review E*, 70(5):056135.
- Futrell, R., Mahowald, K., and Gibson, E. (2015). Large-scale evidence of dependency length minimization in 37 languages. *Proceedings of the National Academy of Sciences*, 112(33):10336–10341.
- Gibson, E. (2000). The dependency locality theory: A distance-based theory of linguistic complexity. *Image, language, brain*, 2000:95–126.
- Gildea, D. and Temperley, D. (2010). Do grammars minimize dependency length? *Cognitive Science*, 34(2):286–310.
- Goldberg, Y. and Elhadad, M. (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*, pages 742–750. Association for Computational Linguistics.
- Gómez-Rodríguez, C. and Ferrer-i-Cancho, R. (2017). Scarcity of crossing dependencies: A direct outcome of a specific constraint? *Physical Review E*, 96(6):062304.
- Gómez-Rodríguez, C. (2017). On the relation between dependency distance, crossing dependencies, and parsing. *Physics of Life Reviews*, 21:200 – 203.
- Gulordava, K. and Merlo, P. (2015). Diachronic trends in word order freedom and dependency length in dependency-annotated corpora of Latin and Ancient Greek. In *Proceedings of the third international conference on dependency linguistics (Depling 2015)*, pages 121–130.
- Gulordava, K. and Merlo, P. (2016). Multi-lingual dependency parsing evaluation: a large-scale analysis of word order properties using artificial data. *Transactions of the Association for Computational Linguistics*, 4:343–356.
- Hudson, R. (2017). Cross-language diversity, head-

- direction and grammars. comment on “Dependency distance: A new perspective on syntactic patterns in natural languages” by Haitao Liu et al. *Physics of life reviews*, 21:204–206.
- Jiang, J. and Liu, H. (2015). The effects of sentence length on dependency distance, dependency direction and the implications—based on a parallel english–chinese dependency treebank. *Language Sciences*, 50:93–104.
- Kirilin, A. and Versley, Y. (2015). What is hard in universal Dependency parsing. In *Proceedings of the 6th Workshop on Statistical Parsing of Morphologically Rich Languages (SPMRL 2015)*, pages 31–38.
- Kulmizev, A., de Lhoneux, M., Gontrum, J., Fano, E., and Nivre, J. (2019). Deep contextualized word embeddings in transition-based and graph-based dependency parsing - a tale of two parsers revisited. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2755–2768, Hong Kong, China, November. Association for Computational Linguistics.
- Liu, H., Xu, C., and Liang, J. (2017). Dependency distance: a new perspective on syntactic patterns in natural languages. *Physics of life reviews*, 21:171–193.
- Liu, H. (2007). Probability distribution of dependency distance. *Glottometrics*, 15:1–12.
- Liu, H. (2008). Dependency distance as a metric of language comprehension difficulty. *Journal of Cognitive Science*, 9(2):159–191.
- Liu, H. (2010). Dependency direction as a means of word-order typology: A method based on dependency treebanks. *Lingua*, 120(6):1567–1578.
- McDonald, R. and Nivre, J. (2011). Analyzing and integrating dependency parsers. *Computational Linguistics*, 37(1):197–230.
- McDonald, R. and Satta, G. (2007). On the complexity of non-projective data-driven dependency parsing. In *Proceedings of the 10th International Conference on Parsing Technologies*, pages 121–132. Association for Computational Linguistics.
- Mielke, S. J., Cotterell, R., Gorman, K., Roark, B., and Eisner, J. (2019). What kind of language is hard to language-model? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4975–4989, Florence, Italy, July. Association for Computational Linguistics.
- Nivre, J., Hall, J., Nilsson, J., Chanev, A., Eryigit, G., Kübler, S., Marinov, S., and Marsi, E. (2007). Malt-Parser: A language-independent system for data-driven dependency parsing. *Natural Language Engineering*, 13(2):95–135.
- Nivre, J., Agić, Ž., Ahrenberg, L., et al. (2017). Universal Dependencies 2.1. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Nivre, J. (2008). Algorithms for deterministic incremental dependency parsing. *Computational Linguistics*, 34(4):513–553.
- Rehbein, I., Steen, J., Do, B.-N., and Frank, A. (2017). Universal Dependencies are hard to parse—or are they? In *Proceedings of the Fourth International Conference on Dependency Linguistics (Depling 2017)*, pages 218–228.
- Rubner, Y., Tomasi, C., and Guibas, L. J. (1998). A metric for distributions with applications to image databases. In *Proceedings of the Sixth International Conference on Computer Vision, ICCV ’98*, pages 59–, Washington, DC, USA. IEEE Computer Society.
- Schwartz, R., Abend, O., and Rappoport, A. (2012). Learnability-based syntactic annotation design. *Proceedings of COLING 2012*, pages 2405–2422.
- Straka, M. and Straková, J. (2017). Tokenizing, POS tagging, Lemmatizing and Parsing UD 2.0 with UDPipe. In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 88–99.
- Temperley, D. and Gildea, D. (2018). Minimizing syntactic dependency lengths: typological/cognitive universal? *Annual Review of Linguistics*, 4:67–80.