

Improved Latent Tree Induction with Distant Supervision via Span Constraints

Zhiyang Xu*, Andrew Drozdov*, Jay Yoon Lee, Tim O’Gorman, Subendhu Rongali, Dylan Finkbeiner, Shilpa Suresh, Mohit Iyyer, and Andrew McCallum

College of Information and Computer Sciences
University of Massachusetts Amherst

{zhiyangxu, adrozdov, jaylee, togorman, srongali, dfinkbeiner, ssuresh, miyyer, mccallum}@cs.umass.edu

Abstract

For over thirty years, researchers have developed and analyzed methods for latent tree induction as an approach for unsupervised syntactic parsing. Nonetheless, modern systems still do not perform well enough compared to their supervised counterparts to have any practical use as structural annotation of text. In this work, we present a technique that uses distant supervision in the form of span constraints (i.e. phrase bracketing) to improve performance in unsupervised constituency parsing. Using a relatively small number of span constraints we can substantially improve the output from DIORA, an already competitive unsupervised parsing system. Compared with full parse tree annotation, span constraints can be acquired with minimal effort, such as with a lexicon derived from Wikipedia, to find exact text matches. Our experiments show span constraints based on entities improves constituency parsing on English WSJ Penn Treebank by more than 5 F1. Furthermore, our method extends to any domain where span constraints are easily attainable, and as a case study we demonstrate its effectiveness by parsing biomedical text from the CRAFT dataset.

1 Introduction

Syntactic parse trees are helpful for various downstream tasks such as speech recognition (Moore et al., 1995), machine translation (Akoury et al., 2019), paraphrase generation (Iyyer et al., 2018), semantic parsing (Xu et al., 2020), and information extraction (Naradowsky, 2014). While supervised syntactic parsers are state-of-the-art models for creating these parse trees, their performance does not transfer well across domains. Moreover, new syntactic annotations are prohibitively expensive; the original Penn Treebank required eight years of annotation (Taylor et al., 2003), and expanding PTB annotation to a new domain can be

* Equal contribution.

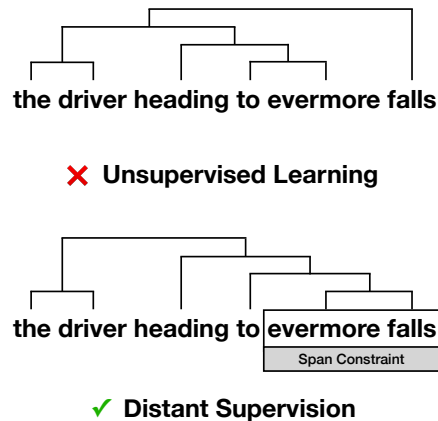


Figure 1: An example sentence and parsing to illustrate distant supervision via span constraints. *Top*: The unsupervised parser predicts a parse tree, but due to natural ambiguity in the text the prediction crosses with a known constraint. *Bottom*: By incorporating the span constraint, the prediction improves and, as a result, recovers the ground truth parse tree. In our experiments, we both inject span constraints directly into parse tree decoding and separately use the constraints only for distant supervision at training time. We find the latter approach is typically more effective.

a large endeavor. For example, the 20k sentences of biomedical treebanking in the CRAFT corpus required 80 annotator hours per week for 2.5 years, include 6 months for annotator training (Verspoor et al., 2011). However, although many domains and many languages lack full treebanks, they do often have access to other annotated resources such as NER, whose spans might provide some partial syntactic supervision. We explore whether *unsupervised* parsing methods can be enhanced with distant supervision from such spans to enable the types of benefits afforded by supervised syntactic parsers without the need for expensive syntactic annotations.

We aim to “bridge the gap” between supervised and unsupervised parsing with distant supervision through span constraints. These span constraints in-

dicates that a certain sequence of words in a sentence form a constituent span in its parse tree, and we obtain these partial ground-truths without explicit user annotation. We take inspiration from previous work incorporating distant supervision into parsing (Haghighi and Klein, 2006; Finkel and Manning, 2009; Ganchev et al., 2010; Cao et al., 2020), and design a novel fully neural system that improves a competitive neural unsupervised parser (DIORA; Drozdov et al. 2019) using span constraints defined on a portion of the training data. In the large majority of cases, the number of spans constraints per sentence is much lower than that specified by a full parse tree. We find that entity spans are effective as constraints, and can readily be acquired from existing data or derived from a gazetteer.

In our experiments, we use DIORA as our baseline and improve upon it by injecting these span constraints as a source of distant supervision. We introduce a new method for training DIORA that leverages the structured SVM loss often used in supervised constituency parsing (Stern et al., 2017; Kitaev and Klein, 2018), but only depends on partial structure. We refer to this method as partially structured SVM (PS-SVM). Our experiments indicate PS-SVM improves upon unsupervised parsing performance as the model adjusts its prediction to incorporate span constraints (depicted in Figure 1). Using ground-truth entities from Ontonotes (Pradhan et al., 2012) as constraints, we achieve more than 5 F1 improvement over DIORA when parsing English WSJ Penn Treebank (Marcus et al., 1993). Using automatically extracted span constraints from an entity-based lexicon (i.e. gazetteer) is an easy alternative to ground truth annotation and gives 2 F1 improvement over DIORA. Importantly, training DIORA with PS-SVM is more effective than simply injecting available constraints into parse tree decoding at test time. We also conduct experiments with different types of span constraints. Our detailed analysis shows that entity-based constraints are similarly useful as the same number of ground truth NP constituent constraints. Finally, we show that DIORA and PS-SVM are helpful for parsing biomedical text, a domain where full parse tree annotation is particularly expensive.

2 Background: DIORA

The Deep Inside-Outside Recursive Autoencoder (DIORA; Drozdov et al., 2019) is an extension of tree recursive neural networks (TreeRNN) that does

not require pre-defined tree structure. It depends on the two primitives $\text{Compose} : \mathbb{R}^{2D} \rightarrow \mathbb{R}^D$ and $\text{Score} : \mathbb{R}^{2D} \rightarrow \mathbb{R}^1$. DIORA is bi-directional — the *inside pass* builds phrase vectors and the *outside pass* builds context vectors. DIORA is trained by predicting words from their context vectors, and has been effective as an unsupervised parser by extracting parse trees from the values computed during the inside pass.

Inside-Outside Typically, a TreeRNN would follow a parse tree to continually compose words or phrases until the entire sentence is represented as a vector, but this requires knowing the tree structure or using some trivial structure such as a balanced binary tree. Instead of using a single structure, DIORA encodes all possible binary trees using a soft-weighting determined by the output of the score function. There are a combinatorial number of valid parse trees for a given sentence — it would be infeasible to encode each of them separately. Instead, DIORA decomposes the problem of representing all valid parse trees by encoding all subtrees over a span into a single phrase vector. For example, each phrase vector is computed in the inside pass according to the following equations:

$$\begin{aligned} h_{i,j}^{in} &= \sum_{k=i+1}^{j-1} s_{i,j,k}^{in} h_{i,j,k}^{in} \\ h_{i,j,k}^{in} &= \text{Compose}(h_{i,k}^{in}, h_{k,j}^{in}) \\ s_{i,j,k}^{in} &= \text{Score}(h_{i,k}^{in}, h_{k,j}^{in}) + s_{i,k}^{in} + s_{k,j}^{in} \end{aligned}$$

The outside pass is computed in a similar way:

$$\begin{aligned} h_{i,j}^{out} &= \sum_{k=i+1}^{j-1} s_{i,j,k}^{out} h_{i,j,k}^{out} \\ h_{j,k}^{out} &= \mathbb{1}_{j < k} \cdot \text{Compose}(h_{j,k}^{in}, h_{i,k}^{out}) + \\ &\quad \mathbb{1}_{j < k} \cdot \text{Compose}(h_{k,i}^{in}, h_{k,j}^{out}) \\ s_{i,j,k}^{out} &= \mathbb{1}_{j < k} \cdot [\text{Score}(h_{j,k}^{in}, h_{i,k}^{out}) + s_{j,k}^{in} + s_{i,k}^{out}] \\ &\quad \mathbb{1}_{j < k} \cdot [\text{Score}(h_{k,i}^{in}, h_{k,j}^{out}) + s_{k,i}^{in} + s_{k,j}^{out}], \end{aligned}$$

where $\mathbb{1}_{j < k}$ is an indicator function that is 1 when the sibling span is on the right, and 0 otherwise (see Figure 2 in Drozdov et al., 2020 for a helpful visualization of the inside and outside pass).

Training DIORA is trained end-to-end directly from raw text and without any parse tree supervision. In our work, we use the same reconstruction objective as in Drozdov et al. (2019). For a sentence x , we optimize the probability of the i -th word x_i using its context (x_{-i}):

$$J_{rec} = -\frac{1}{|x|} \sum_{i=0}^{|x|-1} \log P(x_i|x_{-i}),$$

where $P(\cdot)$ is computed using a softmax layer over a fixed vocab with the outside vector ($h_{i,i}^{out}$) as input.

Parsing DIORA has primarily been used as an unsupervised parser. This requires defining a new primitive $\text{TreeScore} : S(y) = \sum_{i,j,k \in y} s_{i,j,k}^{in}$. A tree y can be extracted from DIORA by solving the search problem that can be done efficiently with the CKY algorithm (Kasami, 1965; Younger, 1967):

$$CKY(x) = \arg \max_y S(y)$$

3 Injecting Span Constraints to DIORA

In this section, we present a method to improve parsing performance by training DIORA such that trees extracted through CKY are more likely to contain known span constraints.

3.1 Test-time injection: Constrained CKY

One option to improve upon CKY is to simply find span constraints and then use a constrained version of CKY (CCKY):

$$CCKY(x, z) = \arg \max_y [S(y) + \epsilon \cdot g(y, z)],$$

where z is a set of known span constraints for x , $g(y, z)$ measures how well the span constraints are satisfied in y , i.e. $g(y, z) = \sum_{i=0}^{|z|-1} \mathbb{1}(z_i \in y)$, and ϵ is an importance weight for the span constraint to guarantee the highest scoring trees are the ones that satisfy the most constraints.¹ Using CCKY rather than CKY typically gives a small boost to parsing performance, but has several downsides described in the remainder of this subsection.

¹To save space, we exclude ϵ hereafter.

Can overfit to constraints DIORA learns to assign weights to the trees that are most helpful for word prediction. For this reason, it is logical to use the weights to find the highest scoring tree. With CCKY, we can find the highest scoring tree that also satisfies the constraints, but this tree could be very different from the original output. Ideally, we would like a method that can incorporate span constraints in a productive way that is not detrimental to the rest of the structure.

Only benefits sentences with constraints If we are dependent on constraints for CCKY, then only sentences that have said constraints will receive any benefit. Ideally, we would like an approach where even sentences without constraints could receive some improvement.

Constraints are required at test time If we are dependent on constraints for CCKY, then we need to find constraints for every sentence at test time. Ideally, we would like an approach where constraints are only needed at the time of training.

Noisy constraints Occasionally a constraint disagrees with a comparable constituency parse tree. In these cases, we would like to have an approach where the model can choose to include only the most beneficial constraints.

3.2 Distant Supervision: Partially Structured SVM

To address the weaknesses of CCKY we present a new training method for DIORA called Partially Structured SVM (PS-SVM).² This is a training objective that can incorporate constraints during training to improve parsing and addresses the aforementioned weaknesses of constrained CKY. PS-SVM follows these steps:

1. Find a *negative* tree (y^-), such as the highest scoring tree predicted by the model:
 $y^- \leftarrow CKY(x)$.
2. Find a *positive* tree (y^+), such as the highest scoring tree that satisfies known constraints:
 $y^+ \leftarrow CCKY(x, z)$.
3. Use the structured SVM with fixed margin to learn to include constraints in the output:
 $J_{PS} = \alpha \cdot \max(0, 1 + S(y^-) - S(y^+))$.

²PS-SVM can be loosely thought of as an application-specific instantiation of Structural SVM with Latent Variables (Yu and Joachims, 2009).

Loss	α	y^-	y^+
NCBL	1	$\arg \max_y S(y)$	$\arg \max_y [S(y) + g(y, z)]$
MIN DIFFERENCE	1	$\arg \max_y S(y)$	$\arg \max_y [S(y) + g(y, z) + g(y, y^-)]$
RESCALE	$g(y^+, y^-)$	$\arg \max_y S(y)$	$\arg \max_y [S(y) + g(y, z)]$
STRUCTURED RAMP	1	$\arg \max_y [S(y) - g(y, z)]$	$\arg \max_y [S(y) + g(y, z)]$

Table 1: Multiple variants of the Partially Structured SVM (PS-SVM) loss, $J_{PS} = \alpha \cdot \max(0, 1 + S(y^-) - S(y^+))$, where z denotes constraint spans and $g(y, z) = \sum_{i=0}^{|z|-1} \mathbb{1}(z_i \in y)$.

3.3 Variants of Partially Structured SVM

The most straightforward application of PS-SVM assigns y^+ to be the highest scoring tree that also incorporates known constraints, and we call this **NAIVE CONSTRAINT-BASED LEARNING (NCBL)**. The shortcoming of NCBL are similar to CCKY, y^+ may be drastically different from the initial prediction y^- and the model may *overfit* to the constraints. With this in mind, an alternative to NCBL is to find y^+ that is high scoring, satisfies the constraints, and has the minimal number of differences with respect to y^- . We refer to this approach as **MIN DIFFERENCE**.

The MIN DIFFERENCE approach gives substantial weight to the initial prediction y^- , which may be helpful for avoiding overfitting to the constraints, but simultaneously is very restrictive on the region of positive trees. In other constraint-based objectives for structured prediction, such as gradient-based inference (Lee et al., 2019), the agreement with constraints is incorporated as a scaling penalty to the gradient step size rather than explicitly restricting the search space of positive examples. Inspired by this, we define another alternative to NCBL called **RESCALE** that scales the step size based on the difference between y^+ and y^- . If the structures are very different, then only use a small step size in order to both prevent overfitting to the constraints and allow for sufficient exploration.

For margin-based learning, for stable optimization a technique known as loss-augmented inference assigns y^- to be the highest scoring and most offending example with respect to the ground truth. When a full structure is not available to assign y^+ , then an alternative option is to use the highest scoring prediction that satisfies the provided partial structure. This approach is called **STRUCTURED RAMP** loss (Chapelle et al., 2009; Gimpel and Smith, 2012; Shi et al., 2021).

In Table 1 we define the 4 variants of PS-SVM. Variants that do not use loss-augmented inference have gradient 0 when y^- contains all constraints.

4 Experimental Setup

In this section, we provide details on data pre-processing, running experiments, and evaluating model predictions. In addition, code to reproduce our experiments and the model checkpoints are available on Github.³

4.1 Training Data and Pre-processing

We train our system in various settings to verify the effectiveness of PS-SVM with span constraints. In all cases, we require access to a text corpus with span constraints.⁴

Ontonotes (CoNLL 2012; Pradhan et al. 2012) consists of ground truth named entity and constituency parse tree labels. In our main experiment (see Table 2), we use the 57,757 ground truth entities from training data as span constraints.

WSJ Penn Treebank (Marcus et al., 1993) consists of ground truth constituency parse tree labels. It is an often-used benchmark for both supervised and unsupervised constituency parsing in English. We also derive synthetic constraints using the ground truth constituents from this data.

MedMentions (Mohan and Li, 2019) is a collection of Pubmed abstracts that have been annotated with UMLS concepts. This is helpful as training data for the biomedical domain. For training we only use the raw text to assist with domain adaptation. We tokenize the text using scispacy.

The Colorado Richly Annotated Full Text (CRAFT) (Cohen et al., 2017) consists of biomedical journal articles that have been annotated with both entity and constituency parse labels. We use CRAFT both for training (with 18,448 entity spans) and evaluation of our model’s performance in the biomedical domain. We sample 3k sentences of training data to use for validation.

³<https://github.com/iesl/distantly-supervised-diora>

⁴Appendix A.1 provides further details about constraints.

4.1.1 Automatically extracted constraints

We experiment with two settings where span constraints are automatically extracted from the training corpus using dictionary lookup in a lexicon. These settings simulate a real world setting where full parse tree annotation is not available, but partial span constraints are readily available.

PMI Constraints We use the phrases defined in the vocab from Mikolov et al. (2013) as a lexicon, treating exact matches found in Ontonotes as constraints. The phrases are learned through word statistics by applying pointwise mutual information (PMI) to find relevant bi-grams, then replacing these bi-grams with a new special token representing the phrase — applied multiple times this technique is used to find arbitrarily long phrases.

Gazetteer We use a list of 1.5 million entity names automatically extracted from Wikipedia (Ratinov and Roth, 2009), which has been effective for supervised entity-centric tasks with both log-linear and neural models (Liu et al., 2019a). We derive constraints by finding exact matches in the Ontonotes corpus that are in the gazetteer. A lexicon containing entity names is often called a gazetteer.

4.2 Training Details

In all cases, we initialize our model’s parameters from pre-trained DIORA (Drozdov et al., 2019). We then continue training using a combination of the reconstruction and PS-SVM loss. Given sentence x and constraints z , the instance loss is:

$$J(x, z) = J_{rec}(x) + J_{PS}(x, z)$$

For the newswire domain, we train for a maximum of 40 epochs on Ontonotes using 6 random seeds and use grid search, taking the best model in each setting according to parsing F1 on the PTB validation set. For biomedical text, since it is a shift in domain from the DIORA pre-training, we first train for 20 epochs using a concatenation of MedMentions and CRAFT data with only the reconstruction loss⁵ (called DIORA_{ft} for “fine-tune”). Then, we train for 40 epochs like previously mentioned, using performance on a subset of 3k random sentences from the CRAFT training data for early stopping. Hyperparameters are in Appendix A.2.

⁵The training jointly with MedMentions and CRAFT is a special case of “intermediate fine-tuning” (Phang et al., 2018).

	F1
General Purpose	
Ordered Neuron† (Shen et al., 2019)	48.1 ±1.0
Compound PCFG† (Kim et al., 2019a)	55.2 ±2.5
DIORA‡ (Drozdov et al., 2019)	56.8
S-DIORA† (Drozdov et al., 2020)	57.6 ±3.2
Constituency Tests	
RoBERTa† (Cao et al., 2020)	62.8 ±1.6
DIORA Span Constraints	
+CCKY	57.5
+PS-SVM _{NCBL}	60.4 ±0.1
+PS-SVM _{MINDIFF}	59.0 ±0.8
+PS-SVM _{RESCALE}	61.2 ±0.6
+PS-SVM _{STRUCTURE RAMP}	59.9 ±1.0

Table 2: Parsing F1 on PTB. The average F1 across random seeds is measured on the test set, and the standard deviation is shown as subscript when applicable. †: Indicates that standard deviation is the approximate lower bound derived from the mean, max, and number of random seeds. ‡: Indicates no average performance available, so the max is reported.

4.3 Evaluation

In all cases, we report Parsing F1 aggregated at the sentence level — F1 is computed separately for each sentence then averaged across the dataset. To be consistent with prior work, punctuation is removed prior to evaluation⁶ and F1 is computed using the eval script provided by Shen et al. (2018).^{7,8} In tables 2, 3, and 4 we average performance across random seeds and report the standard deviation.

Baselines In Table 2, we compare parsing F1 with four *general purpose* unsupervised parsing models that are trained directly from raw text. We also compare with Cao et al. (2020) that uses a small amount of supervision to generate constituency tests used for training — their model has substantially more parameters than our other baselines and is based on RoBERTa (Liu et al., 2019b).

⁶In general, it is less important that subtrees associated with punctuation match the Penn Treebank guideline (Bies et al., 1995) than if the model makes consistent decision with respect to these cases. For this reason, omitting punctuation for evaluation gives a more reliable judgement when parsing is unsupervised.

⁷This script ignores trivial spans, and we use the version provided in <https://github.com/harvardnlp/compound-pcfg>.

⁸We were not able to reproduce the results from the concurrent work Shi et al. (2021), which does not share their parse tree output and uses a slightly different evaluation.

5 Results and Discussion

In our experiments and analysis we aim to address several research questions about incorporating span constraints for the task of unsupervised parsing.

5.1 Is Constrained CKY sufficient?

A natural idea is to constrain the output of DIORA to contain any span constraints (§3.1). We expect this type of *hard constraint* to be ineffective for various reasons: 1) The model is not trained to include constraints, so any predictions that forces their inclusion are inherently noisy; 2) Similar to (1), some constraints are not informative and may be in disagreement with the desired downstream task and the model’s reconstruction loss; and 3) Constraints are required at test time and only sentences with constraints can benefit.

We address these weaknesses by training our model to include the span constraints in its output using PS-SVM. This can be considered a *soft* way to include the constraints, but has other benefits including the following: 1) The model implicitly learns to ignore constraints that are not useful; 2) Constraints are not necessary at test time; and 3) The model improves performance even on sentences that did not have constraints.

The effectiveness of our approach is visible in Table 2 where we use ground truth entity boundaries as constraints. CCKY slightly improves upon DIORA, but our PS-SVM approach has a more substantial impact. We experiment with four variants of PS-SVM (described in §3.3) — RESCALE is most effective, and throughout this text this is the variant of PS-SVM used unless otherwise specified.

5.2 Real world example with low effort constraint collection

Our previous experiments indicate that span constraints are an effective way to improve unsupervised parsing. How can we leverage this method to improve unsupervised parsing in a real world setting? We explore two methods for easily finding span constraints (see Table 3).

We find that PMI is effective as a lexicon, but not as much as the gazetteer. PMI provides more constraints than the gazetteer, but the constraints disagree more frequently with the ground truth structure and a smaller percentage of spans align exactly with the ground truth. The gazetteer approach is better than using CCKY with ground truth entity spans, despite using less than half as many con-

	WSJ	Constraints					
	F1	EM	C	n^z	R_{train}^{pre}	R_{train}^{post}	R_{test}^{post}
DIORA	56.8	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset	\emptyset
+Entity	61.9	96.3	1.9	58,075	79.3	98.9	96.4
+PMI	57.8	43.9	7.4	31,965	75.3	94.4	90.0
+Gazetteer	58.8	51.3	5.0	22,354	80.2	97.0	93.4

Table 3: Parsing F1 on PTB. The max F1 across random seeds is measured on the test set. The corresponding span recall is shown on the Ontonotes $train$ and $test$ data before (R^{pre}) and after (R^{post}) training. The first row shows DIORA performance. Following rows show performance using distant supervision. EM: Exact Match (percent of span constraints that are also constituents); C: Crossing (percent of span constraints that cross a constituent); n^z : Number of span constraints. The constraint-based metrics are not applicable to DIORA and marked with \emptyset .

straints that only align exactly with the ground truth nearly half the time. We use gazetteer in only the most naive way via exact string matching, so we suspect that a more sophisticated yet still high precision approach (e.g. approximate string match) would have more hits and provide more benefit.

For both PMI and Gazetteer, we found that NCBL gave the best performance.

5.3 Impact on consistent convergence

We find that using constraints with PS-SVM considerably decreases the variance on performance compared with previous baselines.⁹ This is not surprising given that latent tree learning (i.e. unsupervised parsing) can converge to many equally viable parsing strategies. By using constraints, we are guiding optimization to converge to a point more aligned with the desired downstream task.

5.4 Are entity spans sufficient as constraints?

Given that DIORA already captures a large percentage of span constraints represented by entities, it is somewhat surprising that including them gives any F1 improvement. That being said, it is difficult to know *a priori* which span constraints are most beneficial and how much improvement to expect. To help understand the benefits of different types of span constraints, we derived synthetic constraints using the most frequent constituent types

⁹Although, most previous work does not explicitly report the standard deviation (STDEV), we can use the mean, max, and number of trials to compute the lower bound on STDEV. This yields 2.5 (Compound PCFG), 3.2 (S-DIORA), and 1.6 (RoBERTa). In contrast, our best setting has STDEV 0.6.

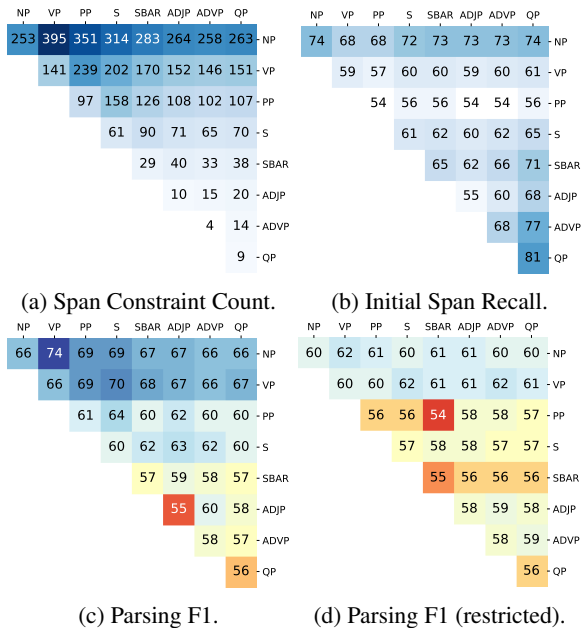


Figure 2: Various statistics when using 1 or 2 constituent types as span constraints on the WSJ training and validation data. (a): The count of each span constraint in the training data (in thousands). (b): The percent of span constraints captured (i.e. span recall) in the validation data. (c): Parsing F1 on the validation data when using the span constraints with PS-SVM. (d): Parsing F1 on the validation data when using PS-SVM, although span constraints have been restricted to match the frequency and nesting behavior of entities.

from ground truth parse trees in Ontonotes (see Figure 2). The constraints extracted this way look very different from the entity constraints in that they often are nested and in general are much more frequent. To make a more fair comparison we prevent nesting and downsample to match the frequency of the entity constraints (see Figure 2d).

From these experiments, we can see NP or VP combined with other constraints usually lead to the best parsing performance (Figure 2c). This is the case even if DIORA had relatively low span recall on a different constraint type (Figure 2b). A reasonable hypothesis is that simply having more constraints leads to better performance, which mirrors the result that the settings with the most constraints perform better overall (Figure 2a). When filtered to match the shape and frequency of entity constraints, we see that performance based on NP constraints is nearly the same as with entities (Figure 2d). This suggests that entity spans are effective as constraints with respect to other types of constraints, but that in general we should aim to gather as many constraints as possible.

	CRAFT	Constraints	
	F1	R_{train}	R_{test}
UB	85.4	82.8	79.2
DIORA	50.7	47.4	44.8
DIORA _{ft}	55.8	72.4	65.9
+CCKY	56.2	99.0	98.6
+PS-SVM	56.8	91.1	85.3

Table 4: Parsing F1 and Span Recall on CRAFT. The max F1 across random seeds is measured on the test set. DIORA_{ft}: Fine-tuned on word prediction to assist domain transfer. UB: The upper bound on performance measured by binarizing the ground truth tree.

5.5 Case Study: Parsing Biomedical Text

The most impactful domain for our method would be unsupervised parsing in a domain where full constituency tree annotation is very expensive, and span constraints are relatively easy to acquire. For this reason, we run experiments using the CRAFT corpus (Verspoor et al., 2011), which contains text from biomedical research. The results are summarized in Tables 4 and 5.

5.5.1 Domain Adaptation: Fine-tuning through Word Prediction

Although CRAFT and PTB are both in English, the text in biomedical research is considerably different compared with text in the newswire domain. When we evaluate the pre-trained DIORA model on the CRAFT test set, we find it achieves 50.7 F1. By simply fine-tuning the DIORA model on biomedical research text using only the word-prediction objective (J_{rec}) we can improve this performance to 55.8 F1 (+5.1 F1; DIORA_{ft} in Table 4). This observation accentuates a beneficial property about unsupervised parsing models like DIORA: for domain adaptation, simply continue training on data from the target domain, which is possible because the word-prediction objective does not require label collection, unlike supervised models.

5.5.2 Incorporating Span Constraints

We use the ground truth entity annotation in the CRAFT training data as a source of distant supervision and continue training DIORA using the PS-SVM objective. By incorporating span constraints this way, we see that parsing performance on the test set improves from 55.8 \rightarrow 56.8 (+1 F1).

For CRAFT, we used grid search over a small set of hyperparameters including loss variants and found that STRUCTURED RAMP performed best.

			DIORA		+CCKY		+PS-SVM	
	n	n^z	F1	R ^z	F1	R ^z	F1	R ^z
CAPTION	1857	1579	55.7	67.6	56.0	98.3	56.0	86.1
HEADING	1149	201	72.0	59.2	72.8	96.5	73.5	83.1
TITLE	29	31	51.4	58.1	53.8	96.8	55.4	71.0
CIT	3	0	40.0	∅	40.0	∅	40.0	∅
S-IMP	1	0	36.8	∅	36.8	∅	31.6	∅
S	5872	5140	53.9	65.9	54.2	98.8	54.9	85.4
NP	136	34	37.1	41.2	40.6	100.0	44.1	52.9
FRAG	39	52	49.3	71.2	49.0	100.0	51.8	84.6
SINV	6	7	50.7	42.9	47.9	85.7	46.3	57.1
SBARQ	5	1	49.5	100.0	49.5	100.0	55.1	100.0
SQ	2	1	28.0	100.0	28.0	100.0	32.9	100.0

Table 5: Parsing F1 on CRAFT test set from the best model bucketed by the sentence’s top-most constituent type. n : Count of sentences. n^z : Count of constraints. R^z: Recall on constraints. ∅: Indicates no constraints.

Performance by Sentence Type In Table 5 we report parsing results bucketed by sentence-type determined by the top-most constituent label. In general, across almost all sentence types, simply constraining the DIORA output to incorporate known spans boosts F1 performance. Training with the PS-SVM objective usually improves F1 further, although the amount depends on the sentence type.

Challenging NP-type Sentences We observe especially low span-recall for sentences with NP as the top-most constituent (Table 5). These are short sentences that exhibit domain-specific structure. Here is a typical sentence and ground truth parse for that case:

((HIF - 1α) KO) - ((skeletal - muscle)
(HIF - 1α) knockout mouse)

Various properties of the above sentence make it difficult to parse. For instance, the sentence construction lacks syntactic cues and there is no verb in the sentence. There is also substantial ambiguity with respect to hyphenation, and the second hyphen is acting as a colon. These properties make it difficult to capture the spans (*skeletal - muscle*) or the second (*HIF - 1α*) despite being constraints.

5.5.3 Parsing of PTB vs. CRAFT

As mentioned in §5.5.1, there is considerable difference in the text between PTB and CRAFT. It follows that there would be a difference in difficulty when parsing these two types of data. After running the parser from Kitaev and Klein (2018) on each dataset, it appears CRAFT is more difficult to parse than PTB. For CRAFT, the unlabeled parsing F1 is 81.3 and the span recall for entities is 37.6. For PTB, the unlabeled parsing F1 is 95.

6 Related Work

Learning from Partially Labeled Corpora Pereira and Schabes (1992) modify the inside-outside algorithm to respect span constraints. Similar methods have been explored for training CRFs (Culotta and McCallum, 2004; Bellare and McCallum, 2007). Rather than modify the weight assignment in DIORA, which is inspired by the inside-outside algorithm, we supervise the tree predicted from the inside-pass.

Concurrent work to ours in distant supervision trains RoBERTa for constituency parsing using answer spans from question-answering datasets and wikipedia hyperlinks (Shi et al., 2021). Although effective, their approach depends entirely on the set of constraints. In contrast, PS-SVM enhances DIORA, which is a model that outputs a parse tree without any supervision.

The span constraints in this work are derived from external resources, and do not necessarily match the parse tree. Constraints may conflict with the parse, which is why CCKY can be less than 100 span recall in Table 4. This approach to model training is often called “distant supervision” (Mintz et al., 2009; Shi et al., 2021). In contrast, “partial supervision” implies gold partial labels are available, which we explore as synthetic data (§5.4), but in general do not make this assumption.

Joint Supervision An implicit way to incorporate constraints is through multi-task learning (MTL; Caruana, 1997). Even when relations between the tasks are not modeled explicitly, MTL has shown promise throughout a range of text processing tasks with neural models (Collobert and Weston, 2008; Swayamdipta et al., 2018; Kuncoro et al., 2020). Preliminary experiments with joint NER did not improve parsing results. This is in-line with DIORA’s relative weakness in representing fine-grained entity types. Modifications of DIORA to improve its semantic representation may prove to make joint NER more viable.

Constraint Injection Methods There exists a rich literature in constraint injection (Ganchev et al., 2010; Chang et al., 2012). Both methods are based on Expectation Maximization (EM) algorithm (Dempster et al., 1977) where the constraint is injected in the E-step of calculating the posterior distribution (Samdani et al., 2012). Another line of work focuses injecting constraint in the M-step (Lee et al., 2019; Mehta et al., 2018) by reflecting

the degree of constraint satisfaction of prediction as the weight of the gradient. Our approach is similar to [Chang et al. \(2012\)](#) as we select the highest scoring output that satisfies constraints and learn from it. PS-SVM_{RESCALE} is based on [Lee et al. \(2019\)](#).

The aforementioned constraint injection methods were usually used as an added loss to the supervised loss function. In this work, we show that the distant supervision through constraint injection is beneficial for unsupervised setting as well.

Structural SVM with Latent Variables The PS-SVM loss we introduce in this work can be loosely thought of as an application-specific instantiation of Structural SVM with Latent Variables ([Yu and Joachims, 2009](#)). Various works have extended Structural SVM with Latent Variables to incorporate constraints for tasks such as sequence labeling ([Yu, 2012](#)) and co-reference resolution ([Chang et al., 2013](#)), although none we have seen focus on unsupervised constituency parsing. Perhaps a more clear distinction is that [Yu and Joachims \(2009\)](#) focuses on latent variables within supervised tasks, and PS-SVM is meant to improve convergence of an unsupervised learning algorithm (i.e., DIORA).

Additional Related Work In Appendix A.3 we list additional work in unsupervised parsing not already mentioned.

7 Conclusion

In this work, we present a method for enhancing DIORA with distant supervision from span constraints. We call this approach Partially Structured SVM (PS-SVM). We find that span constraints based on entities are effective at improving parsing performance of DIORA on English newswire data (+5.1 F1 using ground truth entities, or +2 F1 using a gazetteer). Furthermore, we show PS-SVM is also effective in the domain of biomedical text (+1 F1 using ground truth entities). Our detailed analysis shows that entities are effective as span constraints, giving equivalent benefit as a similar amount of NP-based constraints. We hope our findings will help “bridge the gap” between supervised and unsupervised parsing.

Broader Impact

We hope our work will increase the availability of parse tree annotation for low-resource domains, generated in an unsupervised manner. Compared with full parse tree annotation, span constraints can

be acquired at reduced cost or even automatically extracted.

The gazetteer used in our experiments is automatically extracted from Wikipedia, and our experiments are only for English, which is the language with by far the most Wikipedia entries. Although, similarly sized gazetteers may be difficult to attain in other languages, [Mikheev et al. \(1999\)](#) point out larger gazetteers do not necessarily boost performance, and gazetteers have already proven effective in low-resource domains ([Rijhwani et al., 2020](#)). In any case, we use gazetteers in the most naive way by finding exact text matches. When extending our approach to other languages, an entity recognition model may be a suitable replacement for the gazetteer.

Acknowledgements

We are grateful to our colleagues at UMass NLP and the anonymous reviewers for feedback on drafts of this work. This work was supported in part by the Center for Intelligent Information Retrieval, in part by the Chan Zuckerberg Initiative, in part by the IBM Research AI through the AI Horizons Network, and in part by the National Science Foundation (NSF) grant numbers DMR-1534431, IIS-1514053, CNS-0958392, and IIS-1955567. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

References

- Nader Akoury, Kalpesh Krishna, and Mohit Iyyer. 2019. [Syntactically supervised transformers for faster neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1269–1281, Florence, Italy. Association for Computational Linguistics.
- Kedar Bellare and Andrew McCallum. 2007. Learning extractors from unlabeled text using relevant databases. In *Proceedings of the 2007 AAAI Workshop on information integration on the web*.
- Ann Bies, Mark Ferguson, Karen Katz, Robert MacIntyre, Victoria Tredinnick, Grace Kim, Mary Ann Marcinkiewicz, and Britta Schasberger. 1995. Bracketing guidelines for Treebank II style Penn Treebank project. Technical report, Department of Linguistics, University of Pennsylvania.
- Eric Brill, David Magerman, Mitchell Marcus, and Beatrice Santorini. 1990. [Deducing linguistic structure from the statistics of large corpora](#). In *Speech*

- and Natural Language: *Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.
- Steven Cao, Nikita Kitaev, and Dan Klein. 2020. [Unsupervised parsing via constituency tests](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4798–4808, Online. Association for Computational Linguistics.
- Glenn Carroll and Eugene Charniak. 1992. Two experiments on learning probabilistic dependency grammars from corpora. Technical report, Dept. of Computer Science, Brown University.
- Rich Caruana. 1997. Multitask learning. *Machine learning*, 28(1):41–75.
- Kai-Wei Chang, Rajhans Samdani, and Dan Roth. 2013. [A constrained latent variable model for coreference resolution](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 601–612, Seattle, Washington, USA. Association for Computational Linguistics.
- Ming-Wei Chang, Lev Ratinov, and Dan Roth. 2012. Structured learning with constrained conditional models. *Machine learning*, 88(3):399–431.
- Olivier Chapelle, Chuong B., Choon Teo, Quoc Le, and Alex Smola. 2009. [Tighter bounds for structured estimation](#). In *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc.
- Alexander Clark. 2001. Unsupervised induction of stochastic context-free grammars using distributional clustering. In *CoNLL*.
- Kevin Bretonnel Cohen, Karin Verspoor, Karën Fort, Christopher Funk, Michael Bada, Martha Palmer, and Lawrence Hunter. 2017. The colorado richly annotated full text (CRAFT) corpus: Multi-model annotation in the biomedical domain. In *Handbook of Linguistic Annotation*, page 1379 – 1394. Springer.
- Ronan Collobert and J. Weston. 2008. A unified architecture for natural language processing: deep neural networks with multitask learning. In *ICML '08*.
- Aron Culotta and Andrew McCallum. 2004. [Confidence estimation for information extraction](#). In *Proceedings of HLT-NAACL 2004: Short Papers*, pages 109–112, Boston, Massachusetts, USA. Association for Computational Linguistics.
- A. Dempster, N. Laird, and D. Rubin. 1977. Maximum likelihood from incomplete data via the EM - algorithm plus discussions on the paper. *Journal of the Royal Statistical Society: Series B (Methodological)*.
- Andrew Drozdov, Subendhu Rongali, Yi-Pei Chen, Tim O’Gorman, Mohit Iyyer, and Andrew McCallum. 2020. Unsupervised parsing with S-DIORA: Single tree encoding for deep inside-outside recursive autoencoders. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Andrew Drozdov, Patrick Verga, Mohit Yadav, Mohit Iyyer, and Andrew McCallum. 2019. Unsupervised latent tree induction with deep inside-outside recursive autoencoders. In *NAACL-HLT*.
- Jenny Rose Finkel and Christopher D. Manning. 2009. [Joint parsing and named entity recognition](#). In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, May 31 - June 5, 2009, Boulder, Colorado, USA*, pages 326–334. The Association for Computational Linguistics.
- Kuzman Ganchev, João Graça, Jennifer Gillenwater, and B. Taskar. 2010. Posterior regularization for structured latent variable models. *J. Mach. Learn. Res.*, 11:2001–2049.
- Kevin Gimpel and Noah A. Smith. 2012. [Structured ramp loss minimization for machine translation](#). In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 221–231, Montréal, Canada. Association for Computational Linguistics.
- Aria Haghighi and Dan Klein. 2006. Prototype-driven grammar induction. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*.
- Phu Mon Htut, Kyunghyun Cho, and Samuel Bowman. 2018. [Grammar induction with neural language models: An unusual replication](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 371–373, Brussels, Belgium. Association for Computational Linguistics.
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. [Adversarial example generation with syntactically controlled paraphrase networks](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1875–1885, New Orleans, Louisiana. Association for Computational Linguistics.
- Lifeng Jin, Finale Doshi-Velez, Timothy Miller, William Schuler, and Lane Schwartz. 2018. [Unsupervised Grammar Induction with Depth-bounded PCFG](#). *Transactions of the Association for Computational Linguistics*, 6:211–224.
- T. Kasami. 1965. An efficient recognition and syntax analysis algorithm for context-free languages. Technical Report AFCRL-65-758, Air Force Cambridge Research Laboratory, Bedford, MA†.

- Yoon Kim, Chris Dyer, and Alexander M Rush. 2019a. Compound probabilistic context-free grammars for grammar induction. In *ACL*.
- Yoon Kim, Alexander Rush, Lei Yu, Adhiguna Kuncoro, Chris Dyer, and Gábor Melis. 2019b. [Unsupervised recurrent neural network grammars](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1105–1117, Minneapolis, Minnesota. Association for Computational Linguistics.
- Nikita Kitaev and Dan Klein. 2018. Constituency parsing with a self-attentive encoder. In *Association for Computational Linguistic (ACL)*.
- 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 (ACL-04)*, pages 478–485, Barcelona, Spain.
- Dan Klein and Christopher D. Manning. 2001. Natural language grammar induction using a constituent-context model. In *NeurIPS*.
- Adhiguna Kuncoro, Lingpeng Kong, Daniel Fried, Dani Yogatama, Laura Rimell, Chris Dyer, and Phil Blunsom. 2020. Syntactic structure distillation pre-training for bidirectional encoders. *Transactions of the Association for Computational Linguistics*, 8:776–794.
- Karim Lari and Steve J Young. 1990. The estimation of stochastic context-free grammars using the inside-outside algorithm. *Computer speech & language*, 4(1):35–56.
- Jay Yoon Lee, Sanket Vaibhav Mehta, Michael L. Wick, Jean-Baptiste Tristan, and Jaime G. Carbonell. 2019. [Gradient-based inference for networks with output constraints](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 4147–4154. AAAI Press.
- Tianyu Liu, Jin-Ge Yao, and Chin-Yew Lin. 2019a. [Towards improving neural named entity recognition with gazetteers](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5301–5307, Florence, Italy. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Michael Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692.
- Mitchell Marcus, Beatrice Santorini, and Mary Ann Marcinkiewicz. 1993. Building a large annotated corpus of english: The penn treebank. *Computational linguistics*, 19(2):313–330.
- Stephen Mayhew, Snigdha Chaturvedi, Chen-Tse Tsai, and Dan Roth. 2019. [Named entity recognition with partially annotated training data](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 645–655, Hong Kong, China. Association for Computational Linguistics.
- Sanket Vaibhav Mehta, Jay Yoon Lee, and Jaime Carbonell. 2018. [Towards semi-supervised learning for deep semantic role labeling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4958–4963, Brussels, Belgium. Association for Computational Linguistics.
- Andrei Mikheev, Marc Moens, and Claire Grover. 1999. [Named entity recognition without gazetteers](#). In *Ninth Conference of the European Chapter of the Association for Computational Linguistics*, pages 1–8, Bergen, Norway. Association for Computational Linguistics.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NeurIPS*.
- Mike Mintz, Steven Bills, Rion Snow, and Daniel Jurafsky. 2009. [Distant supervision for relation extraction without labeled data](#). 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 1003–1011, Suntec, Singapore. Association for Computational Linguistics.
- Sunil Mohan and Donghui Li. 2019. Medmentions: A large biomedical corpus annotated with UMLS concepts. In *Automated Knowledge Base Construction (AKBC)*.
- Robert Moore, Douglas Appelt, John Dowding, J. Mark Gawron, and Douglas Moran. 1995. Combining linguistic and statistical knowledge sources in natural-language processing for atis. In *Proceedings of the January 1995 ARPA Spoken Language Systems Technology Workshop*.
- Jason Naradowsky. 2014. *Learning with Joint Inference and Latent Linguistic Structure in Graphical Models*. Ph.D. thesis, University of Massachusetts Amherst.
- Vlad Niculae and Andre Martins. 2020. [LP-SparseMAP: Differentiable relaxed optimization for sparse structured prediction](#). In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 7348–7359. PMLR.

- Fernando Pereira and Yves Schabes. 1992. [Inside-outside reestimation from partially bracketed corpora](#). In *30th Annual Meeting of the Association for Computational Linguistics*, pages 128–135, Newark, Delaware, USA. Association for Computational Linguistics.
- Jason Phang, Thibault Févry, and Samuel R. Bowman. 2018. Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks. *ArXiv*, abs/1811.01088.
- Elias Ponvert, Jason Baldridge, and Katrin Erk. 2011. [Simple unsupervised grammar induction from raw text with cascaded finite state models](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 1077–1086, Portland, Oregon, USA. Association for Computational Linguistics.
- Sameer Pradhan, Alessandro Moschitti, Nianwen Xue, Olga Uryupina, and Yuchen Zhang. 2012. [Conll-2012 shared task: Modeling multilingual unrestricted coreference in ontonotes](#). In *Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning - Proceedings of the Shared Task: Modeling Multilingual Unrestricted Coreference in OntoNotes, EMNLP-CoNLL 2012, July 13, 2012, Jeju Island, Korea*, pages 1–40. ACL.
- Lev Ratinov and Dan Roth. 2009. [Design challenges and misconceptions in named entity recognition](#). In *Proceedings of the Thirteenth Conference on Computational Natural Language Learning (CoNLL-2009)*, pages 147–155, Boulder, Colorado. Association for Computational Linguistics.
- Shruti Rijhwani, Shuyan Zhou, Graham Neubig, and Jaime Carbonell. 2020. [Soft gazetteers for low-resource named entity recognition](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8118–8123, Online. Association for Computational Linguistics.
- Rajhans Samdani, Ming-Wei Chang, and Dan Roth. 2012. Unified expectation maximization. In *NAACL-HLT*.
- Yikang Shen, Zhouhan Lin, Chin-Wei Huang, and Aaron Courville. 2018. Neural language modeling by jointly learning syntax and lexicon. In *ICLR*.
- Yikang Shen, Shawn Tan, Alessandro Sordani, and Aaron Courville. 2019. Ordered neurons: Integrating tree structures into recurrent neural networks. In *International Conference on Learning Representations (ICLR)*.
- Haoyue Shi, Jiayuan Mao, Kevin Gimpel, and Karen Livescu. 2019. Visually grounded neural syntax acquisition. In *Association for Computational Linguistics*.
- Tianze Shi, Ozan İrsoy, Igor Malioutov, and Lillian Lee. 2021. [Learning syntax from naturally-occurring bracketings](#). In *NAACL-HLT*.
- Noah A Smith and Jason Eisner. 2005. Contrastive estimation: Training log-linear models on unlabeled data. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*.
- Benjamin Snyder, Tahira Naseem, and Regina Barzilay. 2009. [Unsupervised multilingual grammar induction](#). 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 73–81, Suntec, Singapore. Association for Computational Linguistics.
- Mitchell Stern, Jacob Andreas, and Dan Klein. 2017. A minimal span-based neural constituency parser. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Swabha Swayamdipta, Sam Thomson, Kenton Lee, Luke Zettlemoyer, Chris Dyer, and Noah A. Smith. 2018. [Syntactic scaffolds for semantic structures](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3772–3782, Brussels, Belgium. Association for Computational Linguistics.
- Ann Taylor, Mitchell Marcus, and Beatrice Santorini. 2003. [The penn treebank: An overview](#). In *Treebanks: Building and Using Parsed Corpora*, pages 5–22, Dordrecht. Springer Netherlands.
- Karin Verspoor, Kevin Cohen, Arrick Lanfranchi, Colin Warner, Helen L. Johnson, Christophe Roeder, Jinho D. Choi, Christopher Funk, Yuriy Malenkiy, Miriam Eckert, Nianwen Xue, William A. Baumgartner Jr., Michael Bada, Martha Palmer, and Lawrence E. Hunter. 2011. A corpus of full-text journal articles is a robust evaluation tool for revealing differences in performance of biomedical natural language processing tools. *BMC Bioinformatics*, 13:207 – 207.
- Adina Williams, Andrew Drozdov, and Samuel R. Bowman. 2018. Do latent tree learning models identify meaningful structure in sentences? *Transactions of the Association for Computational Linguistics*, 6:253–267.
- Dongqin Xu, Junhui Li, Muhua Zhu, Min Zhang, and Guodong Zhou. 2020. [Improving AMR parsing with sequence-to-sequence pre-training](#). In *EMNLP*.
- Daniel H. Younger. 1967. [Recognition and parsing of context-free languages in time n³](#). *Information and Control*, 10(2):189–208.
- Chun-Nam Yu. 2012. [Transductive learning of structural SVMs via prior knowledge constraints](#). In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, volume 22 of *Proceedings of Machine Learning Research*, pages 1367–1376, La Palma, Canary Islands. PMLR.
- Chun-Nam John Yu and Thorsten Joachims. 2009. [Learning structural SVMs with latent variables](#). In *ICML*, pages 1169–1176.

A Appendix

A.1 Constraint Statistics

Here we report a detailed breakdown of span constraints and the associated constituent types. Compared with Shi et al. (2021), span constraints based on entities are less diverse with respect to constituent type. In future work, we plan to use their data combined with DIORA and PS-SVM training. Also, we hypothesize that RoBERTa would be effective as a data augmentation to easily find new constraints.

	Ontonotes			CRAFT
	NER	Gazetter	PMI	NER
Exact match	96.3	51.3	43.9	57.4
Conflict	1.9	5.0	7.4	12.0
NP	9.2	1.7	1.9	4.0
VP	0.0	0.0	0.0	0.0
S	0.1	0.0	0.0	0.0
ADVP	7.5	0.0	1.5	0.2
ADJP	3.1	0.8	2.2	3.3
SBAR	0.0	0.0	0.0	0.0
NML	21.6	11.6	14.9	17.9
QP	46.6	0.0	0.0	0.0
PP	0.1	0.0	0.0	0.0
Total	3.1	1.2	1.7	4.0
Number of sentences	115,811			18,951
Number of ground truth spans	1,878,737			361,394
Span/sentences	0.50	0.19	0.28	0.77

Table 6: Statistics of different type constraints in Ontonotes. The top part shows how each constraint type agree with the ground truth parsing. The middle shows the percentages of each constituency spans found in constraint spans. The bottom part shows the total number of sentences and constraint spans per sentence.

A.2 Hyperparameters

We run a small grid search with multiple random seeds. The following search parameters are fixed for all experiments.

Model Dimension:	400
Optimization Algorithm:	Adam
Hardware:	1x1080ti
Training Time:	$O(24h)$

Also, we search over the 4 variants of PS-SVM (§3.3) when incorporating constraints. We mention the best performing variant of PS-SVM where it is relevant. The best performing setting for each hyperparameter is underlined.

A.2.1 Newswire

For newswire experiments, we train with Ontonotes and validate with PTB.

Learning Rate:	$\underline{2^{-3}}, 1^{-3}$
Max Training Length:	40
Batch Size:	32
Max Epochs:	40
Stopping Criteria:	Validation F1
No. of Random Seeds:	6

Using RESCALE gave the best result with ground truth entity-based constraints, and NCBL gave the best result for PMI and gazetteer-based constraints.

A.2.2 Biomedical Text

First, to assist with domain adaptation, we train using a concatenation of CRAFT and MedMentions (DIORA_{ft}). We sample 3k sentences from CRAFT training data to use for validation.

Learning Rate:	2^{-3}
Max Training Length:	30
Batch Size:	32
Max Epochs:	20
Stopping Criteria:	Validation F1
No. of Random Seeds:	1

Then we incorporate constraints and train only with CRAFT, using the same sample for validation.

Learning Rate:	$2^{-3}, \underline{1^{-3}}, 5^{-4}, 1^{-4}$
Max Training Length:	40
Batch Size:	4, <u>8</u> , 32
Max Epochs:	40
Stopping Criteria:	Validation F1
No. of Random Seeds:	3

Using STRUCTURED RAMP gave the best result.

A.2.3 Other Details

We report validation and test performance where applicable. All of our model output are shared in our github repo for further analysis. Training with PS-SVM uses the same parameters as standard DIORA training — the supervision is directly on the scores computed for the inside-pass and does not require any new parameters.

A.2.4 Use of Validation Data

Shi et al. (2019) point out that validation sets can disproportionately skew performance of unsupervised parsing systems. We re-did early stopping using 100 random sentences and found that the best model remained the same in all cases. This is consistent with the DIORA-related experiments in Shi et al. (2019) that show DIORA performance is robust when only a small number of samples are used for model selection.

A.2.5 Why fine-tune?

To be resource efficient, we use the pre-trained DIORA checkpoint from Drozdov et al. (2019) and fine-tune it for parsing biomedical text. DIORA was trained for 1M gradient updates on nearly 2M sentences from NLI data, taking 3 days using 4x GPUs. MedMentions has ~40k training sentences, CRAFT has only ~40k, and our PS-SVM experiments run in less than 1 day using a single GPU.

A.3 Additional Related Work

In the main text, we mention the most closely related work for training DIORA with our PS-SVM objective. Here we cover other work not discussed. Unsupervised parsing has a long and dense history, and we hope this section provides context to the state of the field, our contribution in this paper, and can serve as a guide for the interested researcher.

History of unsupervised parsing over the last thirty years As early as 1990, researcher were using corpus statistics to induce grammar, not unlike how our span constraints based on PMI are derived (Brill et al., 1990) — at this point the Penn Treebank was still in progress of being annotated. Other techniques focused on optimizing sentence likelihood with probabilistic context-free grammars, although with limited success (Lari and Young, 1990; Carroll and Charniak, 1992; Pereira and Schabes, 1992). Later work exploited the statistics between phrases and their context (Clark, 2001; Klein and Manning, 2001), but the most promising practical progress in this line of work was not seen until 15+ years later.

In the mid 2010s, many papers were published about neural models for language that claimed to induce tree-like structure, albeit none made strong claims about unsupervised parsing. Williams et al. (2018) analyzed these models and discovered a negative result. Despite their tree-structured inductive bias, when measured against ground truth

parse trees from the Penn Treebank these models did only slightly better than random and were not competitive with earlier work grammar induction. Shortly after, Shen et al. (2018) developed a neural language model with a tree-structured attention pattern and Htut et al. (2018) demonstrated its effectiveness at unsupervised parsing, the first positive result for a neural model. In quick succession, more papers were published with improve results and new neural architectures (Shen et al., 2019; Drozdov et al., 2019; Kim et al., 2019a,b; Cao et al., 2020, *inter alia*), some of which we include as baselines in Table 2. Perhaps one of the more interesting work was improved performance of unsupervised parsing with PCFG when parameterized as a neural model (Neural PCFG; Kim et al., 2019a). These results suggest that the modern NLP machinery has made unsupervised parsing more viable, yet it is still not clear which of the newly ubiquitous tools (word vectors, contextual language models, adaptive optimizers, etc.) makes the biggest impact.

Variety of approaches to unsupervised parsing

The majority of the models in the work reported above optimize statistics with respect to the training data (with Cao et al., 2020 as an exception), but many techniques have been explored by now towards the same end. Unsupervised constituency parsing can be done in a variety ways including: exploiting patterns between images and text (Shi et al., 2019), exploiting patterns in parallel text (Snyder et al., 2009), joint induction of dependency and constituency (Klein and Manning, 2004), iterative chunking (Ponvert et al., 2011), contrastive learning (Smith and Eisner, 2005), and more.

Other constraint types In our work we focus on span constraints, especially those based on entities or automatically derived from a lexicon, and encourage those spans to be included in the model’s prediction. Prior knowledge of language can be useful in defining other types of structural constraints. For instance, in Mayhew et al. (2019) the distribution of NER-related tokens is helpful to improve performance for low-resource languages. Perhaps more relevant, Jin et al. (2018) present a version of PCFG with bounded recursion depth. Niculae and Martins (2020) present a flexible optimization framework for incorporating structural constraints such as bounded recursion depth and demonstrate strong results on synthetic data.