

Efficient Dependency-Guided Named Entity Recognition

Zhanming Jie, Aldrian Obaja Muis, Wei Lu

Singapore University of Technology and Design
 8 Somapah Road, Singapore, 487372
 zhanming_jie@mymail.sutd.edu.sg, {aldrian_muis,luwei}@sutd.edu.sg

Abstract

Named entity recognition (NER), which focuses on the extraction of semantically meaningful named entities and their semantic classes from text, serves as an indispensable component for several down-stream natural language processing (NLP) tasks such as relation extraction and event extraction. Dependency trees, on the other hand, also convey crucial semantic-level information. It has been shown previously that such information can be used to improve the performance of NER (Sasano and Kurohashi 2008; Ling and Weld 2012). In this work, we investigate on how to better utilize the structured information conveyed by dependency trees to improve the performance of NER. Specifically, unlike existing approaches which only exploit dependency information for designing local features, we show that certain global structured information of the dependency trees can be exploited when building NER models where such information can provide guided learning and inference. Through extensive experiments, we show that our proposed novel *dependency-guided* NER model performs competitively with models based on conventional semi-Markov conditional random fields, while requiring significantly less running time.

Introduction

Named entity recognition (NER) is one of the most important tasks in the field of natural language processing (NLP). The task focuses on the extraction of named entities together with their semantic classes (such as *organization* or *person*) from text. The extracted named entity information has been shown to be useful in various NLP tasks, including coreference resolution, question answering and relation extraction (Lao and Cohen 2010; Lee et al. 2012; Krishnamurthy and Mitchell 2015).

Dependency trees, on the other hand, were shown to be useful in several semantic processing tasks in NLP such as semantic parsing and question answering (Poon and Domingos 2009; Liang, Jordan, and Klein 2013). The dependency structures convey semantic-level information which was shown to be useful for the NER task. Existing research efforts have exploited such dependency structured information by designing dependency-related local features that can be used in the NER models (Sasano and Kurohashi 2008;

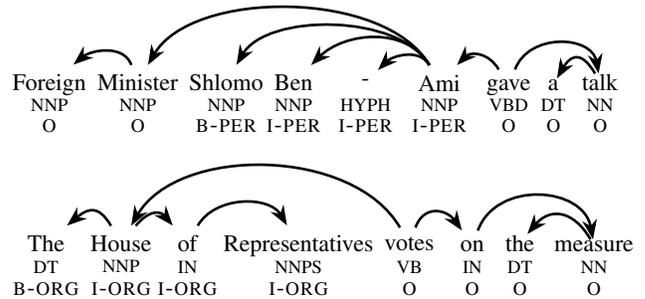


Figure 1: Two sentences annotated with both dependency and named entity information. The edges on top of words represent the dependencies and the labels with IOB encoding are the entity types.

Ling and Weld 2012; Cucchiarelli and Velardi 2001). Figure 1 shows two example phrases annotated with both dependency and named entity information. The local features are usually the head word and its part-of-speech tag at current position. For example, “*Shlomo*” with entity tag B-PER in the first sentence has two local dependency features, head word “*Ami*” and head tag “NNP”. However, such a simple treatment of dependency structures largely ignores the global structured information conveyed by the dependency trees, which can be potentially useful in building NER models.

One key observation we can make in Figure 1 is that named entities are often covered by a single or multiple consecutive dependency arcs. In the first example, the named entity “*Shlomo Ben - Ami*” of type PER (*person*) is completely covered by the single dependency arc from “*Ami*” to “*Shlomo*”. Similarly, the named entity “*The House of Representatives*” of type ORG (*organization*) in the second example is covered by multiple arcs which are adjacent to each other. Such information can potentially be the global features we can obtain from the dependency trees. This leads to the following questions: 1) can such global structured information conveyed by dependency trees be exploited for improved NER, and 2) if so, how to build new NER models where such information can be explicitly incorporated?

With these two questions, in this paper we perform some investigations on how to better utilize the structured infor-

mation conveyed by dependency trees for building novel models for improved named entity recognition. The model assumes the availability of dependency trees before performing NER, which can be obtained from a dependency parser or given as part of the input. Unlike existing approaches that only exploit dependency structures for encoding local features, the model is able to explicitly take into account the global structured information conveyed by dependency trees when performing learning and inference. We call our proposed NER model the *dependency-guided* model (DGM), and build it based on the conventional semi-Markov conditional random fields (semi-CRFs) (Sarawagi and Cohen 2004), a classic model used for information extraction.

Our main contributions can be summarized as follows:

- We present a novel model that is able to explicitly exploit the global structured information conveyed by dependency trees, showing that such information can be effectively integrated into the process of performing NER. To the best of our knowledge, this is the first work that exploits such information for NER.
- Theoretically, we show through average-case time complexity analysis that our model has the same time complexity as that of the linear-chain CRFs, and is better than that of the semi-Markov CRFs.
- Empirically, we demonstrate the benefits of our approach through extensive experiments on benchmark datasets. We show that the resulting model leads to NER results that are competitive with the baseline approach based on semi-Markov CRFs, while requiring significantly less running time.

Furthermore, although in this paper we focus on the task of using the structured information conveyed by dependency trees for improved NER, the underlying idea is general, which we believe can be applied to other tasks that involve building pipeline or joint models for structured prediction. Potentially if we can find the relationship between different NLP tasks, we can build an efficient model for a specific task while using information from other tasks without making the model more complex.

Related Work

Named entity recognition has a long history in the field of natural language processing. One standard approach to NER is to regard the problem as a sequence labeling problem, where each word is assigned a tag, indicating whether the word belongs to part of any named entity or appears outside of all entities. Previous approaches used sequence labeling models such as hidden Markov models (HMMs) (Zhou and Su 2002), maximum entropy Markov models (MEMMs) (McCallum, Freitag, and Pereira 2000), as well as linear-chain (Finkel, Grenager, and Manning 2005) and semi-Markov conditional random fields (CRFs/semi-CRFs) (Sarawagi and Cohen 2004). Muis and Lu (2016b) proposed a weak semi-CRFs model which has a lower complexity than the conventional semi-CRFs model while still having a higher complexity than the linear-chain CRFs model. Our model is proved to have the same time

complexity as linear-chain CRFs model in the average case. The quality of the CRFs model typically depends on the features that are used. While most research efforts exploited standard word-level features (Ratinov and Roth 2009), more sophisticated features can also be used. Ling and Weld (2012) showed that using syntactic-level features from dependency structures in a CRFs-based model can lead to improved NER performance. Such dependency structures were also used in the work by Liu, Huang, and Zhu (2010), where the authors utilized such structures for building a skip-chain variant of the original CRFs model. This shows that some simple structured information conveyed by dependency trees can be exploited for improved NER. In their skip-chain CRFs model, they simply added certain dependency arcs as additional dependencies in the graphical model, resulting in loopy structures. However, such a model did not explicitly explore the relation between entities and global structured information of the dependency trees. The authors also showed that such a model does not outperform a simpler approach that adds additional dependencies between similar words only on top of the original CRFs model. In this work, we also focus on utilizing dependency structures for improving NER. Unlike previous approaches, we focus on exploiting the global structured information conveyed by dependency trees to improve the NER process. Comparing with the semi-CRFs model, our model is not only able to perform competitively in terms of performance, but also more *efficient* in terms of running time.

There are also some existing works that focus on improving the efficiency of NER and other information extraction models. For example, Okanohara et al. (2006) used a separate naive Bayes classifier to filter some entities during training and inference in their semi-CRFs based model. While the filtering process was used to reduce the computational cost of the semi-CRFs model, the model still needs to enumerate all the possible chunks. Yang and Cardie (2012) extended the original semi-CRFs for extracting opinion expressions and used the constituency parse tree information to avoid constructing implausible segments. Lu and Roth (2015) proposed an efficient and scalable model using hypergraph which can handle overlapping entities. Muis and Lu (2016a) extended the hypergraph representation to recognize both contiguous and discontinuous entities.

Background

Before we introduce our models, we would like to have a brief discussion on the relevant background. Specifically, in this section, we review the two classic models that are commonly used for named entity recognition, namely the linear-chain conditional random fields and the semi-Markov conditional random fields models.

Linear-chain CRFs

Conditional random fields, or CRFs (Lafferty, McCallum, and Pereira 2001) is a popular model for structured prediction, which has been widely used in various natural language processing problems, including named entity recognition (McCallum and Li 2003) and semantic role labeling (Cohn and Blunsom 2005).

We focus our discussions on the linear-chain CRFs in this section. The probability of predicting a possible output sequence \mathbf{y} (e.g., a named entity label sequence in our case) given an input \mathbf{x} (e.g., a sentence) is defined as:

$$p(\mathbf{y}|\mathbf{x}) = \frac{\exp(\mathbf{w}^T \mathbf{f}(\mathbf{x}, \mathbf{y}))}{Z(\mathbf{x})} \quad (1)$$

where $\mathbf{f}(\mathbf{x}, \mathbf{y})$ is the feature vector defined over (\mathbf{x}, \mathbf{y}) tuple, \mathbf{w} is the weight vector consisting of parameters used for the model, and $Z(\mathbf{x})$ is the partition function used for normalization, which is defined as follows:

$$Z(\mathbf{x}) = \sum_{\mathbf{y}'} \exp(\mathbf{w}^T \mathbf{f}(\mathbf{x}, \mathbf{y}')) \quad (2)$$

In a (first-order) linear-chain CRF, the partition function for a input of length n can also be written as follows:

$$Z(\mathbf{x}) = \sum_{i=1}^n \sum_{y' \in T} \sum_{y \in T} \exp(\mathbf{w}^T \mathbf{f}(\mathbf{x}, y', y, i)) \quad (3)$$

where $\mathbf{f}(\mathbf{x}, y', y, i)$ is a local feature vector defined at the i -th position of the input sequence. T is the set of all output labels. The above function can be computed efficiently using dynamic programming. It can be seen that the time complexity of a linear-chain CRFs model is $\mathcal{O}(n|T|^2)$.

We aim to minimize the negative joint log-likelihood with L_2 regularization for our dataset:

$$\begin{aligned} \mathcal{L}(\mathbf{w}) = & \sum_i \log \sum_{\mathbf{y}'} \exp(\mathbf{w}^T \mathbf{f}(\mathbf{x}_i, \mathbf{y}')) \\ & - \sum_i \mathbf{w}^T \mathbf{f}(\mathbf{x}_i, \mathbf{y}_i) + \lambda \mathbf{w}^T \mathbf{w} \end{aligned} \quad (4)$$

where $(\mathbf{x}_i, \mathbf{y}_i)$ is the i -th training instance and λ is the L_2 regularization coefficient.

The objective function is convex and we can make use of the L-BFGS (Byrd et al. 1995) algorithm to optimize it. The partial derivative of \mathcal{L} with respect to each parameter w_k is:

$$\frac{\partial \mathcal{L}}{\partial w_k} = \sum_i \left(\mathbf{E}_{p(\mathbf{y}'|\mathbf{x}_i)} [f_k(\mathbf{x}_i, \mathbf{y}')] - f_k(\mathbf{x}_i, \mathbf{y}_i) \right) + 2\lambda w_k$$

Semi-Markov CRFs

The semi-Markov conditional random fields, or semi-CRFs (Sarawagi and Cohen 2004) is an extension of the standard linear-chain CRFs. Different from linear-chain CRFs, which makes use of simple first-order Markovian assumptions, the semi-CRFs assumes that the transitions between words inside a span (e.g., an entity consisting of multiple words) can be non-Markovian. Such an assumption allows more flexible non-Markovian features to be defined. The resulting model was shown to be more effective than its linear-chain counterpart in information extraction tasks.

The partition function in this setting becomes:

$$Z(\mathbf{x}) = \sum_{i=1}^n \sum_{l=1}^L \sum_{y' \in T} \sum_{y \in T} \exp(\mathbf{w}^T \mathbf{f}(\mathbf{x}, y', y, i-l, i)) \quad (5)$$

where $\mathbf{f}(\mathbf{x}, y', y, i-l, i)$ represents the local feature vector at position i with a span of size l . In this case, the time complexity becomes $\mathcal{O}(nL|T|^2)$. The value L is the maximal length of the spans the model can handle.

Consider the sentence “*Lee Ann Womack won Single of the Year award*”. The upper portion of Figure 2 shows all the possible structures that are considered by the partition function with $L = 4$. These structures essentially form a compact lattice representation showing all the possible combinations of entities (with length restriction) that can ever appear in the given sentence. Each orange and red edge is used to represent an entity of type PER and MISC respectively. The black lines are used to indicate words that are not part of any entity. The directed path highlighted in bold blue corresponds to the correct entity information associated with the phrase, where “*Lee Ann Womack*” is an entity of type PER, and “*Single of the Year*” is another entity of type MISC (*miscellaneous*).

Our Model

The primary assumption we would like to make is that the dependency trees of sentences are available when performing the named entity recognition task.

NER with Dependency Features

One approach to exploit information from dependency structures is to design local features based on such dependency structures and make use of conventional models such as semi-CRFs for performing NER. Previous research efforts have shown that such features extracted from dependency structures can be helpful when performing the entity recognition task (Sasano and Kurohashi 2008; Ling and Weld 2012; Cucchiarelli and Velardi 2001).

In Figure 2, we have already used a graphical illustration to show the possible combinations of entities that can ever appear in the given sentence. Each edge in the figure corresponds to one entity (or a single word that is outside of any entity – labeled with O). Features can be defined over such edges, where each feature is assigned a weight. Such features, together with their weights, can be used to score each possible path in the figure. When dependency trees are provided, one can define features involving some local dependency structures. For example, for the word “*Womack*”, one possible feature that can be defined over edges covering this word can be of the form “*Womack ← won*”, indicating there is an incoming arc from the word “*won*” to the current word “*Womack*”. However, we note that such features largely ignore the global structured information as presented by the dependency trees. Specifically, certain useful facts such as the word “*Lee*” is a child of the word “*Womack*” and at the same time a grandchild of the word “*won*” is not captured due to such a simple treatment of dependency structures.

Dependency-Guided NER

The main question we would like to ask is: how should we make good use of the structured information associated with the dependency trees to perform named entity recognition? Since we are essentially interested in NER only, would there

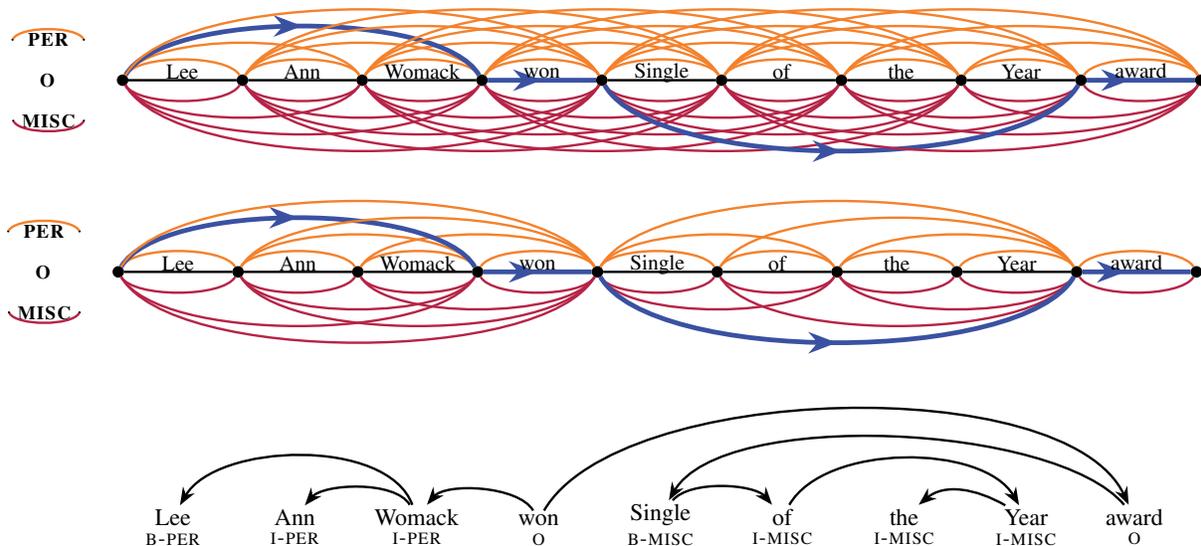


Figure 2: Illustrations of possible combinations of entities for the conventional semi-CRFs model (top) and our DGM model (middle), as well as the example sentence with its dependency structure (bottom).

be some more global structured information in the dependency trees that can be used to guide the NER process?

From the earlier two examples shown in Figure 1 as well as the example shown in Figure 2 we can observe that the named entities tend to be covered by a single or multiple adjacent dependency arcs. This is not a surprising finding as for most named entities which convey certain semantically meaningful information, it is expected that there exist certain internal structures – *i.e.*, dependencies amongst different words within the entities. Words inside each named entity typically do not have dependencies with words outside the entities, except for certain words such as head words which typically have incoming arcs from outside words.

This finding motivates us to exploit such global structured information as presented by dependency trees for performing NER. Following the above observation, we first introduce the following notion:

Definition 1 A *valid span* either consists of a single word, or is a word sequence that is covered by a chain of (undirected) arcs where no arc is covered by another.

For example, in Figure 2, the word sequence “Lee Ann Womack” is a valid span since there exists a single arc from “Womack” to “Lee”. The sequence “Ann Womack won” is also a valid span due to a chain of undirected arcs – one between “Ann” and “Womack” and the other between “Womack” and “won”. Similarly, the single word “Womack” is also a valid span given the above definition.

Based on the above definition, we can build a novel dependency-guided NER model based on the conventional semi-CRFs by restricting the space of all possible combinations of entities to those that strictly contain only valid spans. This leads to the following new partition function:

$$Z(\mathbf{x}) = \sum_{(i,j) \in \mathcal{S}_L(\mathbf{x})} \sum_{y' \in T} \sum_{y \in T} \exp(\mathbf{w}^T \mathbf{f}(\mathbf{x}, y', y, i, j)) \quad (6)$$

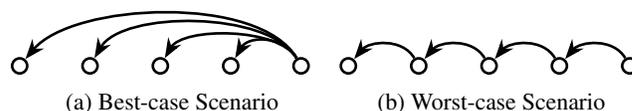


Figure 3: The best-case and worst-case scenarios of DGM.

where $\mathcal{S}(\mathbf{x})$ refers to the set of valid spans for a given input sentence \mathbf{x} , and $\mathcal{S}_L(\mathbf{x})$ refers to its subset that contains only those valid spans whose lengths are no longer than L .

We call the resulting model *dependency-guided model* (DGM). Figure 2 (middle) presents the graphical illustration of all possible paths (combination of entities) with $L = 4$ that our model considers. For example, since the word sequence “Ann Womack won Single” is not a valid span, the corresponding edges covering this word sequence is removed from the original lattice representation in Figure 2 (top). The edge covering the word sequence “of the Year” remains there as it is covered by the arc from “of” to “Year”, thus a valid span. As we can see, the amount of paths that we consider in the new model is substantially reduced.

Time Complexity

We can also analyze the time complexity required for this new model. We show example best-case and worst-case scenarios in Figure 3. For the former case there are $\mathcal{O}(n)$ valid spans. Thus the time complexity in the best case is $\mathcal{O}(n|T|^2)$, the same as the linear-chain CRFs model. For the latter, there are $\mathcal{O}(nL)$ valid spans, leading to the time complexity $\mathcal{O}(nL|T|^2)$ in the worst case, the same as the semi-Markov CRFs model.

Furthermore, we have the following theorem for the average-case time complexity:

Theorem 1 *The average-case time complexity of DGM is $\mathcal{O}(n|T|^2)$.*

Proof We provide two proofs to the above theorem. Below we show one proof and we move the second proof to the supplementary material (S.3.1). In this proof we assume the largest possible value for L , which is n (*i.e.*, we do not limit the length of valid spans). Since the time complexity of our model depends directly on the total number of edges in the DGM graph, which is the same as the number of valid spans, we count the total number of valid spans over all possible structures by defining a bijection between each valid span and a tree with $(n + 1)$ nodes. Here we consider only one entity type since when we have $|T|$ types, the number of edges can be simply calculated by multiplying $|T|^2$ to the number that we will calculate here.

Let $\mathcal{T}(n)$ be the set of all undirected trees with n nodes, let $\tau \in \mathcal{T}(n)$ be an undirected tree with n nodes, let $s_\tau^{u,v}$ be a valid span based on τ covering the words from position u to v , and let $\mathcal{S}(\tau) = \{s_\tau^{u,v} \mid 1 \leq u \leq v \leq n\}$ be the set of valid spans based on τ . Then the set of all valid spans over all possible undirected trees with n nodes is $\mathcal{S}_{\text{DGM}}(n) = \cup_{\tau \in \mathcal{T}(n)} \mathcal{S}(\tau)$. Note that since the definition of valid spans uses undirected arcs of the dependency tree, we use undirected trees to count the number of valid spans.

Now we define the bijection from $\mathcal{S}_{\text{DGM}}(n)$ to $\mathcal{T}(n + 1)$ to show that the two sets have the same number of elements. By definition of valid spans, for each $s_\tau^{u,v} \in \mathcal{S}_{\text{DGM}}(n)$, there is exactly one associated chain of edges in the dependency tree τ , which is the chain of edges starting at u and ending at v , possibly visiting some intermediary nodes in between.

Then, to define the bijection we map $s_\tau^{u,v}$ to a tree $\tau' \in \mathcal{T}(n + 1)$ which is constructed from τ as follows: the edges for the first n nodes coincide with τ , except that the edges involved in the chain of edges are removed, and we make the nodes involved in this chain of edges to be connected to the $(n + 1)$ -th node instead. The process creates a well-defined mapping since the resulting graph is still connected and has n edges over $(n + 1)$ nodes, which means it is a tree with $(n + 1)$ nodes.

Conversely, from a tree $\tau' \in \mathcal{T}(n + 1)$ we can recover the valid span and the undirected tree τ it is based on by removing the $(n + 1)$ -th node and the associated edges, then connecting the nodes that were previously connected to the $(n + 1)$ -th node in increasing order. The smallest and largest positions of the nodes define the starting and ending position of the span, and the resulting tree is the undirected tree τ .

For illustration, Figure 4a shows an example of a valid span $s_\tau^{1,4}$ where $\tau \in \mathcal{T}(4)$ and the nodes and edges involved in defining the valid span are highlighted in red. Figure 4b shows the associated tree $\tau' \in \mathcal{T}(5)$ where the removed edges are put in dashed lines and the new edges in bold.

Therefore the mapping is bijective, and thus the number of total valid spans in $\mathcal{S}_{\text{DGM}}(n)$ is equal to the number of trees in $\mathcal{T}(n + 1)$. It is a known result that the number of possible trees with n nodes is n^{n-2} (Aigner and Ziegler 2010, Chapter 30), which means the total number of valid spans is

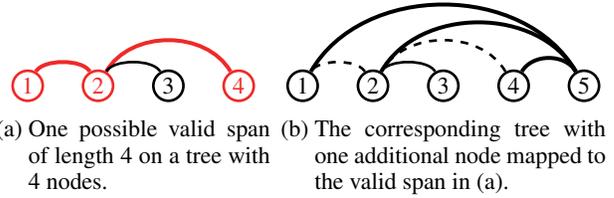


Figure 4: An illustration of the bijection between a valid span $s_\tau^{u,v}$ and a tree $\tau' \in \mathcal{T}(n + 1)$.

$(n + 1)^{n-1}$. Thus, the average number of valid spans is:

$$\frac{\mathcal{S}_{\text{DGM}}(n)}{|\mathcal{T}(n)|} = \frac{(n + 1)^{n-1}}{n^{n-2}} = n \left(1 + \frac{1}{n}\right)^{n-1} \leq n \cdot e \quad (7)$$

which is linear in terms of n , just like the linear-chain CRFs. Here e is the Euler number (≈ 2.718).

This shows that the average-case time complexity of our model is $\mathcal{O}(n|T|^2)$. ■

Now, this is a remarkable result, showing that the complexity of our DGM model in its average case is still linear in the sentence length n even if we set L to its maximal value n . This is also true empirically, which we show in the supplementary material (S.3.2). So, while our model retains the ability to capture non-Markovian features like semi-Markov CRFs, it is more scalable and can be used to extract longer entities.

Besides the standard DGM model defined above, we also consider a variant, where we restrict the chain (of arcs) to be of length 1 (*i.e.*, single arc) only. We call the resulting model DGM-S, where S stands for *single arc*. Such a simplified model will lead to an even simpler lattice representation, resulting in even less running time. However it may also lead to lower performance as such a variant might not be able to capture certain named entities. In one of the later sections, we will conduct experiments on such a model and compare its performance with other models.

Experimental Setup

For experiments, we followed (Finkel and Manning 2009) and used the Broadcast News section of the OntoNotes dataset. Instead of using its earlier 2.0 release, we used the final release – release 5.0 of the dataset, which is available for download¹. The dataset includes the following 6 subsections: ABC, CNN, MNB, NBC, PRI and VOA. Moreover, the current OntoNotes 5.0 release also includes some English P2.5 data, which consists of corpus translated from Chinese and Arabic.² Following (Finkel and Man-

¹<https://catalog.ldc.upenn.edu/LDC2013T19/>

²The OntoNotes 5.0 dataset also contains the earlier release from OntoNotes 1.0 to OntoNotes 4.0. However, we found the number of sentences of the Broadcast News in the current OntoNotes 5.0 dataset does not match the number reported in (Finkel and Manning 2009; 2010), which was based on OntoNotes 2.0. Furthermore, they removed some instances which are inconsistent with their model. We thus cannot conduct experiments to compare against the results reported in their work.

	# Sent.	# Entities		
		ALL	DGM-S	DGM
Train	9,996	18,855	17,584 (93.3%)	18,803 (99.7%)
Test	3,339	5,742	5,309 (92.5%)	5,720 (99.6%)

Table 1: Dataset statistics.

ning 2009), we split the first 75% of the data for training and performed evaluations on the remaining 25%. We set $L = 8$, which can cover all entities in our dataset, and developed the L_2 coefficient using cross-validation (see supplementary material S.1 for details). Following previous works on dependency parsing (Chen and Manning 2014), we preprocessed the dataset using the Stanford CoreNLP³ to convert the constituency trees to basic Stanford dependency (De Marneffe et al. 2006) trees. In our NER experiments, in addition to using the given dependency structures converted from the constituency trees, we also trained a popular transition-based parser, MaltParser⁴ (Nivre, Hall, and Nilsson 2006), using the training set and then used this parser to predict the dependency trees on the test set to be used in our model.

We also looked at the SemEval-2010 Task 1 OntoNotes English corpus⁵, which contains sentences with both dependency and named entity information. Although this dataset is a subset of the OntoNotes dataset, it comes with manually annotated dependency structures. We conducted experiments on this dataset and reported the results in the supplementary material (S.2). The results on this dataset are consistent with the findings reported in this paper.

Features

In order to make comparisons, we implemented a linear-chain CRFs model as well as a semi-Markov CRFs model to serve as baselines. The features used in this paper are basic features which are commonly used in linear-chain CRFs and semi-CRFs. For the linear-chain CRFs, we consider the current word/POS tag, the previous word/POS tag, the current word shape, the previous word shape, prefix/suffix of length up to 3, as well as transition features. For word shape features, we followed (Finkel et al. 2005) to create them. For the semi-CRFs model, we have the following features for each segment: the word/POS tag/word shape before and after the segment, indexed words/POS tags/word shapes in current segment, surface form of the segment, segment length, segment prefix/suffix of length up to 3, the start/end word/tags of current segment, and the transition features.

Inspired by Ling and Weld (2012), we applied the following dependency features for all models: (1) current word in current position/segment and its head word (and its dependency label); (2) current POS tag in current position/segment and its head tag (and dependency label). More details of features can be found in the supplementary material (S.4).

³<http://stanfordnlp.github.io/CoreNLP/>

⁴<http://maltparser.org/>

⁵<https://catalog.ldc.upenn.edu/LDC2011T01/>

Data Statistics

There are in total 18 well-defined named entity types in the OntoNotes 5.0 dataset. Since majority of the entities are from the following three types: PER (*person*), ORG (*organization*) and GPE (*geo-political entities*), following (Finkel and Manning 2009), we merge all the other entity types into one general type, MISC (*miscellaneous*). Table 1 shows the statistics of total number of sentences and entities.

To show the relationships between the named entities and dependency structures, we present the number and percentage of entities that can be handled by our DGM-S model and DGM model respectively. The entities that can be handled by DGM-S should be covered by a single arc as indicated in our model section. As for DGM, the entity spans should be *valid* as in definition 1. Overall, we can see that 93.3% and 92.5% of the entities can be handled by the DGM-S model in training set and test set, respectively. These two numbers for DGM are much higher – 99.7% and 99.6%. With the predicted dependency structures in test set, 91.7% of the entities can be handled by the DGM-S model, while for DGM it is 97.4%.

We provide more detailed statistics for each of the 7 subsections in the supplementary material (S.1). These numbers confirm that it is indeed the case that most named entities do form *valid spans*, even when using predicted dependency trees, and that such global structured information of dependency trees can be exploited for NER.

Results and Discussions

NER Performance

Following previous work (Finkel and Manning 2009), we report standard F-score in this section (detailed results with precision and recall can be found in the supplementary material S.5). Table 2 shows the results of all models when dependency features are exploited. Specifically, it shows results when the given and predicted dependency trees are considered, respectively. Overall, the semi-Markov CRFs, DGM-S and DGM models perform better than the linear-chain CRFs model. Our DGM model obtains an overall F-score at 80.5% and outperforms the semi-CRFs model significantly ($p < 0.001$ with bootstrap resampling (Koehn 2004)). For individual subsection, our DGM also performs the best. On 2 out of the 7 subsections, our model’s improvements over the baseline semi-CRFs model are significant ($p < 0.001$ with bootstrap resampling). For some other subsections like ABC, CNN, MNB, NBC and PRI, DGM has higher F-score than the semi-CRFs model but the improvements are not statistically significant. The results show that comparing with semi-CRFs, our DGM model, which has a lower average-case time complexity, still maintains a competitive performance. Such results confirm that the global structured information conveyed by dependency trees are useful and can be exploited by our DGM model.

The performance of DGM-S is worse than that of semi-CRFs and DGM in general since there are still many named entities that can not be handled by such a simplified model (see Table 1). This model has the same time complexity as the linear-chain CRFs, but performs better than linear-chain

Dependency	Model	ABC	CNN	MNB	NBC	P2.5	PRI	VOA	Overall
Given	Linear-chain CRFs	70.2	75.9	75.7	65.9	70.8	83.2	84.6	77.8
	Semi-Markov CRFs	71.9	78.2	74.7	69.4	73.5	85.1	85.4	79.6
	DGM-S	71.4	77.0	73.4	68.4	72.8	85.1	85.2	79.0
	DGM	72.3	78.6	76.3	69.7	75.5	85.5	86.8	80.5
Predicted	Linear-chain CRFs	68.4	75.4	74.4	66.3	70.8	83.3	83.7	77.3
	Semi-Markov CRFs	71.6	78.0	73.5	71.5	73.7	84.6	85.3	79.5
	DGM-S	70.6	76.4	73.4	68.7	71.3	83.9	84.4	78.2
	DGM	71.9	77.6	75.4	71.4	73.9	84.2	85.1	79.4

Table 2: NER results for all models, when given and predicted dependency trees are used and dependency features are used. Best values and the values which are not significantly different in 95% confidence interval are put in bold.

Dependency	Model	ABC	CNN	MNB	NBC	P2.5	PRI	VOA	Overall
Given	Linear-chain CRFs	66.5	74.1	74.9	65.4	70.8	82.9	82.3	76.3
	Semi-Markov CRFs	72.3	76.6	75.0	69.3	73.7	84.1	83.3	78.5
	DGM-S	69.4	76.1	73.4	68.0	72.5	85.2	85.1	78.6
	DGM	72.7	77.2	75.8	68.5	76.8	86.2	85.5	79.9
Predicted	Linear-chain CRFs	66.5	74.1	74.9	65.4	70.8	82.9	82.3	76.3
	Semi-Markov CRFs	72.3	76.6	75.0	69.3	73.7	84.1	83.3	78.5
	DGM-S	69.1	75.6	73.8	67.2	72.0	84.5	84.2	78.0
	DGM	71.3	76.2	75.9	68.8	74.6	85.1	84.3	78.8

Table 3: NER results for all models, when given and predicted dependency trees are used but dependency features are not used. Best values and the values which are not significantly different in 95% confidence interval are put in bold.

CRFs, largely due to the fact that certain structured information of dependency trees are exploited in DGM-S. We note that such a simplified model obtains similar results as semi-CRFs for the two larger subsections, PRI and VOA. This is largely due to the fact that a larger percentage of the entities in these two subsections can be handled by DGM-S (see supplementary material S.1 for more details). It is noted that the performance of both semi-CRFs and DGM degrades when the predicted dependency trees are used instead of the given. The drop in F-score for DGM is more significant as compared to the semi-CRFs. Nevertheless, their results remain comparable. Such experiments show the importance of considering high quality dependency structures for performing guided NER in our model.

To understand the usefulness of the global structured information of dependency trees better, we conducted further experiments by excluding dependency features from our models. Such results are shown in Table 3. Our DGM model consistently performs competitively with the semi-CRFs model. The only exception occurs when the ABC subsection is considered and the predicted dependency trees are used ($p=0.067$). In general, we can see that when dependency features are excluded, the overall F-score for all models drop substantially. However, we note that for semi-CRFs, the drop in F-score is 1.1% for given dependency trees, and is 1.0% for predicted trees, whereas for DGM, the drops with given and predicted trees are both 0.6%. Overall, such results largely show that our proposed model is able to effectively make use of the global structured information conveyed by dependency trees for NER.

We have also conducted experiments on the widely-used

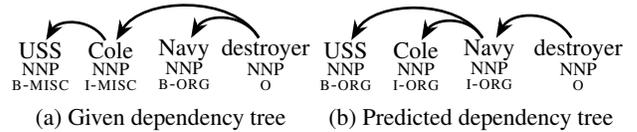


Figure 5: The effect of different dependency parses on the output of DGM. These are taken out from a part of a sentence. The NER result in (a) is correct, while (b) is not.

NER dataset, CoNLL-2003, using the Stanford dependency parser⁶ to generate the dependency trees. Using the same feature set that we describe in this paper, our models do not achieve the state-of-the-art results on this dataset. However, they still perform comparably with the semi-CRFs model, while requiring much lesser running time. Note that since our goal in this paper is to investigate the usefulness of incorporating the dependency structure information for NER, we did not attempt to tune the feature set to get the best result on a specific dataset. Also it is worth remarking that we still obtain a relatively good performance for our DGM model although the dependency parser is not trained within the dataset itself and that a correct dependency structure information is crucial for the DGM model.

Error Analysis

We provide a further analysis of how the dependency parsing performance affects NER based on Table 3. Because our model uses the dependency structure information directly in-

⁶<http://nlp.stanford.edu/software/nndep.shtml>

stead of using them as features, we can analyze the influence of dependency structures on NER more clearly. Specifically, we focus on how the dependency parsing results affect the prediction of our DGM model.

A typical error made by our model taken from the results is shown in Figure 5 where the dependency tree in Figure 5a is given by the conversion of constituent parse tree and the other one in Figure 5b is predicted from the MaltParser. The NER result on the left is correct while the one on the right is incorrect. Based on the predicted dependency structure in Figure 5b, there is no way to predict an entity type for “USS Cole” since this is not a valid span in DGM model. Furthermore, DGM can actually recognize “Navy” as an ORG entity even though the predicted dependency is incorrect. But the model considers “USS Cole Navy” as an entity due to the interference of other entity features (e.g., NNP tag and Capitalized pattern) that “USS Cole” has. While with the given dependency tree, DGM considers “USS Cole” as a valid span and correctly recognizes it as a MISC entity.

Speed Analysis

From Figure 2 we can see that the running time required for each model depends on the number of edges that the model considers. We thus empirically calculated the average number of edges per word each model considers based on our training data. We found that the average number of edges involved in each token is 132 for the semi-CRFs model, while this number becomes 40 and 61 for DGM-S and DGM respectively. A lower number of edges indicates less possible structures to consider, and a reduced running time. See more detailed information in the supplementary material (S.3.2).

The results on training time per iteration (inference time) for all 7 subsections are shown in Figure 6. From the figure we can see that the linear-chain CRFs model empirically runs the fastest. The simple DGM-S model performs comparably with linear-chain CRFs. The semi-CRFs model requires substantially longer time due to the additional factor L (set to 8 in our experiments) in its time complexity. In contrast, our model DGM requires only 47% of the time needed for semi-CRFs for each training iteration, and requires 37% more time than the DGM-S model.

Conclusion

In this paper, we proposed a novel efficient dependency-guided model for named entity recognition. Motivated by the fact that named entities are typically covered by dependency arcs, presenting internal structures, we built a model that is able to explicitly exploit global structured information conveyed by dependency trees for NER. We showed that the model theoretically is better than the semi-Markov CRFs model in terms of time complexity. Experiments show that our model performs competitively with the semi-Markov CRFs model, even though it requires substantially less running time. Future directions include further investigations on the structural relations between dependency trees and named entities, and working towards building integrated models that perform joint prediction of both structures.

We make our code and system available for download at <http://statnlp.org/research/ie/>.

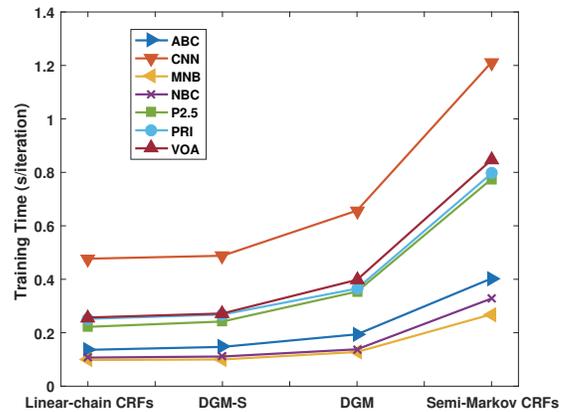


Figure 6: Training time per iteration of all the models.

Acknowledgments

We would like to thank the anonymous reviewers for their constructive and helpful comments. This work is supported by MOE Tier 1 grant SUTDT12015008.

References

- Aigner, M., and Ziegler, G. M. 2010. *Proofs from THE BOOK*. Springer Berlin Heidelberg, 4th edition.
- Byrd, R. H.; Lu, P.; Nocedal, J.; and Zhu, C. 1995. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing* 16(5):1190–1208.
- Chen, D., and Manning, C. D. 2014. A fast and accurate dependency parser using neural networks. In *Proceedings of EMNLP*, 740–750.
- Cohn, T., and Blunsom, P. 2005. Semantic role labelling with tree conditional random fields. In *Proceedings of CoNLL*, 169–172.
- Cucchiarelli, A., and Velardi, P. 2001. Unsupervised named entity recognition using syntactic and semantic contextual evidence. *Computational Linguistics* 27(1):123–131.
- De Marneffe, M.-C.; MacCartney, B.; Manning, C. D.; et al. 2006. Generating typed dependency parses from phrase structure parses. In *Proceedings of LREC*, 449–454.
- Finkel, J. R., and Manning, C. D. 2009. Joint parsing and named entity recognition. In *Proceedings of HLT-NAACL*, 326–334.
- Finkel, J. R., and Manning, C. D. 2010. Hierarchical joint learning: Improving joint parsing and named entity recognition with non-jointly labeled data. In *Proceedings of ACL*, 720–728.
- Finkel, J.; Dingare, S.; Manning, C. D.; Nissim, M.; Alex, B.; and Grover, C. 2005. Exploring the boundaries: gene and protein identification in biomedical text. *BMC bioinformatics* 6(1):1.
- Finkel, J. R.; Grenager, T.; and Manning, C. 2005. Incorporating non-local information into information extraction

- systems by gibbs sampling. In *Proceedings of ACL*, 363–370.
- Koehn, P. 2004. Statistical significance tests for machine translation evaluation. In *Proceedings of EMNLP*, 388–395.
- Krishnamurthy, J., and Mitchell, T. M. 2015. Learning a compositional semantics for freebase with an open predicate vocabulary. *Transactions of the Association for Computational Linguistics* 3:257–270.
- Lafferty, J.; McCallum, A.; and Pereira, F. C. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of ICML*, 282–289.
- Lao, N., and Cohen, W. W. 2010. Relational retrieval using a combination of path-constrained random walks. *Machine learning* 81(1):53–67.
- Lee, H.; Recasens, M.; Chang, A.; Surdeanu, M.; and Jurafsky, D. 2012. Joint entity and event coreference resolution across documents. In *Proceedings of Joint Conference on EMNLP-CoNLL*, 489–500.
- Liang, P.; Jordan, M. I.; and Klein, D. 2013. Learning dependency-based compositional semantics. *Computational Linguistics* 39(2):389–446.
- Ling, X., and Weld, D. S. 2012. Fine-grained entity recognition. In *Proceedings of AAAI*, 94–100.
- Liu, J.; Huang, M.; and Zhu, X. 2010. Recognizing biomedical named entities using skip-chain conditional random fields. In *Proceedings of the Workshop on BioNLP*, 10–18.
- Lu, W., and Roth, D. 2015. Joint mention extraction and classification with mention hypergraphs. In *Proceedings of EMNLP*, 857–867.
- McCallum, A., and Li, W. 2003. Early results for named entity recognition with conditional random fields, feature induction and web-enhanced lexicons. In *Proceedings of HLT-NAACL*, 188–191.
- McCallum, A.; Freitag, D.; and Pereira, F. C. 2000. Maximum entropy markov models for information extraction and segmentation. In *Proceedings of ICML*, 591–598.
- Muis, A. O., and Lu, W. 2016a. Learning to recognize discontinuous entities. In *Proceedings of EMNLP*.
- Muis, A. O., and Lu, W. 2016b. Weak semi-markov crfs for noun phrase chunking in informal text. In *Proceedings of NAACL*, 714–719.
- Nivre, J.; Hall, J.; and Nilsson, J. 2006. Maltparser: A data-driven parser-generator for dependency parsing. In *Proceedings of LREC*, 2216–2219.
- Okanohara, D.; Miyao, Y.; Tsuruoka, Y.; and Tsujii, J. 2006. Improving the scalability of semi-markov conditional random fields for named entity recognition. In *Proceedings of COLING-ACL*, 465–472.
- Poon, H., and Domingos, P. 2009. Unsupervised semantic parsing. In *Proceedings of EMNLP*, 1–10.
- Ratinov, L., and Roth, D. 2009. Design challenges and misconceptions in named entity recognition. In *Proceedings of CoNLL*, 147–155.
- Sarawagi, S., and Cohen, W. W. 2004. Semi-markov conditional random fields for information extraction. In *Proceedings of NIPS*, 1185–1192.
- Sasano, R., and Kurohashi, S. 2008. Japanese named entity recognition using structural natural language processing. In *Proceedings of IJCNLP*, 607–612.
- Yang, B., and Cardie, C. 2012. Extracting opinion expressions with semi-markov conditional random fields. In *Proceedings of Joint Conference on EMNLP-CoNLL*, 1335–1345.
- Zhou, G., and Su, J. 2002. Named entity recognition using an hmm-based chunk tagger. In *Proceedings of ACL*, 473–480.