

Unsupervised Dependency Parsing: Let's Use Supervised Parsers

Phong Le and Willem Zuidema

Institute for Logic, Language, and Computation
University of Amsterdam, the Netherlands
{p.le, zuidema}@uva.nl

Abstract

We present a self-training approach to unsupervised dependency parsing that reuses existing supervised and unsupervised parsing algorithms. Our approach, called ‘iterated reranking’ (IR), starts with dependency trees generated by an unsupervised parser, and iteratively improves these trees using the richer probability models used in supervised parsing that are in turn trained on these trees. Our system achieves 1.8% accuracy higher than the state-of-the-part parser of Spitzkovsky et al. (2013) on the WSJ corpus.

1 Introduction

Unsupervised dependency parsing and its supervised counterpart have many characteristics in common: they take as input raw sentences, produce dependency structures as output, and often use the same evaluation metric (DDA, or UAS, the percentage of tokens for which the system predicts the correct head). Unsurprisingly, there has been much more research on supervised parsing – producing a wealth of models, datasets and training techniques – than on unsupervised parsing, which is more difficult, much less accurate and generally uses very simple probability models. Surprisingly, however, there have been no reported attempts to reuse supervised approaches to tackle the unsupervised parsing problem (an idea briefly mentioned in Spitzkovsky et al. (2010b)).

There are, nevertheless, two aspects of supervised parsers that we would like to exploit in an unsupervised setting. First, we can increase the model ex-

pressiveness in order to capture more linguistic regularities. Many recent supervised parsers use third-order (or higher order) features (Koo and Collins, 2010; Martins et al., 2013; Le and Zuidema, 2014) to reach state-of-the-art (SOTA) performance. In contrast, existing models for unsupervised parsing limit themselves to using simple features (e.g., conditioning on heads and valency variables) in order to reduce the computational cost, to identify consistent patterns in data (Naseem, 2014, page 23), and to avoid overfitting (Blunsom and Cohn, 2010). Although this makes learning easier and more efficient, the disadvantage is that many useful linguistic regularities are missed: an upper bound on the performance of such simple models – estimated by using annotated data – is 76.3% on the WSJ corpus (Spitzkovsky et al., 2013), compared to over 93% actual performance of the SOTA supervised parsers.

Second, we would like to make use of information available from lexical semantics, as in Bansal et al. (2014), Le and Zuidema (2014), and Chen and Manning (2014). Lexical semantics is a source for handling rare words and syntactic ambiguities. For instance, if a parser can identify that “he” is a dependent of “walks” in the sentence “He walks”, then, even if “she” and “runs” do not appear in the training data, the parser may still be able to recognize that “she” should be a dependent of “runs” in the sentence “she runs”. Similarly, a parser can make use of the fact that “sauce” and “John” have very different meanings to decide that they have different heads in the two phrases “ate spaghetti with sauce” and “ate spaghetti with John”.

However, applying existing supervised parsing

techniques to the task of unsupervised parsing is, unfortunately, not trivial. The reason is that those parsers are optimally designed for being trained on manually annotated data. If we use existing unsupervised training methods (like EM), learning could be easily misled by a large amount of ambiguity naturally embedded in unannotated training data. Moreover, the computational cost could rapidly increase if the training algorithm is not designed properly. To overcome these difficulties we propose a framework, iterated reranking (IR), where existing supervised parsers are trained without the need of manually annotated data, starting with dependency trees provided by an existing unsupervised parser as initialiser. Using this framework, we can employ the work of Le and Zuidema (2014) to build a new system that outperforms the SOTA unsupervised parser of Spitzkovsky et al. (2013) on the WSJ corpus.

The contribution of this paper is twofold. First, we show the benefit of using lexical semantics for the unsupervised parsing task. Second, our work is a bridge connecting the two research areas unsupervised parsing and its supervised counterpart. Before going to the next section, in order to avoid confusion introduced by names, it is worth noting that we use *un-trained* existing supervised parsers which will be trained on *automatically annotated* treebanks.

2 Related Work

2.1 Unsupervised Dependency Parsing

The first breakthrough was set by Klein and Manning (2004) with their dependency model with valence (DMV), the first model to outperform the right-branching baseline on the DDA metric: 43.2% vs 33.6% on sentences up to length 10 in the WSJ corpus. Nine years later, Spitzkovsky et al. (2013) achieved much higher DDAs: 72.0% on sentences up to length 10, and 64.4% on all sentences in section 23. During this period, many approaches have been proposed to attempt the challenge.

Naseem and Barzilay (2011), Tu and Honavar (2012), Spitzkovsky et al. (2012), Spitzkovsky et al. (2013), and Marecek and Straka (2013) employ extensions of the DMV but with different learning strategies. Naseem and Barzilay (2011) use semantic cues, which are event annotations from an out-of-domain annotated corpus, in their model during

training. Relying on the fact that natural language grammars must be unambiguous in the sense that a sentence should have very few correct parses, Tu and Honavar (2012) incorporate unambiguity regularisation to posterior probabilities. Spitzkovsky et al. (2012) bootstrap the learning by slicing up all input sentences at punctuation. Spitzkovsky et al. (2013) propose a complete deterministic learning framework for breaking out of local optima using count transforms and model recombination. Marecek and Straka (2013) make use of a large raw text corpus (e.g., Wikipedia) to estimate stop probabilities, using the reducibility principle.

Differing from those works, Bisk and Hockenmaier (2012) rely on Combinatory Categorical Grammars with a small number of hand-crafted general linguistic principles; whereas Blunsom and Cohn (2010) use Tree Substitution Grammars with a hierarchical non-parametric Pitman-Yor process prior biasing the learning to a small grammar.

2.2 Reranking

Our work relies on reranking which is a technique widely used in (semi-)supervised parsing. Reranking requires two components: a k -best parser and a reranker. Given a sentence, the parser generates a list of k best candidates, the reranker then rescores those candidates and picks the one that has the highest score. Reranking was first successfully applied to supervised constituent parsing (Collins, 2000; Charniak and Johnson, 2005). It was then employed in the supervised dependency parsing approaches of Sangati et al. (2009), Hayashi et al. (2013), and Le and Zuidema (2014).

Closest to our work is the work series on semi-supervised constituent parsing of McClosky and colleagues, e.g. McClosky et al. (2006), using self-training. They use a k -best generative parser and a discriminative reranker to parse unannotated sentences, then add resulting parses to the training treebank and re-train the reranker. Different from their work, our work is for unsupervised dependency parsing, without manually annotated data, and uses iterated reranking instead of single reranking. In addition, both two components, k -best parser and reranker, are re-trained after each iteration.

3 The IR Framework

Existing training methods for the unsupervised dependency task, such as Blunsom and Cohn (2010), Gillenwater et al. (2011), and Tu and Honavar (2012), are hypothesis-oriented search with the EM algorithm or its variants: training is to move from a point which represents a model hypothesis to another point. This approach is feasible for optimising models using simple features since existing dynamic programming algorithms can compute expectations, which are sums over all possible parses, or to find the best parse in the whole parse space with low complexities. However, the complexity increases rapidly if rich, complex features are used. One way to reduce the computational cost is to use approximation methods like sampling as in Blunsom and Cohn (2010).

3.1 Treebank-oriented Greedy Search

Believing that the difficulty of using EM is from the fact that treebanks are ‘hidden’, leading to the need of computing sum (or max) over all possible treebanks, we propose a greedy local search scheme based on another training philosophy: treebank-oriented search. The key idea is to explicitly search for concrete treebanks which are used to train parsing models. This scheme thus allows supervised parsers to be trained in an unsupervised parsing setting since there is a (automatically annotated) treebank at any time.

Given \mathcal{S} a set of raw sentences, the search space consists of all possible treebanks $\mathcal{D} = \{d(s) | s \in \mathcal{S}\}$ where $d(s)$ is a dependency tree of sentence s . The target of search is the optimal treebank \mathcal{D}^* that is as good as human annotations. Greedy search with this philosophy is as follows: starting at an initial point \mathcal{D}_1 , we pick up a point \mathcal{D}_2 among its neighbours $\mathbf{N}(\mathcal{D}_1)$ such that

$$\mathcal{D}_2 = \arg \max_{\mathcal{D} \in \mathbf{N}(\mathcal{D}_1)} f_{\mathcal{D}_1}(\mathcal{D}) \quad (1)$$

where $f_{\mathcal{D}_1}(\mathcal{D})$ is an objective function measuring the goodness of \mathcal{D} (which may or may not be conditioned on \mathcal{D}_1). We then continue this search until some stop criterion is satisfied. The crucial factor here is to define $\mathbf{N}(\mathcal{D}_i)$ and $f_{\mathcal{D}_i}(\mathcal{D})$. Below are two special cases of this scheme.

Semi-supervised parsing using reranking (McClosky et al., 2006). This reranking is indeed one-step greedy local search. In this scenario, $\mathbf{N}(\mathcal{D}_1)$ is the Cartesian product of k -best lists generated by a k -best parser, and $f_{\mathcal{D}_i}(\mathcal{D})$ is a reranker.

Unsupervised parsing with hard-EM (Spitkovsky et al., 2010b) In hard-EM, the target is to maximise the following objective function with respect to a parameter set Θ

$$L(\mathcal{S}|\Theta) = \sum_{s \in \mathcal{S}} \max_{d \in Dep(s)} \log P_{\Theta}(d) \quad (2)$$

where $Dep(s)$ is the set of all possible dependency structures of s . The two EM steps are thus

- Step 1: $\mathcal{D}_{i+1} = \arg \max_{\mathcal{D}} P_{\Theta_i}(\mathcal{D})$
- Step 2: $\Theta_{i+1} = \arg \max_{\Theta} P_{\Theta}(\mathcal{D}_{i+1})$

In this case, $\mathbf{N}(\mathcal{D}_i)$ is the whole treebank space and $f_{\mathcal{D}_i}(\mathcal{D}) = P_{\Theta_i}(\mathcal{D}) = P_{\arg \max_{\Theta} P_{\Theta}(\mathcal{D}_i)}(\mathcal{D})$.

3.2 Iterated Reranking

We instantiate the greedy search scheme by iterated reranking which requires two components: a k -best parser P , and a reranker R . Firstly, \mathcal{D}_1 is used to train these two components, resulting in P_1 and R_1 . The parser P_1 then generates a set of lists of k candidates ${}_k\mathcal{D}_1$ (whose Cartesian product results in $\mathbf{N}(\mathcal{D}_1)$) for the set of training sentences \mathcal{S} . The best candidates, according to reranker R_1 , are collected to form \mathcal{D}_2 for the next iteration. This process is halted when a pre-defined stop criterion is met.¹

It is certain that we can, as in the work of Spitkovsky et al. (2010b) and many bootstrapping approaches, employ only parser P . Reranking, however, brings us two benefits. First, it allows us to employ very expressive models like the ∞ -order generative model proposed by Le and Zuidema (2014). Second, it embodies a similar idea to co-training (Blum and Mitchell, 1998): P and R play roles as two views of the data.

¹It is worth noting that, although $\mathbf{N}(\mathcal{D}_i)$ has the size $O(k^n)$ where n is the number of sentences, reranking only needs to process $O(k \times n)$ parses if these sentences are assumed to be independent.

3.3 Multi-phase Iterated Reranking

Training in machine learning often uses *starting big* which is to use up all training data at the same time. However, Elman (1993) suggests that in some cases, learning should start by training simple models on small data and then gradually increase the model complexity and add more difficult data. This is called *starting small*.

In unsupervised dependency parsing, starting small is intuitive. For instance, given a set of long sentences, learning the fact that the head of a sentence is its main verb is difficult because a long sentence always contains many syntactic categories. It would be much easier if we start with only length-one sentences, e.g. “Look!”, since there is only one choice which is usually a verb. This training scheme was successfully applied by Spitzkovsky et al. (2010a) under the name: Baby Step.

We adopt starting small to construct the multi-phase iterated reranking (MPIR) framework. In phase 0, a parser M with a simple model is trained on a set of short sentences $\mathcal{S}^{(0)}$ as in traditional approaches. This parser is used to parse a larger set of sentences $\mathcal{S}^{(1)} \supseteq \mathcal{S}^{(0)}$, resulting in $\mathcal{D}_1^{(1)}$. $\mathcal{D}_1^{(1)}$ is then used as the starting point for the iterated reranking in phase 1. We continue this process until phase N finishes, with $\mathcal{S}^{(i)} \supseteq \mathcal{S}^{(i-1)}$ ($i = 1..N$). In general, we use the resulting reranker in the previous phase to generate the starting point for the iterated reranking in the current phase.

4 Le and Zuidema (2014)’s Reranker

Le and Zuidema (2014)’s reranker is an exception among supervised parsers because it employs an extremely *expressive* model whose features are ∞ -order². To overcome the problem of sparsity, they introduced the inside-outside recursive neural network (IORNN) architecture that can estimate tree-generating models including those proposed by Eisner (1996) and Collins (2003a).

4.1 The ∞ -order Generative Model

Le and Zuidema (2014)’s reranker employs the generative model proposed by Eisner (1996). Intuitively, this model is top-down: starting with ROOT,

²In fact, the order is finite but unbound.




Figure 1: Inside-Outside Recursive Neural Network (IORNN). Black/white rectangles correspond to inner/outer representations.

we generate its left dependents and its right dependents. We then generate dependents for each ROOT’s dependent. The generative process recursively continues until there is no dependent to generate. Formally, this model is described by the following formula

$$P(d(H)) = \prod_{l=1}^L P(H_l^L | \mathcal{C}(H_l^L)) P(d(H_l^L)) \times \prod_{r=1}^R P(H_r^R | \mathcal{C}(H_r^R)) P(d(H_r^R)) \quad (3)$$

where H is the current head, $d(N)$ is the fragment of the dependency parse rooted at N , and $\mathcal{C}(N)$ is the context to generate N . H^L, H^R are respectively H ’s left dependents and right dependents, plus *EOC* (End-Of-Children), a special token to inform that there are no more dependents to generate. Thus, $P(d(\text{ROOT}))$ is the probability of generating the entire dependency structure d .

Le and Zuidema’s ∞ -order generative model is defined as Eisner’s model in which the context $\mathcal{C}^\infty(D)$ to generate D contains *all* of D ’s generated siblings, its ancestors and their siblings. Because of very large fragments that contexts are allowed to hold, traditional count-based methods are impractical (even if we use smart smoothing techniques). They thus introduced the IORNN architecture to estimate the model.

4.2 Estimation with the IORNN

An IORNN (Figure 1) is a recursive neural network whose topology is a tree. What make this network different from traditional RNNs (Socher et al., 2010) is that each tree node u carries two vectors: \mathbf{i}_u - the inner representation, represents the content of the

phrase covered by the node, and \mathbf{o}_u - the outer representation, represents the context around that phrase. In addition, information in an IORNN is allowed to flow not only bottom-up as in RNNs, but also top-down. That makes IORNNs a natural tool for estimating top-down tree-generating models.

Applying the IORNN architecture to dependency parsing is straightforward, along the generative story of the ∞ -order generative model. First of all, the “inside” part of this IORNN is simpler than what is depicted in Figure 1: the inner representation of a phrase is assumed to be the inner representation of its head. This approximation is plausible since the meaning of a phrase is often dominated by the meaning of its head. The inner representation at each node, in turn, is a function of a vector representation for the word (in our case, the word vectors are initially borrowed from Collobert et al. (2011)), the POS-tag and capitalisation feature.

Without loss of generality and ignoring directions for simplicity, they assume that the model is generating dependent u for node h conditioning on context $\mathcal{C}^\infty(u)$ which contains all of u ’s ancestors (including h) and their siblings, and all of previously generated u ’s sisters. Now there are two types of contexts: *full* contexts of heads (e.g., h) whose dependents are being generated, and contexts to generate nodes (e.g., $\mathcal{C}^\infty(u)$). Contexts of the first type are clearly represented by outer representations. Contexts of the other type are represented by *partial outer representations*, denoted by $\bar{\mathbf{o}}_u$. Because the context to generate a node can be constructed recursively by combining the full context of its head and its previously generated sisters, they can compute $\bar{\mathbf{o}}_u$ as a function of \mathbf{o}_h and the inner representations of its previously generated sisters. On the top of $\bar{\mathbf{o}}_u$, they put a softmax layer to estimate the probability $P(x|\mathcal{C}^\infty(u))$.

Training this IORNN is to minimise the cross entropy over all dependents. This objective function is indeed the negative log likelihood $P(\mathcal{D})$ of training treebank \mathcal{D} .

4.3 The Reranker

Le and Zuidema’s (generative) reranker is given by

$$d^* = \arg \max_{d \in_k Dep(s)} P(d)$$

where P (Equation 3) is computed by the ∞ -order generative model which is estimated by an IORNN; and $_k Dep(s)$ is a k -best list.

5 Complete System

Our system is based on the multi-phase IR. In general, any third-party parser for unsupervised dependency parsing can be used in phase 0, and any third-party parser that can generate k -best lists can be used in the other phases. In our experiments, for phase 0, we choose the parser using an extension of the DMV model with stop-probability estimates computed on a large corpus proposed by Marecek and Straka (2013). This system has a moderate performance³ on the WSJ corpus: 57.1% vs the SOTA 64.4% DDA of Spitkovsky et al. (2013). For the other phases, we use the MSTParser⁴ (with the second-order feature mode) (McDonald and Pereira, 2006).

Our system uses Le and Zuidema (2014)’s reranker (Section 4.3). It is worth noting that, in this case, each phase with iterated reranking could be seen as an approximation of hard-EM (see Equation 2) where the first step is replaced by

$$\mathcal{D}_{i+1} = \arg \max_{\mathcal{D} \in \mathbf{N}(\mathcal{D}_i)} P_{\Theta_i}(\mathcal{D}) \quad (4)$$

In other words, instead of searching over the treebank space, the search is limited in a neighbour set $\mathbf{N}(\mathcal{D}_i)$ generated by k -best parser P_i .

5.1 Tuning Parser P

Parser P_i trained on \mathcal{D}_i defines neighbour set $\mathbf{N}(\mathcal{D}_i)$ which is the Cartesian product of the k -best lists in $_k \mathcal{D}_i$. The position and shape of $\mathbf{N}(\mathcal{D}_i)$ is thus determined by two factors: how well P_i can fit \mathcal{D}_i , and k . Intuitively, the lower the fitness is, the more $\mathbf{N}(\mathcal{D}_i)$ goes far away from \mathcal{D}_i ; and the larger k is, the larger

³Marecek and Straka (2013) did not report any experimental result on the WSJ corpus. We use their source code at <http://ufal.mff.cuni.cz/udp> with the setting presented in Section 6.1. Because the parser does not provide the option to parse unseen sentences, we merge the training sentences (up to length 15) to all the test sentences to evaluate its performance. Note that this result is close to the DDA (55.4%) that the authors reported on CoNLL 2007 English dataset, which is a portion of the WSJ corpus.

⁴<http://sourceforge.net/projects/mstparser/>

$\mathbf{N}(\mathcal{D}_i)$ is. Moreover, the diversity of $\mathbf{N}(\mathcal{D}_i)$ is inversely proportional to the fitness. When the fitness decreases, patterns existing in the training treebank become less certain to the parser, patterns that do not exist in the training treebank thus have more chances to appear in k -best candidates. This leads to high diversity of $\mathbf{N}(\mathcal{D}_i)$. We blindly set $k = 10$ in all of our experiments.

With the MSTParser, there are two hyperparameters: `itersMST`, the number of epochs, and `training-kMST`, the k -best parse set size to create constraints during training. `training-kMST` is always 1 because constraints from k -best parses with almost incorrect training parses are useless.

Because `itersMST` controls the fitness of the parser to training treebank \mathcal{D}_i , it, as pointed out above, determines the distance from $\mathbf{N}(\mathcal{D}_i)$ to \mathcal{D}_i and the diversity of the former. Therefore, if we want to encourage the local search to explore more distant areas, we should set `itersMST` low. In our experiments, we test two strategies: (i) MaxEnc, `itersMST = 1`, maximal encouragement, and (ii) MinEnc, `itersMST = 10`, minimal encouragement.

5.2 Tuning Reranker R

Tuning the reranker R is to set values for `dimIORNN`, the dimensions of inner and outer representations, and `itersIORNN`, the number of epochs to train the IORNN. Because the ∞ -order model is very expressive and feed-forward neural networks are universal approximators (Cybenko, 1989), the reranker is capable of perfectly remembering all training parses. In order to avoid this, we set `dimIORNN = 50`, and set `itersIORNN = 5` for *very* early stopping.

5.3 Tuning multi-phase IR

Because Marecek and Straka (2013)’s parser does not distinguish training data from test data, we postulate $\mathcal{S}_0 = \mathcal{S}_1$. Our system has N phases such that $\mathcal{S}_0, \mathcal{S}_1$ contain all sentences up to length $l_1 = 15$, \mathcal{S}_i ($i = 2..N$) contains all sentences up to length $l_i = l_{i-1} + 1$, and \mathcal{S}_N contains all sentences up to length 25. Phase 1 halts after 100 iterations whereas all the following phases run with one iteration. Note that we force the local search in phase 1 to run intensively because we hypothesise that most of the important patterns for dependency parsing can be found within short sentences.

6 Experiments

6.1 Setting

We use the Penn Treebank WSJ corpus: sections 02-21 for training, and section 23 for testing. We then apply the standard pre-processing⁵ for unsupervised dependency parsing task (Klein and Manning, 2004): we strip off all empty sub-trees, punctuation, and terminals (tagged # and \$) not pronounced where they appear; we then convert the remaining trees to dependencies using Collins’s head rules (Collins, 2003b). Both word forms and gold POS tags are used. The directed dependency accuracy (DDA) metric is used for evaluation.

The vocabulary is taken as a list of words occurring more than two times in the training data. All other words are labelled ‘UNKNOWN’ and every digit is replaced by ‘0’. We initialise the IORNN with the 50-dim word embeddings from Collobert et al. (2011)⁶, and train it with the learning rate 0.1,

6.2 Results

We compare our system against recent systems (Table 1 and Section 2.1). Our system with the two encouragement levels, MinEnc and MaxEnc, achieves the highest reported DDAs on section 23: 1.8% and 1.2% higher than Spitzkovsky et al. (2013) on all sentences and up to length 10, respectively. Our improvements over the system’s initialiser (Marecek and Straka, 2013) are 9.1% and 4.4%.

6.3 Analysis

In this section, we analyse our system along two aspects. First, we examine three factors which determine the performance of the whole system: encouragement level, lexical semantics, and starting point. We then search for what IR (with the MaxEnc option) contributes to the overall performance by comparing the quality of the treebank resulted in the end of phase 1 against the quality of the treebank given by its initialier, i.e. Marecek and Straka (2013).

The effect of encouragement level

Figure 2 shows the differences in DDA between using MaxEnc and MinEnc in each phase: we com-

⁵<http://www.cs.famaf.unc.edu.ar/~francolq/en/proyectos/dmccm>

⁶<http://ml.nec-labs.com/senna/>. These word embeddings were *unsupervisedly* learnt from Wikipedia.

System	DDA (@10)
Bisk and Hockenmaier (2012)	53.3 (71.5)
Blunsom and Cohn (2010)	55.7 (67.7)
Tu and Honavar (2012)	57.0 (71.4)
Marecek and Straka (2013)	57.1 (68.8)
Naseem and Barzilay (2011)	59.4 (70.2)
Spitkovsky et al. (2012)	61.2 (71.4)
Spitkovsky et al. (2013)	64.4 (72.0)
Our system (MinEnc)	66.2 (72.7)
Our system (MaxEnc)	65.8 (73.2)

Table 1: Performance on section 23 of the WSJ corpus (all sentences and up to length 10) for recent systems and our system. MinEnc and MaxEnc denote $iters_{SMST} = 10$ and $iters_{SMST} = 1$ respectively.




Figure 2: $DDA_{MaxEnc} - DDA_{MinEnc}$ of all phases on their training sets (e.g., phase 3 with $\mathcal{S}^{(3)}$ containing all training sentences up to length 17).

pute $DDA_{MaxEnc} - DDA_{MinEnc}$ of each phase on its training set (e.g., phase 3 with $\mathcal{S}^{(3)}$ containing all training sentences up to length 17). MinEnc outperforms MaxEnc within phases 1, 2, 3, and 4. However, from phase 5, the latter surpasses the former. It suggests that exploring areas far away from the current point with long sentences is risky. The reason is that long sentences contain more ambiguities than short ones; thus rich diversity, high difference from the current point, but small size (i.e., small k) could easily lead the learning to a wrong path.

The performance of the system with the two encouragement levels on section 23 (Table 1) also suggests the same. MaxEnc strategy helps the system achieve the highest accuracy on short sentences (up to length 10). However, it is less helpful than MinEnc when performing on long sentences.




Figure 3: DDA of phase 1 (MaxEnc), with and without the word embeddings (denoted by w/ sem and wo/ sem, respectively), on training sentences up to length 15 (i.e. $\mathcal{S}^{(1)}$).




Figure 4: DDA of phase 1 (MaxEnc) before and after training with three different starting points provided by three parsers used in phase 0: MS (Marecek and Straka, 2013), GGGPT (Gillenwater et al., 2011), and Harmonic (Klein and Manning, 2004).

The role of lexical semantics

We examine the role of the lexical semantics, which is given by the word embeddings. Figure 3 shows DDAs on training sentences up to length 15 (i.e. $\mathcal{S}^{(1)}$) of phase 1 (MaxEnc) with and without the word-embeddings. With the word-embeddings, phase 1 achieves 71.11%. When the word-embeddings are not given, i.e. the IORNN uses randomly generated word vectors, the accuracy drops 4.2%. It shows that lexical semantics plays a decisive role in the performance of the system.

However, it is worth noting that, even without that knowledge (i.e., with the ∞ -order generative model alone), the DDA of phase 1 is 2% higher than before being trained (66.89% vs 64.9%). It suggests that phase 1 is capable of discovering some useful dependency patterns that are invisible to the parser in phase 0. This, we conjecture, is thanks to high-order features captured by the IORNN.

The importance of the starting point

Starting point is claimed to be important in local search. We examine this by using three different parsers in phase 0: (i) MS (Marecek and Straka,




Figure 5: Precision (top) and recall (bottom) over binned HEAD distance of iterated reranking (IR) and its initializer (MS) on the training sentences in phase 1 (≤ 15 words).

2013), the parser used in the previous experiments, (ii) GGGPT (Gillenwater et al., 2011)⁷ employing an extension of the DMV model and posterior regularization framework for training, and (iii) Harmonic, the harmonic initializer proposed by Klein and Manning (2004).

Figure 4 shows DDAs of phase 1 (MaxEnc) on training sentences up to length 15 with three starting-points given by those parsers. Starting point is clearly very important to the performance of the iterated reranking: the better the starting point is, the higher performance phase 1 has. However, a remarkable point here is that the iterated reranking of phase 1 always finds out more useful patterns for parsing whatever the starting point is in this experiment. It is certainly due to the high order features and lexical semantics, which are not exploited in those parsers.

The contribution of Iterated Reranking

We compare the quality of the treebank resulted in the end of phase 1 against the quality of the treebank given by the initializer Marecek and Straka (2013). Figure 5 shows precision (top) and recall (bottom)

over binned HEAD distance. IR helps to improve the precision on all distance bins, especially on the bins corresponding to long distances (≥ 3). The recall is also improved, except on the bin corresponding to ≥ 7 (but the F1-score on this bin is increased). We attribute this improvement to the ∞ -order model which uses very large fragments as contexts thus be able to capture long dependencies.

Figure 6 shows the correct-head accuracies over POS-tags. IR helps to improve the accuracies over almost all POS-tags, particularly nouns (e.g. NN, NNP, NNS), verbs (e.g. VBD, VBZ, VBN, VBG) and adjectives (e.g. JJ, JJR). However, as being affected by the initializer, IR performs poorly on conjunction (CC) and modal auxiliary (MD). For instance, in the treebank given by the initializer, almost all modal auxiliaries are dependents of their verbs instead of the other way around.

7 Discussion

Our system is different from the other systems shown in Table 1 as it uses an extremely expressive model, the ∞ -order generative model, in which conditioning contexts are very large fragments. Only the work of Blunsom and Cohn (2010), whose resulting grammar rules can contain large tree fragments, shares this property. The difference is that their work needs a pre-defined prior, namely hierarchical non-parametric Pitman-Yor process prior, to avoid large, rare fragments and for smoothing. The IORN of our system, in contrast, does that automatically. It learns by itself how to deal with distant conditioning nodes, which are often less informative than close conditioning nodes on computing $P(x|\mathcal{C}^\infty(u))$. In addition, smoothing is given free: recursive neural nets are able to map ‘similar’ fragments onto close points (Socher et al., 2010) thus an unseen fragment tends to be mapped onto a point close to points corresponding to ‘similar’ seen fragments.

Another difference is that our system exploits lexical semantics via word embeddings, which were learnt unsupervisedly. By initialising the IORN with these embeddings, the use of this knowledge turns out easy and transparent. Spitkovsky et al. (2013) also exploit lexical semantics but in a limited way, using a context-based polysemous unsuper-

⁷code.google.com/p/pr-toolkit




Figure 6: Correct-head accuracies over POS-tags (sorted in the descending order by frequency) of iterated reranking (IR) and its initializer (MS) on the training sentences in phase 1 (≤ 15 words).

vised clustering method to tag words. Although their approach can distinguish polysemes (e.g., ‘cool’ in ‘to cool the selling panic’ and in ‘it is cool’), it is not able to make use of word meaning similarities (e.g., the meaning of ‘dog’ is closer to ‘animal’ than to ‘table’). Naseem and Barzilay (2011)’s system uses semantic cues from an out-of-domain annotated corpus, thus is not fully unsupervised.

We have showed that IR with a generative reranker is an approximation of hard-EM (see Equation 4). Our system is thus related to the works of Spitzkovsky et al. (2013) and Tu and Honavar (2012). However, what we have proposed is more than that: IR is a general framework that we can have more than one option for choosing k -best parser and reranker. For instance, we can make use of a generative k -best parser and a discriminative reranker that are used for supervised parsing. Our future work is to explore this.

The experimental results reveal that starting point is very important to the iterated reranking with the ∞ -order generative model. On the one hand, that is a disadvantage compared to the other systems, which use uninformed or harmonic initialisers. But on the other hand, that is an innovation as our approach is capable of making use of existing systems. The results shown in Figure 4 suggest that if phase 0 uses a better parser which uses less expressive model and/or less external knowledge than our model, such as the one proposed by Spitzkovsky et al. (2013), we can expect even a higher performance. The other systems, except Blunsom and Cohn (2010), however, might not benefit from using good existing

parsers as initializers because their models are not significantly more expressive than others⁸.

8 Conclusion

We have proposed a new framework, iterated reranking (IR), which trains supervised parsers without the need of manually annotated data by using a unsupervised parser as an initialiser. Our system, employing Marecek and Straka (2013)’s unsupervised parser as the initialiser, the k -best MSTParser, and Le and Zuidema (2014)’s reranker, achieved 1.8% DDA higher than the SOTA parser of Spitzkovsky et al. (2013) on the WSJ corpus. Moreover, we also showed that unsupervised parsing benefits from lexical semantics through using word-embeddings.

Our future work is to exploit other existing supervised parsers that fit our framework. Besides, taking into account the fast development of the word embedding research (Mikolov et al., 2013; Pennington et al., 2014), we will try different word embeddings.

Acknowledgments

We thank Remko Scha and three anonymous reviewers for helpful comments. Le thanks Miloš Stanojević for helpful discussion.

⁸In an experiment, we used the Marecek and Straka (2013)’s parser as an initializer for the Gillenwater et al. (2011)’s parser. As we expected, the latter was not able to make use of this.

References

- Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2014. Tailoring continuous word representations for dependency parsing. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics*.
- Yonatan Bisk and Julia Hockenmaier. 2012. Simple robust grammar induction with combinatory categorial grammars. In *AAAI*.
- Avrim Blum and Tom M. Mitchell. 1998. Combining labeled and unlabeled data with co-training. In *COLT*, pages 92–100.
- Phil Blunsom and Trevor Cohn. 2010. Unsupervised induction of tree substitution grammars for dependency parsing. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1204–1213. Association for Computational Linguistics.
- Eugene Charniak and Mark Johnson. 2005. Coarse-to-fine n-best parsing and maxent discriminative reranking. In *ACL*.
- Danqi Chen and Christopher D Manning. 2014. A fast and accurate dependency parser using neural networks. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Michael Collins. 2000. Discriminative reranking for natural language parsing. In *ICML*, pages 175–182.
- Michael Collins. 2003a. Head-driven statistical models for natural language parsing. *Computational linguistics*, 29(4):589–637.
- Michael Collins. 2003b. Head-driven statistical models for natural language parsing. *Computational Linguistics*, 29(4):589–637.
- Ronan Collobert, Jason Weston, Léon Bottou, Michael Karlen, Koray Kavukcuoglu, and Pavel Kuksa. 2011. Natural language processing (almost) from scratch. *The Journal of Machine Learning Research*, 12:2493–2537.
- George Cybenko. 1989. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314.
- Jason M Eisner. 1996. Three new probabilistic models for dependency parsing: An exploration. In *Proceedings of the 16th conference on Computational linguistics-Volume 1*, pages 340–345. Association for Computational Linguistics.
- Jeffrey L Elman. 1993. Learning and development in neural networks: The importance of starting small. *Cognition*, 48(1):71–99.
- Jennifer Gillenwater, Kuzman Ganchev, João Graça, Fernando Pereira, and Ben Taskar. 2011. Posterior sparsity in unsupervised dependency parsing. *The Journal of Machine Learning Research*, 12:455–490.
- Katsuhiko Hayashi, Shuhei Kondo, and Yuji Matsumoto. 2013. Efficient stacked dependency parsing by forest reranking. *Transactions of the Association for Computational Linguistics*, 1(1):139–150.
- Dan Klein and Christopher D. Manning. 2004. Corpus-based induction of syntactic structure: Models of dependency and constituency. In *ACL*, pages 478–485.
- Terry Koo and Michael Collins. 2010. Efficient third-order dependency parsers. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1–11. Association for Computational Linguistics.
- Phong Le and Willem Zuidema. 2014. The inside-outside recursive neural network model for dependency parsing. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- David Marecek and Milan Straka. 2013. Stop-probability estimates computed on a large corpus improve unsupervised dependency parsing. In *ACL (1)*, pages 281–290.
- André FT Martins, Miguel B Almeida, and Noah A Smith. 2013. Turning on the turbo: Fast third-order non-projective turbo parsers. In *Proc. of ACL*.
- David McClosky, Eugene Charniak, and Mark Johnson. 2006. Effective self-training for parsing. In *Proceedings of the main conference on human language technology conference of the North American Chapter of the Association of Computational Linguistics*, pages 152–159. Association for Computational Linguistics.
- Ryan T. McDonald and Fernando C. N. Pereira. 2006. Online learning of approximate dependency parsing algorithms. In *EACL*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, pages 3111–3119.
- Tahira Naseem and Regina Barzilay. 2011. Using semantic cues to learn syntax. In *AAAI*.
- Tahira Naseem. 2014. *Linguistically Motivated Models for Lightly-Supervised Dependency Parsing*. Ph.D. thesis, Massachusetts Institute of Technology.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. *Proceedings of the Empirical Methods in Natural Language Processing (EMNLP 2014)*, 12.
- Federico Sangati, Willem Zuidema, and Rens Bod. 2009. A generative re-ranking model for dependency parsing. In *Proceedings of the 11th International Conference on Parsing Technologies*, pages 238–241. Association for Computational Linguistics.

- Richard Socher, Christopher D. Manning, and Andrew Y. Ng. 2010. Learning continuous phrase representations and syntactic parsing with recursive neural networks. In *Proceedings of the NIPS-2010 Deep Learning and Unsupervised Feature Learning Workshop*.
- Valentin I. Spitzkovsky, Hiyan Alshawi, and Daniel Jurafsky. 2010a. From Baby Steps to Leapfrog: How “Less is More” in unsupervised dependency parsing. In *Proc. of NAACL-HLT*.
- Valentin I. Spitzkovsky, Hiyan Alshawi, Daniel Jurafsky, and Christopher D. Manning. 2010b. Viterbi training improves unsupervised dependency parsing. In *Proceedings of the Fourteenth Conference on Computational Natural Language Learning (CoNLL-2010)*.
- Valentin I. Spitzkovsky, Hiyan Alshawi, and Daniel Jurafsky. 2012. Bootstrapping dependency grammar inducers from incomplete sentence fragments via austere models. In *Proceedings of the 11th International Conference on Grammatical Inference*.
- Valentin I. Spitzkovsky, Hiyan Alshawi, and Daniel Jurafsky. 2013. Breaking out of local optima with count transforms and model recombination: A study in grammar induction. In *EMNLP*, pages 1983–1995.
- Kewei Tu and Vasant Honavar. 2012. Unambiguity regularization for unsupervised learning of probabilistic grammars. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1324–1334. Association for Computational Linguistics.