

SEMI-SUPERVISED SEMANTIC DEPENDENCY PARSING USING CRF AUTOENCODERS

Anonymous authors

Paper under double-blind review

ABSTRACT

Semantic dependency parsing, which aims to find rich bi-lexical relationships, allows words to have multiple dependency heads, resulting in graph-structured representations. We propose an approach to semi-supervised learning of semantic dependency parsers based on the CRF autoencoder framework. Our encoder is a discriminative neural semantic dependency parser that predicts the latent parse graph of the input sentence. Our decoder is a generative neural model that reconstructs the input sentence conditioned on the latent parse graph. Our model is arc-factored and therefore parsing and learning are both tractable. Experiments show our model achieves significant and consistent improvement over the supervised baseline.

1 INTRODUCTION

Semantic dependency parsing (SDP) is a task aiming at discovering sentence-internal linguistic information. The focus of SDP is the identification of predicate-argument relationships for all content words inside a sentence (Oepen et al., 2014; 2015). Compared with syntactic dependencies, semantic dependencies are more general, allowing a word to be either unattached or the argument of multiple predicates. The set of semantic dependencies within a sentence form a directed acyclic graph (DAG), distinguishing SDP from syntactic dependency parsing tasks, where dependencies are usually tree-structured. Extraction of such high-level structured semantic information potentially benefits downstream NLP tasks (Reddy et al., 2017; Schuster et al., 2017).

Several supervised SDP models are proposed in the recent years by modifying syntactic dependency parsers. Their parsing mechanisms are either transition-based (Ribeyre et al., 2014; Kanerva et al., 2015; Wang et al., 2018) or graph-based (Martins & Almeida, 2014; Peng et al., 2017; Dozat & Manning, 2018; Wang et al., 2019).

One limitation of supervised SDP is that labeled SDP data resources are limited in scale and diversity. Due to the rich relationships in SDP, the annotation of semantic dependency graphs is expensive and difficult, calling for professional linguists to design rules and highly skilled annotators to annotate sentences. This limitation becomes more severe with the rise of deep learning, because neural approaches are more data-hungry and susceptible to over-fitting when lacking training data. To alleviate this limitation, we investigate semi-supervised SDP capable of learning from both labeled and unlabeled data.

While a lot of work has been done on supervised SDP, the research of unsupervised and semi-supervised SDP is still lacking. Since parsing results of semantic dependencies are DAGs without the tree-shape restriction, most existing successful unsupervised (Klein & Manning, 2004; I. Spitzkovsky et al., 2010; Jiang et al., 2016; Cai et al., 2017) and semi-supervised (Koo et al., 2008; Druck et al., 2009; Suzuki et al., 2009; Corro & Titov, 2019) learning models for syntactic dependency parsing cannot be applied to SDP directly.

There also exist several unsupervised (Poon & Domingos, 2009; Titov & Klementiev, 2011) and semi-supervised (Das & Smith, 2011; Kočíšký et al., 2016; Yin et al., 2018) methods for semantic parsing, but these models are designed for semantic representations different from dependency graphs, making their adaptation to SDP difficult.

In this work, we propose an end-to-end neural semi-supervised model leveraging both labeled and unlabeled data to learn a dependency graph parser. Our model employs the framework of Condi-

tional Random Field Autoencoder (Ammar et al., 2014), modeling the conditional reconstruction probability given the input sentence with its dependency graph as the latent variable. Our encoder is the supervised model of Dozat & Manning (2018), formulating an SDP task as labeling each arc in a directed graph with a simple neural network. Analogous to a CRF model (Sutton et al., 2012), our encoder is capable of computing the probability of a dependency graph conditioned on the input sentence. The decoder is a generative model based on recurrent neural network language model (Mikolov et al., 2010), which formulates the probability of generating the input sentence, but we take into account the information given by the dependency parse graphs when generating the input.

Our model is arc-factored, i.e., the encoding, decoding and reconstructing probabilities can all be factorized into the product of arc-specific quantities, making both learning and parsing tractable. A unified learning objective is defined that takes advantage of both labeled and unlabeled data. Compared with previous semi-supervised approaches based on Variational Autoencoder (Kingma & Welling, 2013), our learning process does not involve sampling, promising better stability.

We evaluate our model on SemEval 2015 Task 18 Dataset (English) (Open et al., 2015) and find that our model consistently outperforms the state-of-the-art supervised baseline. We also conduct detailed analysis showing the benefits of different amounts of unlabeled data.

2 MODEL

Our model is based on the CRF autoencoder framework (Ammar et al., 2014) which provides a unified fashion for structured predictors to leverage both labeled and unlabeled data. A CRF autoencoder aims to produce a reconstruction of the input \hat{X} from the original input X with an intermediate latent structure Y . It is trained to maximize the conditional reconstruction probability $P(\hat{X} = X|X)$ with the latent variable Y marginalized. Ideally, successful reconstruction implies that the latent structure captures important information of the input.

We adopt the following notations when describing our model. We represent a vector in lowercase bold, e.g., \mathbf{s} , and use a superscript for indexing, e.g., \mathbf{s}^i for the i -th vector. We represent a scalar in lowercase italics, e.g., s , and use a subscript for indexing, e.g., s_i for the i -th element of vector \mathbf{s} . An uppercase italic letter such as Y denotes a matrix. A lower case letter with a subscript pair such as $y_{i,j}$ refers to the element of matrix Y at row i and column j . An uppercase bold letter, e.g., \mathbf{U} , stands for a tensor. We maintain this convention when indexing, e.g., \mathbf{y}_i is the i -th row of matrix Y .

In our model, the input is a natural language sentence consisting of a sequence of words. A sentence with m words is represented by $\mathbf{s} = (s_0, s_1, s_2, \dots, s_m)$, where s_0 is a special token TOP. The latent variable produced by our encoder is a dependency parse graph of the input sentence, represented as a matrix of booleans $Y \in \{0, 1\}^{(m+1) \times (m+1)}$, where $y_{i,j} = 1$ indicates that there exists a dependency arc pointing from word s_i to word s_j . The reconstructed output generated by our decoder is a word sequence $\hat{\mathbf{s}} = (\hat{s}_1, \hat{s}_2, \dots, \hat{s}_m)$.

Our encoder with parameters Θ computes $P_\Theta(Y|\mathbf{s})$, the probability of generating a dependency parse graph Y given a sentence \mathbf{s} . Our decoder with parameters Λ computes $P_\Lambda(\hat{\mathbf{s}}|Y)$, the probability of reconstructing sentence $\hat{\mathbf{s}}$ conditioned on the parse graph Y . The encoder and decoder in combination specify the following conditional distribution.

$$P_{\Theta, \Lambda}(\hat{\mathbf{s}}, Y|\mathbf{s}) = P_\Theta(Y|\mathbf{s})P_\Lambda(\hat{\mathbf{s}}|Y)$$

To compute the conditional probability $P(\hat{\mathbf{s}}|\mathbf{s})$, we sum out the latent variable Y .

$$P_{\Theta, \Lambda}(\hat{\mathbf{s}}|\mathbf{s}) = \sum_{Y \in \mathcal{Y}} P_{\Theta, \Lambda}(\hat{\mathbf{s}}, Y|\mathbf{s})$$

where \mathcal{Y} is the set of all possible dependency parse graphs of \mathbf{s} . During training, we set $\hat{\mathbf{s}} = \mathbf{s}$ and maximize the conditional reconstruction probability $P(\hat{\mathbf{s}}|\mathbf{s})$.

Note that throughout our model, we only consider dependency arc predictions (i.e., whether an arc exists between each word pair). Arc-labels will be learned separately as described in Section 3. We leave the incorporation of arc-label prediction in our model for future work.

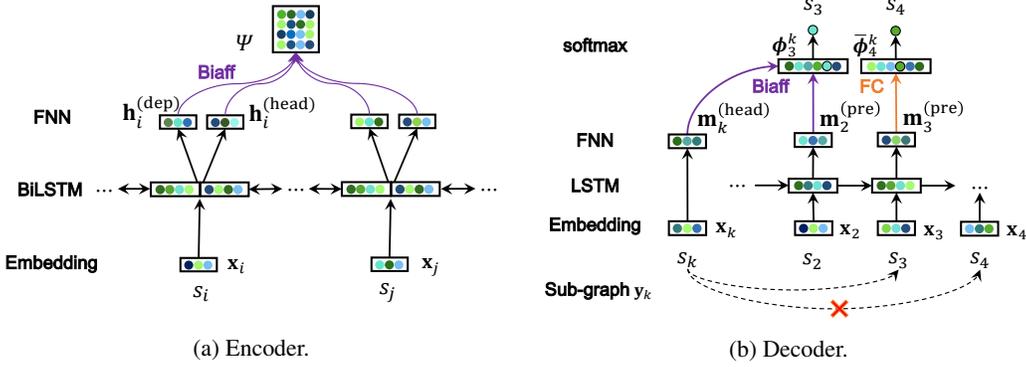


Figure 1: **(a)** Illustration of the encoder structure, following the design of Dozat & Manning (2018). **(b)** Illustration of the decoder generating \hat{s}^k from the k -th neural generator, guided by sub-graph y_k . Dashed arcs at the bottom represent dependency arcs. A cross over arc (s_k, s_4) indicates the absence of this arc.

2.1 ENCODER

Our encoder can be any arc-factored discriminative SDP model. Here we adopt the model of Dozat & Manning (2018), which formulates the semantic dependency parsing task as independently labeling each arc in a directed complete graph. To predict whether or not a directed arc (s_i, s_j) exists, the model computes contextualized representations of s_i and s_j and feeds them into a binary classifier.

The architecture of our encoder is shown in Fig.1a. Word, part-of-speech tag (for short, POS tag), and lemma embeddings¹ of each word in the input sentence are concatenated and fed into a multi-layer bi-directional LSTM to get a contextualized representation of the word.

$$\begin{aligned} \mathbf{x}_i &= \mathbf{e}_i^{(\text{word})} \oplus \mathbf{e}_i^{(\text{tag})} \oplus \mathbf{e}_i^{(\text{lemma})} \\ R &= \text{BiLSTM}(X) \end{aligned} \quad (1)$$

where $\mathbf{e}_i^{(\text{word})}$, $\mathbf{e}_i^{(\text{tag})}$ and $\mathbf{e}_i^{(\text{lemma})}$ are notations for the word, POS tag and lemma embedding respectively, concatenated (\oplus) to form an embedding \mathbf{x}_i for word s_i . Stacking \mathbf{x}_i for $i = 0, 1, \dots, m$ forms matrix X .

The contextualized word representation is then fed into two single-layer feedforward neural networks (FNN) with different parameters to produce two vectors: one for the representation of the word as a dependency head and the other for the representation of the word as a dependent. They are denoted as $\mathbf{h}_i^{(\text{head})}$ and $\mathbf{h}_i^{(\text{dep})}$ respectively.

$$\mathbf{h}_i^{(\text{head})} = \text{FNN}^{(\text{enc-head})}(\mathbf{r}_i) \quad \mathbf{h}_i^{(\text{dep})} = \text{FNN}^{(\text{enc-dep})}(\mathbf{r}_i)$$

Finally, a *biaffine* function is applied to every arc between word pairs (s_i, s_j) to obtain an arc-existence score $\psi_{i,j}$.

$$\psi_{i,j} = \mathbf{h}_i^{(\text{head})} W \mathbf{h}_j^{(\text{dep})} + b$$

where W is a square matrix of size $d \times d$ (d is the size of vector $\mathbf{h}_i^{(\text{head})}$ and $\mathbf{h}_j^{(\text{dep})}$), and b is a scalar. The likelihood of every arc's presence given a sentence, $P(y_{i,j} = 1|\mathbf{s})$, can be computed by applying a *sigmoid* function on score $\psi_{i,j}$. The arc-absence probability $P(y_{i,j} = 0|\mathbf{s})$ is evidently $1 - P(y_{i,j} = 1|\mathbf{s})$.

To conclude, the probability of producing a dependency parse graph Y from the encoder given an input sentence \mathbf{s} can be computed as below.

$$P(Y|\mathbf{s}) = \prod_{i,j} P(y_{i,j}|\mathbf{s})$$

¹Unless stated otherwise, our model makes use of lemma embeddings by default.

2.2 DECODER

Our generative decoder is based on recurrent neural network language models (Mikolov et al., 2010), but we take dependency relationships into account during reconstruction. Our inspiration sources from the decoder with a Graph Convolutional Network (GCN) used by Corro & Titov (2019) to incorporate tree-structured syntactic dependencies when generating sentences, but our decoder differs significantly from theirs in that ours handles parse graphs and is arc-factored.

As mentioned above, semantic dependency parsing allows a word to have multiple dependency heads. If we generate a word conditioned on multiple heads, then it becomes difficult to make the decoder arc-factored and hence inference and learning becomes less tractable. Instead, we propose to generate a word for multiple times, each time conditioned on a different head.

Specifically, we split dependency graph Y of a sentence $\mathbf{s} = (s_0, s_1, \dots, s_m)$ with m words and a TOP token into $m + 1$ parts:

$$Y = [\mathbf{y}_0; \mathbf{y}_1; \mathbf{y}_2; \dots; \mathbf{y}_m]$$

Each \mathbf{y}_i is the i -th row of Y , representing a sub-graph where arcs are rooted at the i -th word of the sentence \mathbf{s} . Mathematically, we have $\mathbf{y}_i = \{y_{i,j} | j \in (1, 2, \dots, m)\}$.

We then generate $m+1$ sentences $(\hat{\mathbf{s}}^0, \hat{\mathbf{s}}^1, \hat{\mathbf{s}}^2, \dots, \hat{\mathbf{s}}^m)$ using $m+1$ neural generators. The generation of sentence $\hat{\mathbf{s}}^k$ is guided by the k -th sub-graph \mathbf{y}_k . Each generator is a left-to-right LSTM language model and computes $P_\Lambda(\hat{s}_i^k | \hat{\mathbf{s}}_{1:i-1}^k, y_{k,i})$, the probability of generating each word conditioned on its preceding words and whether \mathbf{y}_k contains a dependency arc to the word. We share parameters among all the $m + 1$ generators.

Fig.1b shows an example for computing the generative probability of \hat{s}_i^k by the k -th generator ($k \in \{0, 1, \dots, m\}$) that incorporates the information of the k -th sub-graph \mathbf{y}_k . Recall that \mathbf{y}_k contains only dependencies rooted at s_k . Below we describe how to compute the generative probability of each word \hat{s}_i^k with and without the dependency arc (s_k, s_i) respectively.

Generative probability with a dependency Suppose there is a dependency arc from s_k to s_i , we need to compute the generative probability $P_\Lambda(\hat{s}_i^k | \hat{\mathbf{s}}_{1:i-1}^k, y_{k,i} = 1)$. The LSTM in the k -th generator takes the embedding of the previous word s_{i-1} computed through Eq.1 as its input and outputs the hidden state \mathbf{g}_{i-1} , which is fed into an FNN to produce a representation $\mathbf{m}_{i-1}^{(\text{pre})}$. Meanwhile, the embedding of the k -th word (also computed through Eq.1) is fed into another FNN to get its representation $\mathbf{m}_k^{(\text{head})}$ as a dependency head.

$$G = \text{LSTM}(X)$$

$$\mathbf{m}_{i-1}^{(\text{pre})} = \text{FNN}^{(\text{dec-pre})}(\mathbf{g}_{i-1}) \quad (2)$$

$$\mathbf{m}_k^{(\text{head})} = \text{FNN}^{(\text{dec-head})}(\mathbf{x}_k) \quad (3)$$

$\mathbf{m}_k^{(\text{head})}$ and $\mathbf{m}_{i-1}^{(\text{pre})}$ are fed into a *bilinear* function to obtain a vocabulary-size score vector ϕ_i^k .

$$\phi_i^k = \mathbf{m}_k^{(\text{head})} \mathbf{U} \mathbf{m}_{i-1}^{(\text{pre})} \quad (4)$$

Here, \mathbf{U} is a tensor of size $d \times V \times d$, where V is the vocabulary size and d is the size of vector $\mathbf{m}_k^{(\text{head})}$ and $\mathbf{m}_{i-1}^{(\text{pre})}$. To conserve parameters, the tensor \mathbf{U} is diagonal (i.e., $u_{i,k,j} = 0$ wherever $i \neq j$). A *softmax* function can then be applied to ϕ_i^k , from which we pick the generative probability of \hat{s}_i^k .

Generative probability without a dependency Suppose there is no dependency arc from s_k to s_i . In this case, reconstruction of \hat{s}_i^k resembles a normal recurrent neural network language model. The representation $\mathbf{m}_{i-1}^{(\text{pre})}$ from Eq.2 is fed into a fully connected layer to get $\bar{\phi}_i^k$, a vector of vocabulary size containing generative scores of all the words.

$$\bar{\phi}_i^k = \text{FC}(\mathbf{m}_{i-1}^{(\text{pre})}) \quad (5)$$

The generative probability $P_\Lambda(\hat{s}_i^k | \hat{\mathbf{s}}_{1:i-1}^k, y_{k,i} = 0)$ can then be computed by applying a *softmax* function on $\bar{\phi}_i^k$ and selecting the corresponding probability of \hat{s}_i^k . Since we simply reconstruct word

s_i without considering the dependency arc information, this probability is exactly the same in the $m + 1$ generators and only needs to be computed once.

To conclude the overall design of our decoder, it is worth noting that in $m + 1$ generation processes, parameters among all LSTMs are shared, as well as those among all FNNs² and FCs. Still, embeddings in Eq.1 are shared among both encoder and decoder.

With $P(\hat{s}_i^k | \hat{s}_{1:i-1}^k, y_{k,i})$ computed for $i = 1, \dots, m, k = 0, 1, \dots, m$, the probability of generating $\hat{s}^0, \hat{s}^1, \hat{s}^2, \dots, \hat{s}^m$ from dependency graph Y can be computed through:

$$P_\Lambda(\hat{s}^0, \hat{s}^1, \dots, \hat{s}^m | Y) = \prod_{k=0}^m P_\Lambda(\hat{s}^k | y_k) = \prod_{k=0}^m \prod_{i=1}^m P_\Lambda(\hat{s}_i^k | \hat{s}_{1:i-1}^k, y_{k,i})$$

In our model, we are only interested in the case where all the $m + 1$ sentences are the same. In addition, to balance the influence of the encoder and the decoder, we take the geometric mean of the $m + 1$ probabilities. The final decoding probability is defined as follows.

$$P_\Lambda(\hat{s} | Y) := \prod_{i=1}^m \prod_{k=0}^m \sqrt[m+1]{P_\Lambda(\hat{s}_i | \hat{s}_{1:i-1}, y_{k,i})}$$

Note that this is not a properly normalized probability distribution, but in practice we find it sufficient for semi-supervised SDP.

2.3 PARSING

Given parameters $\{\Theta, \Lambda\}$ of our encoder and decoder, we can parse a sentence \mathbf{s} by finding a $Y \in \mathcal{Y}(\mathbf{s})$ which maximizes probability $P(\hat{\mathbf{s}} = \mathbf{s}, Y | \mathbf{s})$, where $\mathcal{Y}(\mathbf{s})$ is the set of all parse graphs of sentence \mathbf{s} .

$$\begin{aligned} Y^* &= \arg \max_{Y \in \mathcal{Y}(\mathbf{s})} \log P_{\Theta, \Lambda}(\hat{\mathbf{s}}, Y | \mathbf{s}) = \arg \max_{Y \in \mathcal{Y}(\mathbf{s})} \log P_\Lambda(\hat{\mathbf{s}} | Y) P_\Theta(Y | \mathbf{s}) \\ &= \arg \max_{Y \in \mathcal{Y}(\mathbf{s})} \log \prod_{i,j} \sqrt[m+1]{P_\Lambda(\hat{s}_j | \hat{s}_{1:j-1}, y_{i,j})} P_\Theta(y_{i,j} | \mathbf{s}) \\ &= \arg \max_{Y \in \mathcal{Y}(\mathbf{s})} \sum_{i,j} \left(\frac{1}{m+1} \log P_\Lambda(\hat{s}_j | \hat{s}_{1:j-1}, y_{i,j}) + \log P_\Theta(y_{i,j} | \mathbf{s}) \right) \end{aligned} \quad (6)$$

Since the probability is arc-factored, we can determine the existence of each dependency arc independently by picking the value of $y_{i,j}$ that maximizes the corresponding term. The time complexity of our parsing algorithm is $O(m^2)$ for a sentence with m words.

3 LEARNING

Since we want to train our model in a semi-supervised manner, we design loss functions for labeled and unlabeled data respectively. For each training sentence \mathbf{s} , the overall loss function is defined as a combination of supervised loss \mathcal{L}_l and unsupervised loss \mathcal{L}_u .

$$\mathcal{L}(\mathbf{s}) = \iota(\mathbf{s}) * \mathcal{L}_l(\mathbf{s}) + (1 - \iota(\mathbf{s})) * \rho \mathcal{L}_u(\mathbf{s}) \quad (7)$$

where an indicator $\iota(\mathbf{s}) \in \{0, 1\}$ specifies whether training sentence \mathbf{s} is labeled or not and a tunable constant ρ balances the two losses.

Supervised Loss For any labeled sentence (\mathbf{s}, Y^*) , where \mathbf{s} stands for a sentence and Y^* stands for a gold parse graph, we can compute the discriminative loss.

$$\mathcal{L}_l(\mathbf{s}) = -\log P_{\Theta, \Lambda}(\hat{\mathbf{s}} = \mathbf{s}, Y^* | \mathbf{s}) \quad (8)$$

²FNN^(dec-pre) and FNN^(dec-head) never share parameters between each other, since their usages are different.

Following the derivation of Eq.6, we have:

$$\log P_{\Theta,\Lambda}(\hat{\mathbf{s}}, Y^*|\mathbf{s}) = \sum_{i,j} \left(\frac{1}{m+1} \log P_{\Lambda}(\hat{s}_j|\hat{\mathbf{s}}_{1:j-1}, y_{i,j}^*) + \log P_{\Theta}(y_{i,j}^*|\mathbf{s}) \right) \quad (9)$$

Gold parses also provide a label for each dependency. We follow Dozat & Manning (2018) and model dependency labels with a purely supervised module on top of the BiLSTM layer of the encoder. Its parameters are learned by optimizing a cross-entropy loss function.

Unsupervised Loss For any unlabeled sentence \mathbf{s} , we maximize the conditional reconstruction probability $P(\hat{\mathbf{s}} = \mathbf{s}|\mathbf{s})$. The unsupervised loss is:

$$\begin{aligned} \mathcal{L}_u(\mathbf{s}) &= -\log P_{\Theta,\Lambda}(\hat{\mathbf{s}}|\mathbf{s}) = -\log \sum_{Y \in \mathcal{Y}(\mathbf{s})} P_{\Theta,\Lambda}(Y, \hat{\mathbf{s}}|\mathbf{s}) \\ &= -\log \sum_{Y \in \mathcal{Y}(\mathbf{s})} P_{\Lambda}(\hat{\mathbf{s}}|Y)P_{\Theta}(Y|\mathbf{s}) = -\sum_{i,j} \log \sum_{y_{i,j} \in \{0,1\}} \left(P_{\Theta}(y_{i,j}|\mathbf{s}) \times \sqrt[m+1]{P_{\Lambda}(\hat{s}_j|\hat{\mathbf{s}}_{1:j-1}, y_{i,j})} \right) \end{aligned} \quad (10)$$

Derivations of Eq.10 are provided in Appendix A. Given a dataset containing both labeled and unlabeled sentences, our model can be trained end-to-end by optimizing the loss function Eq.7 over the combined dataset using any gradient based method.

4 EXPERIMENTS

4.1 SETTINGS

Dataset We examine the performance of our model on the English corpus of the SDP 2014 & 2015: Broad Coverage Semantic Dependency Parsing dataset (Oepen et al., 2015). The corpus is composed of three distinct and parallel semantic dependency annotations (DM, PAS, PSD) of Sections 00-21 of the WSJ Corpus, as well as a balanced sample of twenty files from the Brown Corpus. More information of this dataset is shown in Table 3 in Appendix B. We evaluate the performance of models through two metrics: Unlabeled F1 score (UF1) and Labeled F1 score (LF1). UF1 measures the accuracy of the binary classification of arc existence, while LF1 measures the correctness of each arc-label as well.

Network Configuration For our encoder, we adopt the hyper-parameters of Dozat & Manning (2018). Following Dozat & Manning (2018), words or lemmas whose occurrences are less than 7 times within the training set are treated as UKN. For our decoder, we set the number of layer(s) of uni-directional LSTM to 1, whose recurrent hidden size is 600. For $\mathbf{FNN}^{(\text{dec-head})}$ and $\mathbf{FNN}^{(\text{dec-pre})}$, the output sizes are both 400, activated by a $\tanh(\cdot)$ function.

Learning Our loss function (Eq.7) is optimized by the Adam+AMSGrad optimizer (Reddi et al., 2018), with hyper-parameters β_1, β_2 kept the same as those of Dozat & Manning (2018). The interpolation constant ρ is tuned with the size of unlabeled data. A detailed table of hyper-parameter values is provided in Appendix C. The training time for one batch with our autoencoder is 2–3 times of that of Dozat & Manning (2018) because of the extra decoder.

4.2 VARYING SIZE OF UNLABELED DATA

In our first experiment (with the DM annotations only), we fix the amount of labeled data and continuously incorporate more unlabeled data into the training set. Specifically, we randomly sample 10% of the whole dataset as labeled data. Unlabeled data are then sampled from the remaining part (with their gold parses removed), with a proportion increasing from 0% to 90% of the complete dataset. For unlabeled data, we find that long sentences do not help in improving F1 scores and therefore in this and all the subsequent experiments we remove unlabeled sentences longer than 20 to reduce the running time and memory usage.

Experimental results are visualized in Fig.2. First, we observe that in the purely supervised setting (i.e., +0% unlabeled data), our model already outperforms the baseline (Dozat & Manning, 2018)

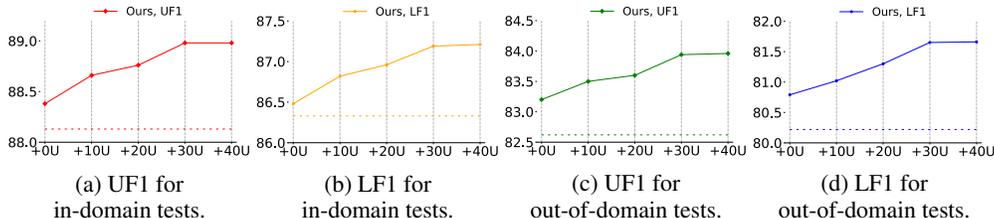


Figure 2: Results with fixed amount of labeled data and varying amount of unlabeled data. Experiments are carried out on DM representations. **+10U** represents using 10% unlabeled data, and so on. Numbers on Y-axes represent F1 scores. Dashed horizontal lines are results of the supervised baseline (Dozat & Manning, 2018) trained on labeled data only. Solid lines are our results.

Models		Labeled:Unlabeled									
		0.1:9.9		1:9		3:7		5:5		10:0	
		UF1	LF1	UF1	LF1	UF1	LF1	UF1	LF1	UF1	LF1
id	D&M	75.15	70.8	88.13	86.33	91.61	90.41	92.91	92.03	94.1	93.34
	Ours-Sup	75.22	70.66	88.38	86.48	92.02	90.87	93.15	92.24	94.23	93.49
	Ours-Semi	76.96	72.43	89.01	87.19	92.22	91.15	93.09	92.25	-	-
ood	D&M	70.03	65.19	82.62	80.22	86.87	85.14	88.51	87.03	89.97	88.84
	Ours-Sup	70.39	65.70	83.20	80.79	87.29	85.49	88.84	87.49	90.23	89.15
	Ours-Semi	72.00	67.15	83.96	81.61	87.54	85.86	88.94	87.55	-	-

Table 1: Experimental results with varying proportions of labeled and unlabeled data. **D&M** stands for the supervised model of Dozat & Manning (2018) trained on labeled data only. **Ours-Sup** stands for our model trained on labeled data only. **Ours-Semi** stands for our model trained on both labeled and unlabeled data.

and the advantage of our model is larger on the out-of-domain test set than on the in-domain test set. Since our encoder is exactly the baseline model, this shows the benefit of adding the decoder for joint learning and parsing even in the supervised setting.

Second, with an increasing size of unlabeled dataset from 0% to 30%, we see a clear increase in performance of our model, suggesting the benefit of semi-supervised learning with our model. However, we observe little improvement when the size of unlabeled data exceeds 40% (not shown in the figure), indicating a possible upper bound of the effectiveness of unlabeled data.

4.3 VARYING PROPORTION OF UNLABELED DATA

In our second experiment (again with the DM annotations), we use the full training set and vary the proportion of labeled and unlabeled data. Experimental results are shown in Table 1. Our semi-supervised model shows the largest advantage over the supervised models with the 0.1:9.9 proportion (which contains only 339 labeled sentences). With the increased proportion of labeled data, the performance of all the models goes up, but the advantage of our semi-supervised model vanishes. This demonstrates the diminishing effectiveness of unlabeled data when adding more labeled data. Another worth-noting observation is that the superiority of our semi-supervised model is much stronger on the out-of-domain tests and does not vanish even with the 5:5 proportion. This suggests good generalizability of our semi-supervised model.

4.4 ON ALL REPRESENTATIONS

In the previous two experiments, we evaluate our model on the DM representation. Here we evaluate our model on all the three representations: DM, PAS and PSD. We slightly tune the hyper-parameters based on the optimal values from the previous experiments of the DM representation. We use 10% of the sentences as labeled data and the rest 90% of the sentences as unlabeled data. For the completeness of our experiment, we follow Dozat & Manning (2018) and examine four different word representations: **basic** (i.e., using only word and POS tag embeddings), **+Lemma** (i.e., using word,

	Models	DM		PAS		PSD		Avg	
		UF1	LF1	UF1	LF1	UF1	LF1	UF1	LF1
id, +Lemma	D&M	88.13	86.33	91.99	90.61	88.24	73.20	89.45	83.38
	Ours-Sup	88.38	86.48	92.27	91.00	88.28	73.33	89.64	83.60
	Ours-Semi	89.01	87.19	92.10	90.88	88.64	73.69	89.92	83.92
ood, +Lemma	D&M	82.62	80.22	88.25	86.22	85.07	70.68	85.31	79.04
	Ours-Sup	83.20	80.79	88.68	86.79	85.03	71.20	85.64	79.59
	Ours-Semi	83.96	81.61	88.83	86.85	85.32	71.56	86.04	80.01

Table 2: Experimental results on all the three representations. **D&M** stands for the supervised model of Dozat & Manning (2018). **Ours-Sup** stands for our model trained on labeled data only. **Ours-Semi** stands for our model trained on both labeled and unlabeled data.

POS tag and lemma embeddings), **+Char** (i.e., using word, POS tag and character embeddings) and **+Lemma+Char** (i.e. using word, POS tag, lemma and character embeddings).

Table 2 shows the experimental results of **+Lemma**, the default word representation. The results of the other word representations show very similar trends (see Appendix D). We observe significant improvement of our semi-supervised model over the two supervised baselines on both DM and PSD representations. However, it is surprising to find that on the PAS representation, our semi-supervised model exhibits little advantage over its supervised counterpart. One possible explanation, as Dozat & Manning (2018) also noted, is that PAS is the easiest of the three representations (as can be seen by comparing the scores of the three representations in Table 2) and our supervised model may already reach the performance ceiling.

5 RELATED WORK

Work on unsupervised or semi-supervised dependency parsing, to the best of our knowledge, is dominated by tree-structured parsing (Koo et al., 2008; Druck et al., 2009; Suzuki et al., 2009). Recently, Corro & Titov (2019) introduced an approximate inference method with a Variational Autoencoder (Kingma et al., 2014) for semi-supervised syntactic dependency parsing. Our decoder is inspired by their work, but differs from theirs in that our decoder handles parse graphs and is arc-factored. Cai et al. (2017) used the framework of CRF Autoencoder (Ammar et al., 2014) to perform unsupervised syntactic dependency parsing. The same framework has been used by Zhang et al. (2017) for semi-supervised sequence labeling. Our work also adopts the CRF Autoencoder framework, but with both the encoder and the decoder redesigned for semantic dependency parsing.

Existing unsupervised and semi-supervised approaches to semantic parsing focused on semantic representations different from dependency graphs, e.g., general-purpose logic forms (Sondheimer & Nebel, 1986) and formal meaning representations (Bordes et al., 2012). Therefore, these approaches cannot be applied to SDP directly. Poon & Domingos (2009) presented the first unsupervised semantic parser to transform dependency trees into quasi-logical forms with Markov logic. Following this work, Titov & Klementiev (2011) proposed a non-parametric Bayesian model for unsupervised semantic parsing using hierarchical Pitman-Yor process (Teh, 2006). Das & Smith (2011) described a semi-supervised approach to frame-semantic parsing. Kočíšký et al. (2016) proposed a semi-supervised semantic parsing approach making use of unpaired logical forms. Recently, Yin et al. (2018) proposed a variational autoencoding model for semi-supervised semantic parsing of tree-structured semantic representations.

6 CONCLUSION

In this work, we proposed a semi-supervised learning model for semantic dependency parsing using CRF Autoencoders. Our model is composed of a discriminative neural encoder producing a dependency graph conditioned on an input sentence, and a generative neural decoder for input reconstruction based on the dependency graph. The model works in an arc-factored fashion, promising end-to-end learning and efficient parsing. We evaluated our model under both full-supervision settings and semi-supervision settings. Our model outperforms the baseline on multiple target representations. By adding unlabeled data, our model exhibits further performance improvements.

REFERENCES

- Waleed Ammar, Chris Dyer, and Noah A. Smith. Conditional Random Field Autoencoders for Unsupervised Structured Prediction. *Advances in Neural Information Processing Systems*, 4: 3311–3319, 2014.
- Antoine Bordes, Xavier Glorot, Jason Weston, and Yoshua Bengio. Joint Learning of Words and Meaning Representations for Open-text Semantic Parsing. In *Artificial Intelligence and Statistics*, pp. 127–135, 2012.
- Jiong Cai, Yuanyuan Jiang, and Kewei Tu. CRF Autoencoder for Unsupervised Dependency Parsing. In *EMNLP*, 2017.
- Caio Corro and Ivan Titov. Differentiable Perturb-and-Parse: Semi-Supervised Parsing with a Structured Variational Autoencoder. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=BJlgNh0qKQ>.
- Dipanjan Das and Noah A Smith. Semi-Supervised Frame-Semantic Parsing for Unknown Predicates. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pp. 1435–1444. Association for Computational Linguistics, 2011.
- Timothy Dozat and Christopher D. Manning. Simpler but More Accurate Semantic Dependency Parsing. In *ACL*, 2018.
- Gregory Druck, Gideon S. Mann, and Andrew McCallum. Semi-Supervised Learning of Dependency Parsers using Generalized Expectation Criteria. In *IJCNLP-ACL*, 2009.
- Valentin I. Spitzkovsky, Hiyan Alshawi, Daniel Jurafsky, and Christopher Manning. Viterbi Training Improves Unsupervised Dependency Parsing. *CoNLL 2010 - Fourteenth Conference on Computational Natural Language Learning, Proceedings of the Conference*, pp. 9–17, 07 2010.
- Yong Jiang, Wenjuan Han, and Kewei Tu. Unsupervised Neural Dependency Parsing. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 763–771, 2016.
- Jenna Kanerva, Juhani Luotolahti, and Filip Ginter. Turku: Semantic Dependency Parsing as A Sequence Classification. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pp. 965–969, 2015.
- Diederik Kingma and Max Welling. Auto-Encoding Variational Bayes. *ICLR*, 12 2013.
- Durk P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. Semi-Supervised Learning with Deep Generative Models. In *Advances in neural information processing systems*, pp. 3581–3589, 2014.
- Dan Klein and Christopher D Manning. Corpus-based Induction of Syntactic Structure: Models of Dependency and Constituency. In *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*, pp. 478. Association for Computational Linguistics, 2004.
- Tomáš Kočiský, Gábor Melis, Edward Grefenstette, Chris Dyer, Wang Ling, Phil Blunsom, and Karl Moritz Hermann. Semantic Parsing with Semi-Supervised Sequential Autoencoders. In *EMNLP*, 2016.
- Terry Koo, Xavier Carreras, and Michael Collins. Simple Semi-Supervised Dependency Parsing. In *Proceedings of ACL-08*, 2008.
- André FT Martins and Mariana SC Almeida. Priberam: A Turbo Semantic Parser with Second Order Features. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pp. 471–476, 2014.
- Tomáš Míkolov, Martin Karafiát, Lukáš Burget, Jan Černocký, and Sanjeev Khudanpur. Recurrent Neural Network Based Language Model. In *Eleventh annual conference of the international speech communication association*, 2010.

- Stephan Oepen, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Dan Flickinger, Jan Hajić, Angelina Ivanova, and Yi Zhang. Semeval 2014 task 8: Broad-Coverage Semantic Dependency Parsing. In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pp. 63–72, 2014.
- Stephan Oepen, Marco Kuhlmann, Yusuke Miyao, Daniel Zeman, Silvie Cinkov, Dan Flickinger, Jan Hajić, and Zdeka Ureov. SemEval 2015 Task 18: Broad-Coverage Semantic Dependency Parsing. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pp. 915–926, Denver, CO, USA, 2015.
- Hao Peng, Sam Thomson, and Noah A Smith. Deep Multitask Learning for Semantic Dependency Parsing. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2017.
- Hoifung Poon and Pedro Domingos. Unsupervised Semantic Parsing. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 1-Volume 1*, pp. 1–10. Association for Computational Linguistics, 2009.
- Sashank J. Reddi, Satyen Kale, and Sanjiv Kumar. On the Convergence of Adam and Beyond. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=ryQu7f-RZ>.
- Siva Reddy, Oscar Tackstrom, Slav Petrov, Mark Steedman, and Mirella Lapata. Universal Semantic Parsing. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017. URL <http://aclweb.org/anthology/D/D17/D17-1009.pdf>.
- Corentin Ribeyre, Éric Villemonte de La Clergerie, and Djamel Seddah. Alpage: Transition-based Semantic Graph Parsing with Syntactic Features. In *International Workshop on Semantic Evaluation*, 2014.
- Sebastian Schuster, Éric Villemonte de La Clergerie, Marie Candito, Benoît Sagot, Christopher Manning, and Djamel Seddah. Paris and Stanford at EPE 2017: Downstream Evaluation of Graph-based Dependency Representations. In *EPE 2017-The First Shared Task on Extrinsic Parser Evaluation*, pp. 47–59, 2017.
- Norman K Sondheimer and Bernhard Nebel. A Logical-Form and Knowledge-Base Design for Natural Language Generation. In *Proceedings of the workshop on Strategic computing natural language*, pp. 231–241. Association for Computational Linguistics, 1986.
- Charles Sutton, Andrew McCallum, et al. An Introduction to Conditional Random Fields. *Foundations and Trends® in Machine Learning*, 4(4):267–373, 2012.
- Jun Suzuki, Hideki Isozaki, Xavier Carreras, and Michael Collins. An Empirical Study of Semi-Supervised Structured Conditional Models for Dependency Parsing. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 2-Volume 2*, pp. 551–560. Association for Computational Linguistics, 2009.
- Yee Whye Teh. A Hierarchical Bayesian Language Model Based on Pitman-Yor Processes. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, pp. 985–992. Association for Computational Linguistics, 2006.
- Ivan Titov and Alexandre Klementiev. A Bayesian Model for Unsupervised Semantic Parsing. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pp. 1445–1455. Association for Computational Linguistics, 2011.
- Xinyu Wang, Jingxian Huang, and Kewei Tu. Second-Order Semantic Dependency Parsing with End-to-End Neural Networks. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4609–4618, Florence, Italy, July 2019. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/P19-1454>.

Yuxuan Wang, Wanxiang Che, Jiang Guo, and Ting Liu. A Neural Transition-Based Approach for Semantic Dependency Graph Parsing. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

Pengcheng Yin, Chunting Zhou, Junxian He, and Graham Neubig. StructVAE: Tree-Structured Latent Variable Models for Semi-Supervised Semantic Parsing. In *ACL*, 2018.

Xiao Zhang, Yong Jiang, Hao Peng, Kewei Tu, and Dan Goldwasser. Semi-Supervised Structured Prediction with Neural CRF Autoencoder. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 1701–1711, 2017.

A DETAILED DERIVATION

Derivation of the marginalized probability over all possible dependency graphs of a sentence with m words for Eq.10 is shown below.

$$\begin{aligned}
 \log \sum_{Y \in \mathcal{Y}} P(s, Y | s) &= \log \sum_{Y \in \mathcal{Y}} P(s | Y) P(Y | s) \\
 &= \log \sum_{Y \in \mathcal{Y}} \left(\prod_{i,j} P(y_{i,j} | s) \sqrt[m+1]{P(\hat{s}_j | \hat{s}_{1:j-1}, y_{i,j})} \right) \\
 &= \log \sum_{y_{0,1}} \sum_{y_{0,2}} \cdots \sum_{y_{m,m}} \left(P(y_{i,j} | s) \sqrt[m+1]{P(\hat{s}_j | \hat{s}_{1:j-1}, y_{i,j})} \right) \\
 &= \log \left(\left(\sum_{y_{0,1}} P(y_{0,1} | s) \sqrt[m+1]{P(\hat{s}_1 | \hat{s}_0, y_{0,1})} \right) \cdots \left(\sum_{y_{m,m}} P(y_{m,m} | s) \sqrt[m+1]{P(\hat{s}_m | \hat{s}_{1:m-1}, y_{m,m})} \right) \right) \\
 &= \log \prod_{i,j} \left(\sum_{y_{i,j}} P(y_{i,j} | s) \sqrt[m+1]{P(\hat{s}_j | \hat{s}_{1:j-1}, y_{i,j})} \right) \\
 &= \sum_{i,j} \log \left(\sum_{y_{i,j}} \left(P(y_{i,j} | s) \sqrt[m+1]{P(\hat{s}_j | \hat{s}_{1:j-1}, y_{i,j})} \right) \right)
 \end{aligned}$$

B DATASET

Table 3 shows the source and size of each part of the English corpus of the SDP 2014 & 2015: Broad Coverage Semantic Dependency Parsing dataset (Open et al., 2015).

	Source	Sentences	Tokens
train	WSJ Sec.00-20	35,657	802,717
test (id)	WSJ Sec.21	1,410	31,948
test (ood)	Brown	1,849	31,583

Table 3: The sources and sizes of the training and test sets of the SDP 2014 & 2015 (English) dataset. We extract WSJ Section 20 (1,692 sentences) from the train set for development purpose. **id** stands for in-domain testing, while **ood** stands for out-of-domain testing.

C HYPER-PARAMETERS

Here we provide a summary of hyper-parameters in our experiments, as shown in Table 4.

D EXPERIMENTS ON ALL REPRESENTATIONS

Complete experimental results on DM, PAS and PSD under the setting of Section 4.4 (i.e., **basic**, **+Char**, **+Lemma** and **+Lemma+Char**) are shown in Table 6.

E STABILITY OF OUR MODEL

To test the stability of our model, we repeat the experiment of Section 4.4 on the DM annotation for three times, each time with a differently sampled dataset. Table 5 shows the results. We observe consistent advantage of our model over the baseline on all the three datasets.

Hidden Layer	Hidden Sizes
Word/GloVe/POS/Lemma/Char	100
GloVe Linear	125
Encoder BiLSTM	3*600
Encoder FNN(head)	3*400
Encoder FNN(dep)	3*400
Decoder UniLSTM	1*600
Decoder FNN(head)	1*400
Decoder FNN(pre)	1*400
Dropouts	Dropout Prob.
Word/GloVe/POS/Lemma	20%
Encoder FNN	25 %
BiLSTM (FF/recur)	45%/25%
Optimizer & Loss	Value
Adam β_1	0
Adam β_2	0.95
Learning rate	$1e^{-3}$
L2 regularization	$3e^{-9}$

Table 4: Summary of hyper-parameters.

	Models	Random1		Random2		Random3		Avg	
		UF1	LF1	UF1	LF1	UF1	LF1	UF1	LF1
id, +Lemma	D&M	88.13	86.33	88.19	86.38	88.80	87.20	88.37	86.64
	Ours-Sup	88.38	86.48	88.65	86.94	88.93	87.27	88.65	86.90
	Ours-Semi	89.01	87.19	89.09	87.33	89.27	87.46	89.12	87.33
ood, +Lemma	D&M	82.62	80.22	83.32	80.96	83.20	80.84	83.05	80.67
	Ours-Sup	83.20	80.79	83.79	81.34	83.58	81.24	83.52	81.12
	Ours-Semi	83.96	81.61	84.25	81.93	83.88	81.44	84.03	81.66

Table 5: Experimental results on three randomly sampled datasets. **D&M** stands for the supervised model of Dozat & Manning (2018). **Ours-Sup** stands for our model trained on labeled data only. **Ours-Semi** stands for our model trained on both labeled and unlabeled data.

	Models	DM		PAS		PSD		Avg	
		UF1	LF1	UF1	LF1	UF1	LF1	UF1	LF1
id, basic	D&M	86.92	84.77	91.57	90.18	87.31	71.71	88.60	82.22
	Ours-Sup	87.64	85.29	92.06	90.75	87.67	72.25	89.12	82.76
	Ours-Semi	88.32	85.95	91.93	90.65	88.27	73.30	89.51	83.30
ood, basic	D&M	81.52	78.77	88.22	86.20	84.05	69.53	84.60	78.17
	Ours-Sup	82.23	79.36	88.52	86.55	84.16	69.80	84.97	78.57
	Ours-Semi	83.03	80.04	88.46	86.53	84.68	70.72	85.39	79.10
id, +Char	D&M	87.38	85.35	91.85	90.46	87.64	72.28	88.96	82.70
	Ours-Sup	87.81	85.80	92.08	90.85	87.91	72.68	89.27	83.11
	Ours-Semi	88.20	85.87	91.97	90.74	88.12	73.19	89.43	83.27
ood, +Char	D&M	82.14	79.55	88.30	86.25	84.43	69.92	84.96	78.57
	Ours-Sup	82.34	79.71	88.62	86.71	84.60	70.34	85.19	78.92
	Ours-Semi	82.60	79.84	88.66	86.81	84.69	70.83	85.32	79.16
id, +Lemma	D&M	88.13	86.33	91.99	90.61	88.24	73.20	89.45	83.38
	Ours-Sup	88.38	86.48	92.27	91.00	88.28	73.33	89.56	83.54
	Ours-Semi	89.01	87.19	92.10	90.88	88.64	73.69	89.85	83.85
id, +Lemma	D&M	88.13	86.33	91.99	90.61	88.24	73.20	89.45	83.38
	Ours-Sup	88.38	86.48	92.27	91.00	88.28	73.33	89.64	83.60
	Ours-Semi	89.01	87.19	92.10	90.88	88.64	73.69	89.92	83.92
ood, +Lemma	D&M	82.62	80.22	88.25	86.22	85.07	70.68	85.31	79.04
	Ours-Sup	83.20	80.79	88.68	86.79	85.03	71.20	85.64	79.59
	Ours-Semi	83.96	81.61	88.83	86.85	85.32	71.56	86.04	80.01
ood, +Lemma+Char	D&M	83.26	80.96	88.51	86.54	85.18	71.02	85.65	79.51
	Ours-Sup	83.59	81.48	88.80	86.91	85.39	71.18	85.93	79.86
	Ours-Semi	84.29	82.02	88.78	86.85	85.49	71.58	86.19	80.15

Table 6: Experimental results on all the three SDP representations with four different word representations. **D&M** stands for the supervised model of Dozat & Manning (2018). **Ours-Sup** stands for our model trained on labeled data only. **Ours-Semi** stands for our model trained on both labeled and unlabeled data.