

Synthetic Data Made to Order: The Case of Parsing

Dingquan Wang and Jason Eisner

Department of Computer Science, Johns Hopkins University

{wdd, jason}@cs.jhu.edu

Abstract

To approximately parse an unfamiliar language, it helps to have a treebank of a similar language. But what if the closest available treebank still has the wrong word order? We show how to (stochastically) permute the constituents of an existing dependency treebank so that its surface part-of-speech statistics approximately match those of the target language. The parameters of the permutation model can be evaluated for quality by dynamic programming and tuned by gradient descent (up to a local optimum). This optimization procedure yields trees for a new artificial language that resembles the target language. We show that delexicalized parsers for the target language can be successfully trained using such “made to order” artificial languages.

1 Introduction

Dependency parsing is a core task in natural language processing (NLP). Given a sentence, a dependency parser produces a dependency tree, which specifies the typed head-modifier relations between pairs of words. While *supervised* dependency parsing has been successful (McDonald and Pereira, 2006; Nivre, 2008; Kiperwasser and Goldberg, 2016), unsupervised parsing can hardly produce useful parses (Mareček, 2016). So it is extremely helpful to have some treebank of supervised parses for training purposes.

1.1 Past work: Cross-lingual transfer

Unfortunately, manually constructing a treebank for a new target language is expensive (Böhmová et al., 2003). As an alternative, *cross-lingual transfer* parsing (McDonald et al., 2011) is sometimes possible, thanks to the recent development of multi-lingual treebanks (McDonald et al., 2013; Nivre et al., 2015; Nivre et al., 2017). The idea is to parse the sentences of the target language with a supervised parser trained on the treebanks of one or more source languages. Although the parser cannot be expected to know the words of the target language, it can make do with parts of

speech (POS) (McDonald et al., 2011; Täckström et al., 2013; Zhang and Barzilay, 2015) or cross-lingual word embeddings (Duong et al., 2015; Guo et al., 2016; Ammar et al., 2016). A more serious challenge is that the parser may not know how to handle the word *order* of the target language, unless the source treebank comes from a closely related language (e.g., using German to parse Luxembourgish). Training the parser on trees from *multiple* source languages may mitigate this issue (McDonald et al., 2011) because the parser is more likely to have seen target part-of-speech sequences somewhere in the training data. Some authors (Rosa and Žabokrtský, 2015a,b; Wang and Eisner, 2016) have shown additional improvements by preferring source languages that are “close” to the target language, where the closeness is measured by distance between POS language models trained on the source and target corpora.

1.2 This paper: Tailored synthetic data

We will focus on delexicalized dependency parsing, which maps an input POS tag sequence to a dependency tree. We evaluate single-source transfer—train a parser on a single source language, and evaluate it on the target language. This is the setup of Zeman and Resnik (2008) and Sjøgaard (2011a).

Our novel ingredient is that *rather than seek a close source language that already exists, we create one*. How? Given a dependency treebank of a possibly distant source language, we stochastically permute the children of each node, according to some distribution that *makes the permuted language close to the target language*.

And how do we find this distribution? We adopt the tree-permutation model of Wang and Eisner (2016). We design a dynamic programming algorithm which, for any given distribution p in Wang and Eisner’s family, can compute the *expected counts* of all POS bigrams in the permuted source treebank. This allows us to evaluate p by computing the divergence between the bigram POS language model formed by these expected counts,

and the one formed by the *observed counts* of POS bigrams in the unparsed target language. In order to find a p that locally minimizes this divergence, we adjust the model parameters by stochastic gradient descent (SGD).

1.3 Key limitations in this paper

Better measures of surface closeness between two languages might be devised. However, even counting the expected POS N -grams is moderately expensive, taking time exponential in N if done exactly. So we compute only these local statistics, and only for $N = 2$. We certainly need $N > 1$ because the 1-gram distribution is not affected by permutation at all. $N = 2$ captures useful bigram statistics: for example, to mimic a verb-final language with prenominal modifiers, we would seek constituent permutations that result in matching its relatively high rate of VERB–PUNCT and ADJ–NOUN bigrams. While $N > 2$ might have improved the results, it was too slow for our large-scale experimental design. §7 discusses how richer measures could be used in the future.

We caution that throughout this paper, we assume that our corpora are annotated with gold POS tags, even in the target language (which lacks any gold training trees). This is an idealized setting that has often been adopted in work on unsupervised and cross-lingual transfer. §7 discusses a possible avenue for doing without gold tags.

2 Modeling Surface Realization

We begin by motivating the idea of tree permutation. Let us suppose that the dependency tree for a sentence starts as a labeled graph—a tree in which siblings are not yet ordered with respect to their parent or one another. Each language has some systematic way to *realize* its unordered trees as surface strings:¹ it imposes a particular order on the tree’s word tokens. More precisely, a language specifies a *distribution* $p(\text{string} \mid \text{unordered tree})$ over a tree’s possible realizations.

As an engineering matter, we now make the strong assumption that the unordered dependency trees are similar across languages. That is, we suppose that different languages use similar underlying syntactic/semantic graphs, but differ in how they realize this graph structure on the surface.

¹Modeling this process was the topic of the recent Surface Realization Shared Task (Mille et al., 2018). Most relevant is work on *tree linearization* (Filippova and Strube, 2009; Futrell and Gibson, 2015; Puzikov and Gurevych, 2018).

Thus, given a gold POS corpus \mathbf{u} of the unknown target language, we may hope to explain its distribution of surface POS bigrams as the result of applying some target-language surface realization model to the distribution of cross-linguistically “typical” unordered trees. To obtain samples of the latter distribution, we use the treebanks of one or more *other* languages. The present paper evaluates our method when only a single source treebank is used. In the future, we could try tuning a mixture of all available source treebanks.

2.1 Realization is systematic

We presume that the target language applies the same stochastic realization model to all trees. All that we can optimize is the parameter vector of this model. Thus, we deny ourselves the freedom to realize each individual tree in an *ad hoc* way. To see why this is important, suppose the target language is French, whose corpus \mathbf{u} contains many NOUN–ADJ bigrams. We could achieve such a bigram from the unordered source tree

the cake made Sue sleepy by ordering

it to yield the cake sleepy made Sue .

However, that realization is not in fact appropriate for French, so that ordered tree would not be a useful training tree for French. Our approach should disprefer this tempting but incorrect realization, because any model with a high probability of this realization would, if applied *systematically* over the whole corpus, also yield sentences like He sleepy made Sue, with unwanted PRON–ADJ bigrams that would not match the surface statistics of French. We hope our approach will instead choose the realization model that is correct for French, in which the NOUN–ADJ bigrams arise instead from source trees where the ADJ is a dependent of the NOUN, yielding (e.g.)

the cake tasty pleased Sue . This has the same POS sequence as the example above (as it happens), but now assigns the correct tree to it.

2.2 A parametric realization model

As our family of realization distributions, we adopt the log-linear model used for this purpose by Wang and Eisner (2016). The model assumes that the root node a of the unordered dependency tree selects an ordering $\pi(a)$ of the n_a nodes consisting

of a and its $n_a - 1$ dependent children. The procedure is repeated recursively at the child nodes. This method can produce only projective trees.

Each node a draws its ordering $\pi(a)$ independently according to

$$p_{\theta}(\pi | a) = \frac{1}{Z(a)} \exp \sum_{1 \leq i < j \leq n_a} \theta \cdot \mathbf{f}(\pi, i, j) \quad (1)$$

which is a distribution over the $n_a!$ possible orderings. $Z(a)$ is a normalizing constant. \mathbf{f} is a feature vector extracted from the ordered pair of nodes π_i, π_j , and θ is the model’s parameter vector of feature weights. See Appendix A for the feature templates, which are a subset of those used by Wang and Eisner (2016). These features are able to examine the tree’s node labels (POS tags) and edge labels (dependency relations). Thus, when a is a verb, the model can assign a positive weight to “subject precedes verb” or “subject precedes object,” thus preferring orderings with these features.

Following Wang and Eisner (2016, §3.1), we choose new orderings for the noun and verb nodes only,² preserving the source treebank’s order at all other nodes a .

2.3 Generating training data

Given a source treebank B and some parameters θ , we can use equation (1) to randomly sample realizations of the trees in B . The effect is to reorder dependent phrases within those trees. The resulting permuted treebank B' can be used to train a parser for the target language.

2.4 Choosing parameters θ

So how do we choose θ that works for the target language? Suppose \mathbf{u} is a corpus of target-language POS sequences, using the same set of POS tags as B . We evaluate parameters θ according to whether POS tag sequences in B' will be distributed like POS tag sequences in \mathbf{u} .

To do this, first we estimate a bigram language model \hat{q} from the actual distribution q of POS sequences observed in \mathbf{u} . Second, let p_{θ} denote the distribution of POS sequences that we expect to see in B' , that is, POS sequences obtained by

²Specifically, the 93% of nodes tagged with NOUN, PROP, PRON or VERB in Universal Dependencies format. In retrospect, this restriction was unnecessary in our setting, but it skipped only 4.4% of nodes on average (from 2% to 11% depending on language). The remaining nodes were nouns, verbs, or childless.

stochastically realizing observed trees in B according to θ . We estimate another bigram model \hat{p}_{θ} from this distribution p_{θ} .

We then try to set θ , using SGD, to minimize a divergence $D(\hat{p}_{\theta}, \hat{q})$ that we will define below.

2.4.1 Estimation of bigram models

Estimating \hat{q} is straightforward: $\hat{q}(t | s) = c_q(st)/c_q(s)$, where $c_q(st)$ is the count of POS bigram st in the average³ sentence of \mathbf{u} and $c_q(s) = \sum_{t'} c_q(st')$. We estimate \hat{p}_{θ} in the same way, where $c_p(st)$ denotes the *expected* count of st in a random POS sequence $\mathbf{y} \sim p_{\theta}$. This is equivalent to choosing $\hat{q}, \hat{p}_{\theta}$ to minimize the KL-divergences $\text{KL}(q || \hat{q}), \text{KL}(p_{\theta} || \hat{p}_{\theta})$. It ensures that each model’s expected bigram counts match those in the POS sequences.

However, these maximum-likelihood estimates might overfit on our finite data, \mathbf{u} and B . We therefore smooth both models by first adding $\lambda = 0.1$ to all bigram counts $c_q(st)$ and $c_p(st)$.⁴

2.4.2 Divergence of bigram models

We need a metric to evaluate θ . If p and q are bigram language models over POS sequences \mathbf{y} (sentences), their Kullback-Leibler divergence is

$$\text{KL}(p || q) \stackrel{\text{def}}{=} \mathbb{E}_{\mathbf{y} \sim p} [\log p(\mathbf{y}) - \log q(\mathbf{y})] \quad (2)$$

$$= \sum_{s,t} c_p(st) \cdot (\log p(t | s) - \log q(t | s)) \quad (3)$$

where \mathbf{y} ranges over POS sequences and st ranges over POS bigrams. These include bigrams where $s = \text{BOS}$ (“beginning of sequence”) or $t = \text{EOS}$ (“end of sequence”), which are boundary tags that we take to surround \mathbf{y} .

All quantities in equation (3) can be determined directly from the (expected) bigram counts given by c_p and c_q . No other model estimation is needed.

A concern about equation (3) is that a single bigram st that is badly underrepresented in q may contribute an arbitrarily large term $\log \frac{p(t|s)}{q(t|s)}$. To limit this contribution to at most $\log \frac{1}{\alpha}$, for some small $\alpha \in (0, 1)$, we define $\text{KL}_{\alpha}(p || q)$ by a variant of equation (3) in which $q(t | s)$ has been replaced by $\tilde{q}(t | s) \stackrel{\text{def}}{=} \alpha p(t | s) + (1 - \alpha)q(t | s)$.⁵

³A more familiar definition of c_q would use the *total* count in \mathbf{u} . Our definition, which yields the same bigram probabilities, is analogous to our definition of c_p . This c_p is needed for $\text{KL}(p || q)$ in (3), and c_q symmetrically for $\text{KL}(q || p)$.

⁴Ideally one should tune λ to minimize the language model perplexity on held-out data (e.g., by cross-validation).

⁵This is inspired by the α -skew divergence of Lee (1999).

Our final divergence metric $D(\hat{p}_\theta, \hat{q})$ defines D as a linear combination of exclusive and inclusive KL_α divergences, which respectively emphasize p_θ 's precision and recall at matching q 's bigrams:

$$D(p, q) = (1-\beta) \cdot \frac{\text{KL}_{\alpha_1}(p \parallel q)}{\mathbb{E}_{\mathbf{y} \sim p}[|\mathbf{y}|]} + \beta \cdot \frac{\text{KL}_{\alpha_2}(q \parallel p)}{\mathbb{E}_{\mathbf{y} \sim q}[|\mathbf{y}|]} \quad (4)$$

where $\beta, \alpha_1, \alpha_2$ are tuned by cross-validation to maximize the downstream parsing performance. The division by average sentence length converts KL from nats per sentence to nats per word,⁶ so that the KL values have comparable scale even if B has much longer or shorter sentences than \mathbf{u} .

3 Algorithms

3.1 Efficiently computing expected counts

We now present a polynomial-time algorithm for computing the expected bigram counts c_p under p_θ (or equivalently \hat{p}_θ), for use above. This averages expected counts from each unordered tree $\mathbf{x} \in B$. Algorithm 1 in the supplement gives pseudocode.

The insight is that rather than sampling a single realization of \mathbf{x} (as B' does), we can use dynamic programming to sum efficiently over all of its exponentially many realizations. This gives an exact answer. It algorithmically resembles tree-to-string machine translation, which likewise considers the possible reorderings of a source tree and incorporates a language model by similarly tracking their surface N -grams (Chiang, 2007, §5.3.2).

For each node a of the tree \mathbf{x} , let the POS string \mathbf{y}_a be the realization of the subtree rooted at a . Let $c_a(st)$ be the expected count of bigram st in \mathbf{y}_a , whose distribution is governed by equation (1). We allow $s = \text{BOS}$ or $t = \text{EOS}$ as defined in §2.4.2.

The c_a function can be represented as a sparse map from POS bigrams to reals. We compute c_a at each node a of \mathbf{x} in a bottom-up order. The final step computes c_{root} , giving the expected bigram counts in \mathbf{x} 's realization \mathbf{y} (that is, c_p in §2.4).

We find c_a as follows. Let $n = n_a$ and recall from §2.2 that $\pi(a)$ is an ordering of a_1, \dots, a_n , where a_1, \dots, a_{n-1} are the child nodes of a , and a_n is a dummy node representing a 's head token.

⁶2001). Indeed, we may regard $\text{KL}_\alpha(p \parallel q)$ as the α -skew divergence between the unigram distributions $p(\cdot \mid s)$ and $q(\cdot \mid s)$, averaged over all s in proportion to $c_p(s)$. In principle, we could have used the α -skew divergence between the distributions $p(\cdot)$ and $q(\cdot)$ over POS sequences \mathbf{y} , but computing that would have required a sampling-based approximation (§7).

⁶Recall that the units of negated log-probability are called bits for log base 2, but nats for log base e .

Also, let a_0 and a_{n+1} be dummy nodes that always appear at the start and end of any ordering.

For all $0 \leq i \leq n$ and $1 \leq j \leq n+1$, let $p_a(i, j)$ denote the expected count of the $a_i a_j$ node bigram—the probability that $\pi(a)$ places node a_i immediately before node a_j . These *node bigram probabilities* can be obtained by enumerating all possible orderings π , a matter we return to below.

It is now easy to compute c_a :

$$\begin{aligned} c_a(st) &= c_a^{\text{within}}(st) + c_a^{\text{between}}(st) \quad (5) \\ c_a^{\text{within}}(st) &= \begin{cases} \sum_{i=1}^n c_{a_i}(st) & \text{if } s \neq \text{BOS}, t \neq \text{EOS} \\ 0 & \text{otherwise} \end{cases} \\ c_a^{\text{across}}(st) &= \sum_{i=0}^n \sum_{j=1}^{n+1} p_a(i, j) c_{a_i}(s \text{ EOS}) c_{a_j}(\text{BOS } t) \end{aligned}$$

That is, c_a inherits all non-boundary bigrams st that fall within its child constituents (via c_a^{within}). It also counts bigrams st that cross the boundary between consecutive nodes (via c_a^{across}), where nodes a_i and a_j are consecutive with probability $p_a(i, j)$.

When computing c_a via (5), we will have already computed $c_{a_1}, \dots, c_{a_{n-1}}$ bottom-up. As for the dummy nodes, a_n is realized by the length-1 string h where h is the head token of node a , while a_0 and a_{n+1} are each realized by the empty string. Thus, c_{a_n} simply assigns count 1 to the bigrams $\text{BOS } h$ and $h \text{ EOS}$, and c_{a_0} and $c_{a_{n+1}}$ each assign expected count 1 to BOS EOS . (Notice that thus, $c_a^{\text{across}}(st)$ counts \mathbf{y}_a 's boundary bigrams—the bigrams st where $s = \text{BOS}$ or $t = \text{EOS}$ —when $i = 0$ or $j = n+1$ respectively.)

3.2 Efficient enumeration over permutations

The main challenge above is computing the node bigram probabilities $p_a(i, j)$. These are marginals of $p(\pi \mid a)$ as defined by (1), which unfortunately is intractable to marginalize: there is no better way than enumerating all $n!$ permutations.

That said, there is a particularly efficient way to enumerate the permutations. The Steinhaus-Johnson-Trotter (SJT) algorithm (Sedgewick, 1977) does so in $O(1)$ time per permutation, obtaining each permutation by applying a single swap to the previous one. Only the features that are affected by this swap need to be recomputed. For our features (Appendix A), this cuts the runtime per permutation from $O(n^2)$ to $O(n)$.

Furthermore, the single swap of adjacent nodes only changes 3 bigrams (possibly including boundary bigrams). As a result, it is possible to

obtain the marginal probabilities with $O(1)$ additional work per permutation. When a node bigram is destroyed, we increment its marginal probability by the total probability of permutations encountered since the node bigram was last created. This can be found as a difference of partial sums. The final partial sum is the normalizing constant $Z(a)$, which can be applied at the end. Pseudocode is given in supplementary material as Algorithm 2.

When we train the parameters θ (§2.4), we must back-propagate through the whole computation of equation (4), which depends on tag bigram counts $c_a(st)$, which depend via (5) on expected node bigram counts $p_a(i, j)$, which depend via Algorithm 2 on the permutation probabilities $p(\pi | a)$, which depend via (1) on the feature weights θ .

4 Heuristics

4.1 Pruning high-degree trees

As a further speedup, we only train on trees with number of words < 40 and $\max_a n_a \leq 5$, so $n_a! \leq 120$.⁷ We then produce the synthetic treebank B' (§2.3) by drawing a single realization of each tree in B for which $\max_a n_a \leq 7$. This requires sampling from up to $7! = 5040$ candidates per node, again using SJT.⁸

That is, in this paper we run exact algorithms (§3), but only on a subset of B . The subset is not necessarily representative. An improvement would use importance sampling, with a proposal distribution that samples the slower trees less often during SGD but upweights them to compensate.

§7 suggests a future strategy that would run on all trees in B via approximate, sampling-based algorithms. The exact methods would remain useful for calibrating the approximation quality.

4.2 Minibatch estimation of c_p

To minimize (4), we use the Adam variant of SGD (Kingma and Ba, 2014), with learning rate 0.01 chosen by cross-validation (§5.1).

SGD requires a stochastic estimate of the gradient of the training objective. Ordinarily this is done by replacing an expectation over the entire training set with an expectation over a minibatch.

⁷We found that this threshold worked much better than ≤ 4 and about as well as the much slower ≤ 6 .

⁸This pruning heuristic retains 36.1% of the trees (averaging over the 20 development treebanks (§5.1)) for training, and 66.6% for actual realization. The latter restriction follows Wang and Eisner (2016, §4.2): they too discarded trees with nodes having $n_a \geq 8$.

Equation (2) with $p = \hat{p}_\theta$ is indeed an expectation over sentences of B . It can be stochastically estimated as (3) where c_p gives the expected bigram counts averaged over only the sentences in a minibatch of B . These are found using §3’s algorithms with the current θ . Unfortunately, the term $\log p(t | s)$ depends on bigram counts that should be derived from the *entire* corpus B in the same way. Our solution is to simply reuse the *minibatch* estimate of c_p for the latter counts. We use a large minibatch of 500 sentences from B so that this *drop-in estimate* does not introduce too much bias into the stochastic gradient: after all, we only need to estimate bigram statistics on 17 POS types.⁹

By contrast, the c_q values that are used for the expectation in the second term of (4) and in $\log q(t | s)$ do not change during optimization, so we simply compute them once from all of \mathbf{u} .

4.3 Informed initialization

Unfortunately the objective (4) is not convex, so the optimizer is sensitive to initialization (see §5.3 below for empirical discussion). Initializing $\theta = \mathbf{0}$ (so that $p(\pi | a)$ is uniform) gave poor results in pilot experiments. Instead, we initially choose θ to be the realization parameters of the source language, as estimated from the source treebank B . This is at least a linguistically realistic θ , although it may not be close to the target language.¹⁰

For this initial estimation, we follow Wang and Eisner (2016) and perform supervised training on B of the log-linear realization model (1), by maximizing the conditional log-likelihood of B , namely $\sum_{(\mathbf{x}, \mathbf{t}) \in B} \log p_\theta(\mathbf{t} | \mathbf{x})$, where (\mathbf{x}, \mathbf{t}) are an unordered tree and its observed ordering in B . This initial objective is convex.¹¹

5 Experiments

We performed a large-scale experiment requiring hundreds of thousands of CPU-hours. To our knowledge, this is the largest study of parsing transfer yet attempted.

⁹We also used the minibatch to estimate the average sentence length $\mathbb{E}_{\mathbf{y} \sim p}[|\mathbf{y}|]$ in (4), although here we could have simply used all of B since this value does not change.

¹⁰As an improvement, one could also try initial realization parameters for B that are estimated from treebanks of *other* languages. Concretely, the optimizer could start by selecting a “galactic” treebank from Wang and Eisner (2016) that is already close to the target language, according to (4), and try to make it even closer. We leave this to future work.

¹¹Unfortunately, we did not regularize it, which probably resulted in initializing some parameters too close to $\pm\infty$ for the optimizer to change them meaningfully.

5.1 Data and setup

As our main dataset, we use Universal Dependencies version 1.2 (Nivre et al., 2015)—a set of 37 dependency treebanks for 33 languages, with a unified POS-tag set and relation label set.

Our evaluation metric was unnormalized attachment score (UAS) when parsing a target treebank with a parser trained on a (possibly permuted) source treebank. For both evaluation and training, we used only the training portion of each treebank.

Our parser was Yara (Rasooli and Tetreault, 2015), a fast and accurate transition-based dependency parser that can be rapidly retrained. We modified Yara to ignore the input words and use only the input *gold* POS tags (see §1.3). To train the Yara parser on a (possibly permuted) source treebank, we first train on 80% of the trees and use the remaining 20% to tune Yara’s hyperparameters. We then retrain Yara on 100% of the source trees and evaluate it on the target treebank.

Similar to Wang and Eisner (2017), we use 20 treebanks (18 distinct languages) as *development data*, and hold out the remaining 17 treebanks for the final evaluation. We chose the hyperparameters $(\alpha_1, \alpha_2, \beta)$ of (4) to maximize the target-language UAS, averaged over all 376 transfer experiments where the source and target treebanks were development treebanks of different languages.¹² (See Appendix C for details.)

The next few sections perform some exploratory analysis on these 376 experiments. Then, for the final test in §5.4, we will evaluate UAS on all 337 transfer experiments where the source is a development treebank and the target is a test treebank of a different language.¹³

5.2 Exploratory analysis

We have assumed that a smaller divergence between source and target treebanks results in better transfer parsing accuracy. Figure 1 shows that these quantities are indeed correlated, both for the original source treebanks and for their “made to order” permuted versions.

¹²We have $19 \times 20 = 380$ pairs in total, minus the four excluded pairs (grc, grc_proiel), (grc_proiel, grc), (la_proiel, la_itt) and (la_itt, la_proiel). Unlike Wang and Eisner (2017), we exclude duplicated languages in development and testing.

¹³Specifically, there are 3 duplicated sets: {grc, grc_proiel}, {la, la_proiel, la_itt}, and {fi, fi_ftb}. Whenever one treebank is used as the target language, we exclude the other treebanks in the same set.

¹⁵According to the family (and sub-family) information at <http://universaldependencies.org>.

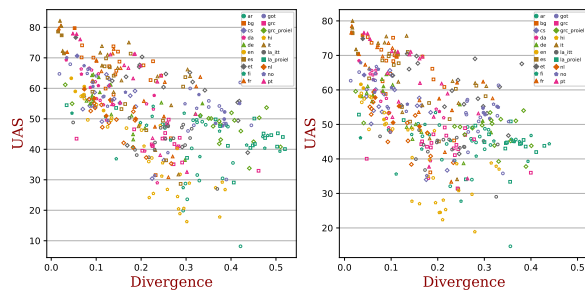


Figure 1: UAS is higher when divergence is lower. Each point represents a pair of source and target languages, whose shape and color identify the treebank of the target language (see legend). The marker is solid if the source and target languages belong to the same language family.¹⁵ The left graph uses the original source treebank (Kendall’s $\tau = -0.41$), while the right graph uses its permuted version ($\tau = -0.39$).

Thus, we hope that the optimizer will find a systematic permutation that reduces the divergence. Does it? Yes: Figures 5 and 6 in the supplementary material show that the optimizer almost always manages to reduce the objective on training data, as expected.

One concern is that our divergence metric might misguide us into producing dysfunctional languages whose trees cannot be easily recovered from their surface strings, i.e., they have no good parser. In such a language, the word order might be extremely free (e.g., $\theta = \mathbf{0}$), or common constructions might be syntactically ambiguous. Fortunately, Appendix D shows that our synthetic languages appear natural with respect to their parsability.

The above findings are promising. So does permuting the source language in fact result in better transfer parsing of the target language? We experiment on the 376 development pairs.

The solid lines in Figure 2 show our improvements on the dev data, with a simpler scatterplot given by in Figure 7 in the supplementary material. The upshot is that the synthetic source treebanks yield a transfer UAS of 52.92 on average. This is not yet a result on held-out test data: recall that 52.92 was the best transfer UAS achieved by any hyperparameter setting. That said, it is 1.00 points better than transferring from the original source treebanks, a significant difference (paired permutation test by language pair, $p < 0.01$).

Figure 2 shows that this average improvement is mainly due to the many cases where the source and target languages come from different families.

Permutation tends to improve source languages that were doing badly to start with. However, it tends to hurt a source language that is already in the target language family.

A hypothetical experiment shows that permuting the source does have good *potential* to help (or at least not hurt) in both cases. The dashed lines in Figure 2—and the scatterplot in Figure 8—show the potential of the method, by showing the improvement we would get from permuting each source treebank using an “oracle” realization policy—the supervised realization parameters θ that are estimated from the actual target treebank. The usefulness of this oracle-permuted source varies depending on the source language, but it is usually much better than the automatically-permuted version of the same source.

This shows that large improvements would be possible if we could only find the best permutation policy allowed by our model family. The question for future work is whether such gains can be achieved by a more sensitive permutation model than (1), a better divergence objective than (4), or a better search algorithm than §4.2. Identifying the best available source treebank, or the best mixture of all source treebanks, would also help greatly.

5.3 Sensitivity to initializer

Figure 2 makes clear that performance of the synthetic source treebanks is strongly correlated with that of their original versions. Most points in Figure 7 lie near the diagonal (Kendall’s $\tau = 0.85$). Even with oracle permutation in Figure 8, the correlation remains strong ($\tau = 0.59$), suggesting that the choice of source treebank is important even beyond its effect on search initialization.

We suspected that when “made to order” source treebanks (more than the oracle versions) have performance close to their original versions, this is in part because the optimizer can get stuck near the initializer (§4.3). To examine this, we experimented with random restarts, as follows. In addition to informed initialization (§4.3), we optimized from 5 other starting points $\theta \sim \mathcal{N}(\mathbf{0}, I)$. From these 6 runs, we selected the final parameters that achieved the best divergence (4). As shown by Figure 9 in the supplement, greater gains appear to be possible with more aggressive search methods of this sort, which we leave to future work. We could also try non-random restarts based on the realization parameters of other languages, as suggested in footnote 10.

5.4 Final evaluation on the test languages

For our final evaluation (§5.1), we use the same hyperparameters (Appendix C) and report on single-source transfer to the 17 held-out treebanks.

The development results hold up in Figure 3. Using the synthetic languages yields 50.36 UAS on average—1.75 points over the baseline, which is significant (paired permutation test, $p < 0.01$).

In the supplementary material (Appendix E), we include some auxiliary experiments on multi-source transfer.

6 Related Work

6.1 Unsupervised parsing

Unsupervised parsing has remained challenging for decades (Mareček, 2016). Classical grammar induction approaches (Lari and Young, 1990; Carroll and Charniak, 1992; Klein and Manning, 2004; Headden III et al., 2009; Naseem et al., 2010) estimate a generative grammar to explain the sentences, for example by the Expectation-Maximization (EM) algorithm, and then use it to parse. Some such approaches try to improve the grammar model. For example, Klein and Manning (2004)’s dependency model with valence was the first to beat a trivial baseline; later improvements considered higher-order effects and punctuation (Headden III et al., 2009; Spitzkovsky et al., 2012). Other approaches try to avoid search error, using strategies like convexified objectives (Wang et al., 2008; Gimpel and Smith, 2012), informed initialization (Klein and Manning, 2004; Mareček and Straka, 2013), search bias (Smith and Eisner, 2005, 2006; Naseem et al., 2010; Gillenwater et al., 2010), branch-and-bound search (Gormley and Eisner, 2013), and switching objectives (Spitzkovsky et al., 2013).

The alternative of cross-lingual transfer has recently flourished thanks to the development of consistent cross-lingual datasets of POS-tagged (Petrov et al., 2012) and dependency-parsed (McDonald et al., 2013) sentences. McDonald et al. (2011) showed a significant improvement over grammar induction by simply using the delexicalized parser trained on other language(s). Subsequent improvements have come from re-weighting source languages (Søgaard, 2011b; Rosa and Žabokrtský, 2015a,b; Wang and Eisner, 2016), adapting the model to the target language using WALS (Dryer and Haspelmath, 2013) features (Naseem et al., 2012; Täckström et al., 2013;




Figure 2: Unlabeled attachment scores (UAS) from 376 pairs of development treebanks. Each column represents a target treebank, and each polyline within that column shows transfer from variants of a different source treebank. The three points on the polyline (from left to right) represent the target UAS for parsers trained on three sources: the original source treebank, the “made to order” permutation that attempts to match surface statistics of the target treebank, and an oracle permutation that uses a realization model trained on the target language. We use solid markers and purple lines if the transfer is *within-family* (source and target treebank from the same language family), and hollow and olive for cross-family transfer. The black polyline in each column is the mean of the others. The table in the lower left gives summary results; the number in each column header gives the number of points summarized. For each column, we boldface the better result between the “Synthetic” and “Original”, or both if they are not significantly different (paired permutation test, $p < 0.01$). We also show the oracle permutation result in row “Oracle”.

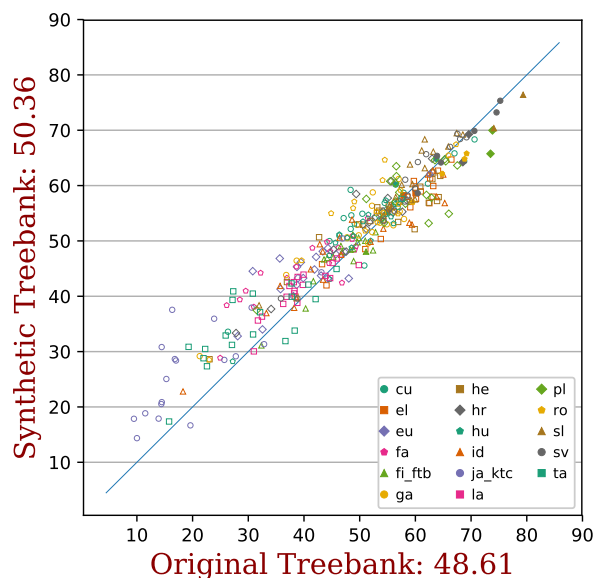


Figure 3: UAS on 337 language pairs from the training languages to the test languages.

Zhang and Barzilay, 2015; Ammar et al., 2016), and improving the lexical representations via multilingual word embeddings (Duong et al., 2015; Guo et al., 2016; Ammar et al., 2016) and synthetic data generation (§6.2).

6.2 Synthetic data generation

Our novel proposal ties into the recent interest in data augmentation in supervised machine learning. In unsupervised parsing, the most widely

adopted synthetic data method has been *annotation projection*, which generates synthetic analyses of target-language sentences by “projecting” the analysis from a source-language translation. Of course, this requires bilingual corpora as an additional resource. Annotation projection was proposed by Yarowsky et al. (2001), gained promising results on sequence labelling tasks, and was later developed for unsupervised parsing (Hwa et al., 2005; Ganchev et al., 2009; Smith and Eisner, 2009; Tiedemann, 2014; Ma and Xia, 2014; Tiedemann et al., 2014). Recent work in this vein has mainly focused on improving the synthetic data, including reweighting the training trees (Agić et al., 2016) or pruning those that cannot be aligned well (Rasooli and Collins, 2015, 2017; Lacroix et al., 2016).

On the other hand, Wang and Eisner (2016) proposed to permute source language treebanks using word order realization models trained on other source languages. They generated on the order of 50,000 synthetic languages by “mixing and matching” a few dozen source languages. Their idea was that with a large set of synthetic languages, they could use them as supervised examples to train an unsupervised structure discovery system that could analyze any new language. Systems built with this dataset were competitive in single-source parser transfer (Wang and Eisner, 2016), typology

prediction (Wang and Eisner, 2017), and parsing unknown languages (Wang and Eisner, 2018).

Our work in this paper differs in that our synthetic treebanks are “made to order.” Rather than combine aspects of different treebanks and hope to get at least one combination that is close to the target language, we “combine” the source treebank with a POS corpus of the target language, which guides our customized permutation of the source.

Beyond unsupervised parsing, synthetic data has been used for several other tasks. In NLP, it has been used for complex tasks such as question-answering (QA) (Serban et al., 2016) and machine reading comprehension (Weston et al., 2016; Hermann et al., 2015; Rajpurkar et al., 2016), where highly expressive neural models are used and not enough real data is available to train them. In the playground of supervised parsing, Gulordava and Merlo (2016) conduct a controlled study on the parsibility of languages by generating treebanks with short dependency length and low variability of word order.

7 Conclusion & Future Work

We have shown how cross-lingual transfer parsing can be improved by permuting the source treebank to better resemble the target language on the surface (in its distribution of gold POS bigrams). The code is available at <https://github.com/wddabc/ordersynthetic>. Our work is grounded in the notion that by trying to explain the POS bigram counts in a target corpus, we can discover a stochastic realization policy for the target language, which correctly “translates” the source trees into appropriate target trees.

We formulated an objective for evaluating such a policy, based on KL-divergence between bigram models. We showed that the objective could be computed efficiently by dynamic programming, thanks to the limitation to bigram statistics.

Experimenting on the Universal Dependencies treebanks v1.2, we showed that the synthetic treebanks were—on average—modestly but significantly better than the corresponding real treebanks for single-source transfer (and in Appendix E, on multi-source transfer).

On the downside, Figure 7 shows that with our current method, permuting the source language to be more like the target language is helpful (on average) only when the source language is from a different language family. This contrast would be

even more striking if we had a better optimizer: Figure 9 shows that SGD’s initialization bias limits its permutation’s benefit for cross-family training, as well as its harm for within-family training.

Several opportunities for future work have already been mentioned throughout the paper. We are also interested in experimenting with richer families of permutation distributions, as well as “conservative” distributions that tend to prefer the original source order. We could use entropy regularization (Grandvalet and Bengio, 2005) to encourage more “deterministic” patterns of realization in the synthetic languages.

We would also like to consider more sensitive divergence measures that go beyond bigrams, for example using recurrent neural network language models (RNNLMs) for \hat{q} and \hat{p}_θ . This means abandoning our exact dynamic programming methods; we would also like to abandon exact exhaustive enumeration in order to drop §4.1’s bounds on n . Fortunately, there exist powerful MCMC methods (Eisner and Tromble, 2006) that can sample from interesting distributions over the space of $n!$ permutations, even for large n . Thus, we could approximately sample from p_θ by drawing permuted versions of each tree in B .

Given this change, a very interesting direction would be to graduate from POS language models to word language models, using cross-lingual unsupervised word embeddings (Ruder et al., 2017). This would eliminate the need for the gold POS tags that we unrealistically assumed in this paper (which are typically unavailable for a low-resource target language). Furthermore, it would enable us to harness richer lexical information beyond the 17 UD POS tags. After all, even a (gold) POS corpus might not be sufficient to determine the word order of the target language: “NOUN VERB NOUN” could be either subject-verb-object or object-verb-subject. However, “water drink boy” is presumably object-verb-subject. Thus, using cross-lingual embeddings, we would try to realize the unordered source trees so that their word strings, with few edits, can achieve high probability under a neural language model of the target.

Acknowledgements

This work was supported by National Science Foundation Grants 1423276 & 1718846. We are grateful to the state of Maryland for the Maryland Advanced Research Computing Center, a crucial resource. We thank Shijie Wu and Adithya Renduchintala for early discussion, Argo lab members for further discussion, and the 3 reviewers for quality comments.

References

- Željko Agić, Anders Johannsen, Barbara Plank, Héctor Martínez Alonso, Natalie Schluter, and Anders Søgaard. 2016. Multilingual projection for parsing truly low-resource languages. *Transactions of the Association for Computational Linguistics*, 4:301–312.
- Waleed Ammar, George Mulcaire, Miguel Ballesteros, Chris Dyer, and Noah Smith. 2016. Many languages, one parser. *Transactions of the Association of Computational Linguistics*, 4:431–444.
- Alena Böhmová, Jan Hajič, Eva Hajičová, and Barbora Hladká. 2003. The Prague dependency treebank. In *Treebanks*, pages 103–127. Springer.
- Glenn Carroll and Eugene Charniak. 1992. Two experiments on learning probabilistic dependency grammars from corpora. In *Statistically-Based Natural Language Processing Techniques: Papers from the Workshop*, pages 1–13, Menlo Park: AAAI Press. AAAI. Technical Report WS-92-01.
- David Chiang. 2007. Hierarchical phrase-based translation. *Computational Linguistics*, 33(2):201–228.
- Matthew S. Dryer and Martin Haspelmath, editors. 2013. *The World Atlas of Language Structures Online*. Max Planck Institute for Evolutionary Anthropology, Leipzig. <http://wals.info/>.
- Long Duong, Trevor Cohn, Steven Bird, and Paul Cook. 2015. Cross-lingual transfer for unsupervised dependency parsing without parallel data. In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*, pages 113–122.
- Jason Eisner and Roy W. Tromble. 2006. Local search with very large-scale neighborhoods for optimal permutations in machine translation. In *Proceedings of the HLT-NAACL Workshop on Computationally Hard Problems and Joint Inference in Speech and Language Processing*, pages 57–75.
- Katja Filippova and Michael Strube. 2009. Tree linearization in English: Improving language model based approaches. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Short Papers*, pages 225–228.
- Richard Futrell and Edward Gibson. 2015. Experiments with generative models for dependency tree linearization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1978–1983.
- Kuzman Ganchev, Jennifer Gillenwater, and Ben Taskar. 2009. Dependency grammar induction via bitext projection constraints. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 369–377.
- Jennifer Gillenwater, Kuzman Ganchev, Joo Graa, Fernando Pereira, and Ben Taskar. 2010. Sparsity in dependency grammar induction. In *Proceedings of the ACL 2010 Conference Short Papers*, pages 194–199.
- Kevin Gimpel and Noah A. Smith. 2012. Concavity and initialization for unsupervised dependency parsing. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 577–581.
- Matthew Gormley and Jason Eisner. 2013. Nonconvex global optimization for latent-variable models. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 444–454.
- Yves Grandvalet and Yoshua Bengio. 2005. Semi-supervised learning by entropy minimization. In L. K. Saul, Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems 17*, pages 529–536. MIT Press.
- Kristina Gulordava and Paola Merlo. 2016. Multilingual dependency parsing evaluation: A large-scale analysis of word order properties using artificial data. *Transactions of the Association for Computational Linguistics*, 4:343–356.
- Jiang Guo, Wanxiang Che, David Yarowsky, Haifeng Wang, and Ting Liu. 2016. A representation learning framework for multi-source transfer parsing. In *AAAI*, pages 2734–2740.
- William P. Headden III, Mark Johnson, and David McClosky. 2009. Improving unsupervised dependency parsing with richer contexts and smoothing. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 101–109.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Advances in Neural Information Processing Systems*, pages 1684–1692.
- Rebecca Hwa, Philip Resnik, Amy Weinberg, Clara Cabezas, and Okan Kolak. 2005. Bootstrapping parsers via syntactic projection across parallel texts. *Natural Language Engineering*, 11(3):311–325.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Eliyahu Kiperwasser and Yoav Goldberg. 2016. Simple and accurate dependency parsing using bidirectional LSTM feature representations. *Transactions of the Association of Computational Linguistics*, 4:313–327.

- Dan Klein and Christopher Manning. 2004. Corpus-based induction of syntactic structure: Models of dependency and constituency. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics*, pages 478–485.
- Ophélie Lacroix, Lauriane Aufrant, Guillaume Wisniewski, and François Yvon. 2016. Frustratingly easy cross-lingual transfer for transition-based dependency parsing. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1058–1063.
- Karim Lari and Steve J. Young. 1990. The estimation of stochastic context-free grammars using the inside-outside algorithm. *Computer Speech and Language*, 4(1):35–56.
- Lillian Lee. 1999. Measures of distributional similarity. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics*, pages 25–32.
- Lillian Lee. 2001. On the effectiveness of the skew divergence for statistical language analysis. In *Proceedings of AISTATS*.
- Xuezhe Ma and Fei Xia. 2014. Unsupervised dependency parsing with transferring distribution via parallel guidance and entropy regularization. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1337–1348.
- David Mareček and Milan Straka. 2013. Stop-probability estimates computed on a large corpus improve unsupervised dependency parsing. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 281–290.
- David Mareček. 2016. Twelve years of unsupervised dependency parsing. In *Proceedings of the 16th ITAT Conference on Information Technologies—Applications and Theory*, pages 56–62.
- Ryan McDonald, Joakim Nivre, Yvonne Quirnbach-Brundage, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith Hall, Slav Petrov, Hao Zhang, Oscar Täckström, Claudia Bedini, Núria Bertomeu Castelló, and Jungmee Lee. 2013. Universal dependency annotation for multilingual parsing. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 92–97.
- Ryan McDonald and Fernando Pereira. 2006. *Discriminative Learning and Spanning Tree Algorithms for Dependency Parsing*. Ph.D. thesis, University of Pennsylvania.
- Ryan McDonald, Slav Petrov, and Keith Hall. 2011. Multi-source transfer of delexicalized dependency parsers. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 62–72.
- Simon Mille, Anja Belz, Bernd Bohnet, Yvette Graham, Emily Pitler, and Leo Wanner. 2018. The first multilingual surface realisation shared task (SR’18): Overview and evaluation results. In *Proceedings of the 1st Workshop on Multilingual Surface Realization (MSR), 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1–12.
- Tahira Naseem, Regina Barzilay, and Amir Globerson. 2012. Selective sharing for multilingual dependency parsing. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 629–637.
- Tahira Naseem, Harr Chen, Regina Barzilay, and Mark Johnson. 2010. Using universal linguistic knowledge to guide grammar induction. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1234–1244.
- Joakim Nivre. 2008. Algorithms for deterministic incremental dependency parsing. *Computational Linguistics*, 34(4):513–553.
- Joakim Nivre, Željko Agić, Lars Ahrenberg, Maria Jesus Aranzabe, Masayuki Asahara, Aitziber Atutxa, Miguel Ballesteros, John Bauer, Kepa Bengoetxea, Riyaz Ahmad Bhat, Eckhard Bick, Cristina Bosco, Gosse Bouma, Sam Bowman, Marie Candito, Gülşen Cebiroğlu Eryiğit, Giuseppe G. A. Celano, Fabricio Chalub, Jinho Choi, Çağrı Çöltekin, Miriam Connor, Elizabeth Davidson, Marie-Catherine de Marneffe, Valeria de Paiva, Arantza Diaz de Ilarraza, Kaja Dobrovoljc, Timothy Dozat, Kira Drozanova, Puneet Dwivedi, Marhaba Eli, Tomaž Erjavec, Richárd Farkas, Jennifer Foster, Cláudia Freitas, Katarína Gajdošová, Daniel Galbraith, Marcos Garcia, Filip Ginter, Iakes Goenaga, Koldo Gojenola, Memduh Gökırmak, Yoav Goldberg, Xavier Gómez Guinovart, Berta González Saavedra, Matias Groni, Normunds Grūzītis, Bruno Guillaume, Nizar Habash, Jan Hajič, Linh Hà Mỹ, Dag Haug, Barbora Hladká, Petter Hohle, Radu Ion, Elena Irimia, Anders Johansson, Fredrik Jørgensen, Hüner Kaşıkara, Hiroshi Kanayama, Jenna Kanerva, Natalia Kotlyba, Simon Krek, Veronika Laippala, Phng Lê Hông, Alessandro Lenci, Nikola Ljubešić, Olga Lyashevskaya, Teresa Lynn, Aibek Makazhanov, Christopher Manning, Cătălina Mărănduc, David Mareček, Héctor Martínez Alonso, André Martins, Jan Mašek, Yuji Matsumoto, Ryan McDonald, Anna Missilä, Verginica Mititelu, Yusuke Miyao, Simonetta Montemagni, Amir More, Shunsuke Mori, Bohdan Moskalevskyi, Kadri Muischnek, Nina Mustafina, Kaili Müürisepp, Lng Nguyễn Thị, Huyễn Nguyễn Thị Minh, Vitaly Nikolaev, Hanna Nurmi, Stina Ojala, Petya Osenova, Lilja Øvrelid, Elena Pascual, Marco Passarotti, Cene Augusto Perez, Guy Perrier, Slav Petrov, Jussi Piitulainen, Barbara Plank, Martin Popel, Lauma Pretkalniņa, Prokopis Prokopidis, Tiina Puolakainen, Sampo Pyysalo, Alexandre Rademaker, Loganathan Ramasamy, Livy Real, Laura Rituma,

- Rudolf Rosa, Shadi Saleh, Manuela Sanguinetti, Baiba Saulīte, Sebastian Schuster, Djamé Seddah, Wolfgang Seeker, Mojgan Seraji, Lena Shakurova, Mo Shen, Dmitry Sichinava, Natalia Silveira, Maria Simi, Radu Simionescu, Katalin Simkó, Mária Šimková, Kiril Simov, Aaron Smith, Alane Suhr, Umut Sulubacak, Zsolt Szántó, Dima Taji, Takaaki Tanaka, Reut Tsarfaty, Francis Tyers, Sumire Uematsu, Larraitz Uribe, Gertjan van Noord, Viktor Varga, Veronika Vincze, Jonathan North Washington, Zdeněk Žabokrtský, Amir Zeldes, Daniel Zeman, and Hanzhi Zhu. 2017. Universal dependencies 2.0. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Joakim Nivre et al. 2015. Universal dependencies 1.2. LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics, Charles University in Prague. Data available at <http://universaldependencies.org>.
- Slav Petrov, Dipanjan Das, and Ryan McDonald. 2012. A universal part-of-speech tagset. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC-2012)*. European Language Resources Association (ELRA).
- Yevgeniy Puzikov and Iryna Gurevych. 2018. BinLin: A simple method of dependency tree linearization. In *Proceedings of the First Workshop on Multilingual Surface Realisation*, pages 13–28.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392.
- Mohammad Sadegh Rasooli and Michael Collins. 2015. Density-driven cross-lingual transfer of dependency parsers. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 328–338.
- Mohammad Sadegh Rasooli and Michael Collins. 2017. Cross-lingual syntactic transfer with limited resources. *Transactions of the Association for Computational Linguistics*, 5:279–293.
- Mohammad Sadegh Rasooli and Joel R. Tetreault. 2015. Yara parser: A fast and accurate dependency parser. *Computing Research Repository*, arXiv:1503.06733 (version 2).
- Rudolf Rosa and Zdeněk Žabokrtský. 2015a. KLcpos3 — a language similarity measure for delexicalized parser transfer. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 243–249.
- Rudolf Rosa and Zdeněk Žabokrtský. 2015b. MST-Parser model interpolation for multi-source delexicalized transfer. In *Proceedings of the 14th International Conference on Parsing Technologies*, pages 71–75.
- Sebastian Ruder, Ivan Vulić, and Anders Søgaard. 2017. A survey of cross-lingual word embedding models. *Computing Research Repository*, arXiv:1706.04902.
- Robert Sedgewick. 1977. Permutation generation methods. *ACM Computing Surveys*, 9(2):137–164.
- Iulian Vlad Serban, Alberto García-Durán, Caglar Gulcehre, Sungjin Ahn, Sarath Chandar, Aaron Courville, and Yoshua Bengio. 2016. Generating factoid questions with recurrent neural networks: The 30M factoid question-answer corpus. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 588–598.
- David A. Smith and Jason Eisner. 2009. Parser adaptation and projection with quasi-synchronous grammar features. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 822–831.
- Noah A. Smith and Jason Eisner. 2005. Guiding unsupervised grammar induction using contrastive estimation. In *International Joint Conference on Artificial Intelligence (IJCAI) Workshop on Grammatical Inference Applications*, pages 73–82.
- Noah A. Smith and Jason Eisner. 2006. Annealing structural bias in multilingual weighted grammar induction. In *Proceedings of the International Conference on Computational Linguistics and the Association for Computational Linguistics (COLING-ACL)*, pages 569–576.
- Anders Søgaard. 2011a. Data point selection for cross-language adaptation of dependency parsers. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 682–686. Association for Computational Linguistics.
- Anders Søgaard. 2011b. Data point selection for cross-language adaptation of dependency parsers. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 682–686.
- Valentin I Spitkovsky, Hiyan Alshawi, and Daniel Jurafsky. 2012. Three dependency-and-boundary models for grammar induction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 688–698. Association for Computational Linguistics.
- Valentin I. Spitkovsky, Hiyan Alshawi, and Daniel Jurafsky. 2013. Breaking out of local optima with count transforms and model recombination: A study

- in grammar induction. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing (EMNLP 2013)*.
- Oscar Täckström, Ryan McDonald, and Joakim Nivre. 2013. Target language adaptation of discriminative transfer parsers. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1061–1071.
- Jörg Tiedemann. 2014. Rediscovering annotation projection for cross-lingual parser induction. In *Proceedings of the 25th International Conference on Computational Linguistics (COLING): Technical Papers*, pages 1854–1864. Dublin City University and Association for Computational Linguistics.
- Jörg Tiedemann, Željko Agić, and Joakim Nivre. 2014. Treebank translation for cross-lingual parser induction. In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning*, pages 130–140.
- Dingquan Wang and Jason Eisner. 2016. The Galactic Dependencies treebanks: Getting more data by synthesizing new languages. *Transactions of the Association of Computational Linguistics*, 4:491–505. Data available at <https://github.com/gdtreebank/gdtreebank>.
- Dingquan Wang and Jason Eisner. 2017. Fine-grained prediction of syntactic typology: Discovering latent structure with *supervised* learning. *Transactions of the Association for Computational Linguistics (TACL)*, 5.
- Dingquan Wang and Jason Eisner. 2018. Surface statistics of an unknown language indicate how to parse it. *Transactions of the Association for Computational Linguistics (TACL)*. To appear.
- Qin Iris Wang, Dale Schuurmans, and Dekang Lin. 2008. Semi-supervised convex training for dependency parsing. In *Proceedings of ACL-HLT*, pages 532–540.
- Jason Weston, Antoine Bordes, Sumit Chopra, and Tomas Mikolov. 2016. Towards AI-complete question answering: A set of prerequisite toy tasks. In *Proceedings of the International Conference on Learning Representations*.
- David Yarowsky, Grace Ngai, and Richard Wicentowski. 2001. Inducing multilingual text analysis tools via robust projection across aligned corpora. In *Proceedings of the First International Conference on Human Language Technology Research*.
- Daniel Zeman and Philip Resnik. 2008. Cross-language parser adaptation between related languages. In *Proceedings of the IJCNLP-08 Workshop on NLP for Less Privileged Languages*.
- Yuan Zhang and Regina Barzilay. 2015. Hierarchical low-rank tensors for multilingual transfer parsing. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1857–1867.