

Neural Relation Extraction within and across Sentence Boundaries

Pankaj Gupta,^{1,2} Subburam Rajaram,¹ Hinrich Schütze,² Thomas Runkler¹

¹Corporate Technology, Machine-Intelligence (MIC-DE), Siemens AG Munich, Germany

²CIS, University of Munich (LMU) Munich, Germany

{pankaj.gupta, subburam.rajaram}@siemens.com | pankaj.gupta@campus.lmu.de

Abstract

Past work in relation extraction mostly focuses on binary relation between entity pairs *within single sentence*. Recently, the NLP community has gained interest in relation extraction in entity pairs *spanning multiple sentences*. In this paper, we propose a novel architecture for this task: inter-sentential dependency-based neural networks (iDepNN). iDepNN models the shortest and augmented dependency paths via recurrent and recursive neural networks to extract relationships within (intra-) and across (inter-) sentence boundaries. Compared to SVM and neural network baselines, iDepNN is more robust to false positives in relationships spanning sentences. We evaluate our models on four datasets from newswire (MUC6) and medical (BioNLP shared task) domains that achieve state-of-the-art performance and show a better balance in precision and recall for inter-sentential relationships. We perform better than 11 teams participating in the BioNLP shared task 2016 and achieve a gain of 5.2% (0.587 vs 0.558) in F_1 over the winning team. We also release the cross-sentence annotations for MUC6.

Introduction

The task of relation extraction (RE) aims to identify semantic relationship between a pair of nominals or entities e_1 and e_2 in a given sentence S . Due to a rapid growth in information, it plays a vital role in knowledge extraction from unstructured texts and serves as an intermediate step in a variety of NLP applications in newswire, web and high-valued biomedicine (Bahcall 2015) domains. Consequently, there has been increasing interest in relation extraction, particularly in augmenting existing knowledge bases.

Progress in relation extraction is exciting; however most prior work (Zhang et al. 2006; Kambhatla 2004; Vu et al. 2016a; Gupta, Schütze, and Andrassy 2016) is limited to single sentences, i.e., *intra-sentential* relationships, and ignores relations in entity pairs spanning sentence boundaries, i.e., *inter-sentential*. Thus, there is a need to move beyond single sentences and devise methods to extract relationships spanning sentences. For instance, consider the sentences:

Paul Allen has started a company and named [Vern Raburn] _{e_1} its President. The company, to be called [Paul Allen Group] _{e_2} will be based in Bellevue, Washington.

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

The two sentences together convey the fact that the entity e_1 is associated with e_2 , which cannot be inferred from either sentence alone. The missed relations impact the system performance, leading to poor recall. But precision is equally important; e.g., in high-valued biomedicine domain, significant inter-sentential relationships must be extracted, especially in medicine that aims toward accurate diagnostic testing and precise treatment, and extraction errors can have severe negative consequences. In this work, we present a neural network (NN) based approach to precisely extract relationships within and across sentence boundaries, and show a better balance in precision and recall with an improved F_1 .

Previous work on cross-sentence relation extraction used coreferences to access entities that occur in a different sentence (Gerber and Chai 2010; Yoshikawa et al. 2011) without modeling inter-sentential relational patterns. Swampillai and Stevenson (2011) described a SVM-based approach to both intra- and inter-sentential relations. Recently, Quirk and Poon (2016) applied distant supervision to cross-sentence relation extraction of entities using binary logistic regression (non-neural network based) classifier and Peng et al. (2017) applied sophisticated graph long short-term memory networks to cross-sentence n-ary relation extraction. However, it still remains challenging due to the need for coreference resolution, noisy text between the entity pairs spanning multiple sentences and lack of labeled corpora.

Bunescu and Mooney (2005), Nguyen, Matsuo, and Ishizuka (2007) and Mintz et al. (2009) have shown that the shortest dependency path (SDP) between two entities in a dependency graph and the dependency subtrees are the most useful dependency features in relation classification. Further, Liu et al. (2015) developed these ideas using Recursive Neural Networks (*RecNNs*, Socher et al. (2014)) and combined the two components in a precise structure called Augmented Dependency Path (ADP), where each word on a SDP is attached to a dependency subtree; however, limited to single sentences. In this paper, we aspire from these methods to extend shortest dependency path across sentence boundary and effectively combine it with dependency subtrees in NNs that can capture semantic representation of the structure and boost relation extraction spanning sentences.

The *contributions* are: **(1)** Introduce a novel dependency-based neural architecture, named as inter-sentential Dependency-based Neural Network (iDepNN) to ex-

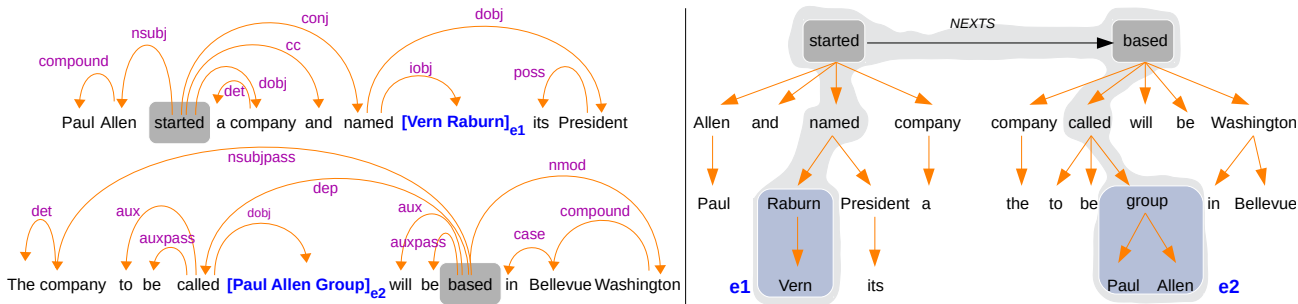


Figure 1: Left: Sentences and their dependency graphs. Right: Inter-sentential Shortest Dependency Path (iSDP) across sentence boundary. Connection between the roots of adjacent sentences by *NEXTS*.

tract relations within and across sentence boundaries by modeling shortest and augmented dependency paths in a combined structure of bidirectional RNNs (biRNNs) and RecNNs. **(2)** Evaluate different linguistic features on four datasets from newswire and medical domains, and report an improved performance in relations spanning sentence boundary. We show amplified precision due to robustness towards false positives, and a better balance in precision and recall. We perform better than 11 teams participating in the BioNLP shared task 2016 and achieve a gain of 5.2% (0.587 vs 0.558) in F_1 over the winning team. **(3)** Release relation annotations for the MUC6 dataset for intra- and inter-sentential relationships. *Code, data and supplementary* are available at <https://github.com/pgcool/Cross-sentence-Relation-Extraction-iDepNN>.

Methodology

Inter-sentential Dependency-Based Neural Networks (iDepNN)

Dependency-based neural networks (DepNN) (Bunescu and Mooney 2005; Liu et al. 2015) have been investigated for relation extraction between entity pairs limited to single sentences, using the dependency information to explore the semantic connection between two entities. In this work, we introduce *iDepNN*, the inter-sentential Dependency-based Neural Network, an NN that models relationships between entity pairs spanning sentences, i.e., inter-sentential within a document. We refer to the iDepNN that only models the shortest dependency path (SDP) spanning sentence boundary as *iDepNN-SDP* and to the iDepNN that models augmented dependency paths (ADPs) as *iDepNN-ADP*; see below. biRNNs (bidirectional RNNs, Schuster and Paliwal (1997)) and RecNNs (recursive NNs, Socher et al. (2012)) are the backbone of iDepNN.

Modeling Inter-sentential Shortest Dependency Path (iDepNN-SDP): We compute the inter-sentential Shortest Dependency Path (iSDP) between entities spanning sentence boundaries for a relation. To do so, we obtain the dependency parse tree for each sentence using the Stanford-CoreNLP dependency parser (Manning et al. 2014). We then use NetworkX (Hagberg, Swart, and S Chult 2008) to represent each token as a node and the dependency relation as a link between the nodes. In the case of multiple sentences,

the root node of the parse tree of a sentence is connected to the root of the subsequent tree, leading to the shortest path from one entity to another across sentences.

Figure 1 (Left) shows dependency graphs for the example sentences where the two entities *e1* and *e2* appear in nearby sentences and exhibit a relationship. Figure 1 (Right) illustrates that the dependency trees of the two adjacent sentences and their roots are connected by *NEXTS* to form an *iSDP*, an inter-Sentential Dependency Path, (highlighted in gray) between the two entities. The shortest path spanning sentence boundary is seen as a sequence of words between two entities. Figure 2 shows how a biRNN (Schuster and Paliwal 1997; Vu et al. 2016b) uses iSDP to detect relation between *e1* and *e2*, positioned one sentence apart.

Modeling Inter-sentential Dependency Subtrees: To effectively represent words on the shortest dependency path within and across sentence boundary, we model dependency subtrees assuming that each word w can be seen as the word itself and its children on the dependency subtree. The notion of representing words using subtree vectors within the dependency neural network (DepNN) is similar to (Liu et al. 2015); however, our proposed structures are based on iSDPs and ADPs that span sentences.

To represent each word w on the subtree, its word embedding vector $\mathbf{x}_w \in \mathcal{R}^d$ and subtree representation $\mathbf{c}_w \in \mathcal{R}^{d'}$ are concatenated to form its final representation $\mathbf{p}_w \in \mathcal{R}^{d+d'}$. We use 200-dimensional pretrained GloVe embeddings (Pennington, Socher, and Manning 2014). The subtree representation of a word is computed through recursive transformations of the representations of its children words. A RecNN is used to construct subtree embedding \mathbf{c}_w , traversing bottom-up from its leaf words to the root for entities spanning sentence boundaries, as shown in Figure 2. For a word which is a leaf node, i.e., it does not have a subtree, we set its subtree representation as \mathbf{c}_{LEAF} . Figure 2 illustrates how subtree-based word representations are constructed via iSDP.

Each word is associated with a dependency relation r , e.g., $r = \text{dobj}$, during the bottom-up construction of the subtree. For each r , a transformation matrix $\mathbf{W}_r \in \mathcal{R}^{d' \times (d+d')}$ is learned. The subtree embedding is computed as:

$$\mathbf{c}_w = f\left(\sum_{q \in \text{Children}(w)} \mathbf{W}_{R(w,q)} \cdot \mathbf{p}_q + \mathbf{b}\right) \text{ and } \mathbf{p}_q = [\mathbf{x}_q, \mathbf{c}_q]$$

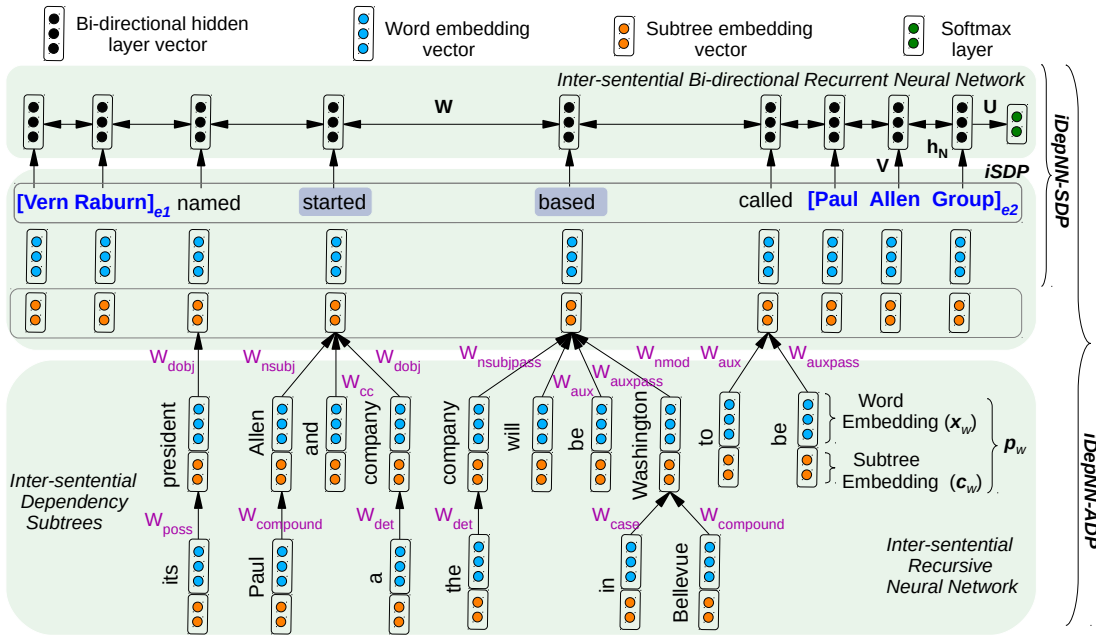


Figure 2: Inter-sentential Dependency-based Neural Network variants: iDepNN-SDP and iDepNN-ADP

where $R_{(w,q)}$ is the dependency relation between word w and its child word q and $\mathbf{b} \in \mathcal{R}^d$ is a bias. This process continues recursively up to the root word such as the word “named” on the iSDP in the figure.

Modeling Inter-sentential Augmented Dependency Path (iDepNN-ADP): Following Liu et al. (2015), we combine the two components: iSDP and *dependency subtrees* spanning sentence boundaries to form a combined structure which we name as inter-sentential Augmented Dependency Path (*iDepNN-ADP*). As shown in Figure 2, each word on iSDP is attached to its subtree representation \mathbf{c}_w . An attached subtree enriches each word on the iSDP with additional information about how this word functions in specific sentence to form a more precise structure for classifying relationships within and across sentences.

To capture the semantic representation of *iDepNN-ADP*, we first adopt a RecNN to model the dependency subtrees for each word on the iSDP. Then, we design a biRNN to obtain salient semantic features on the iSDP. The overall structure of *iDepNN-ADP* (Figure 2) is built upon the combination of recursive and recurrent NNs spanning sentences.

Learning: We develop a biRNN over the two structures: iDepNN-SDP and iDepNN-ADP, and pass the last hidden vector \mathbf{h}_N (in the iSDP word sequence, Figure 2) to a softmax layer whose output is the probability distribution \mathbf{y} over relation labels R , as $\mathbf{y} = \text{softmax}(\mathbf{U} \cdot \mathbf{h}_N + \mathbf{b}_y)$ where $\mathbf{U} \in \mathcal{R}^{R \times H}$ is the weight matrix connecting hidden vector of dimension H to output of dimension R and $\mathbf{b}_y \in \mathcal{R}^R$ is the bias. \mathbf{h}_N is the last hidden vector of the biRNN.

To compute semantic representation \mathbf{h}_w for each word w on the iSDP, we adopt the *Connectionist biRNN* (Vu et al. 2016a) that combines the forward and backward pass by adding their hidden layers (\mathbf{h}_{f_t} and \mathbf{h}_{b_t}) at each time step

t and also adds a weighted connection to the previous combined hidden layer \mathbf{h}_{t-1} to include all intermediate hidden layers into the final decision.

$$\begin{aligned} \mathbf{h}_{f_t} &= f(\mathbf{V} \cdot \mathbf{i}_t + \mathbf{W} \cdot \mathbf{h}_{f_{t-1}}) \\ \mathbf{h}_{b_t} &= f(\mathbf{V} \cdot \mathbf{i}_{N-t+1} + \mathbf{W} \cdot \mathbf{h}_{b_{t+1}}) \\ \mathbf{h}_t &= f(\mathbf{h}_{f_t} + \mathbf{h}_{b_t} + \mathbf{W} \cdot \mathbf{h}_{t-1}) \end{aligned}$$

where $\mathbf{V} \in \mathcal{R}^{H \times |i|}$, N is the total number of words on iSDP and \mathbf{i}_t the input vector at t , defined by:

$$\text{iDepNN-SDP} : \mathbf{i}_t = [\mathbf{x}_t, \mathbf{L}_t] \quad \text{iDepNN-ADP} : \mathbf{i}_t = [\mathbf{p}_t, \mathbf{L}_t]$$

where \mathbf{L}_t are lexical level features (e.g., part-of-speech tag, position indicators, entity types) for each word at t . Observe, in order to minimize the number of parameters, we share the same weight matrix \mathbf{W} in three parts: forward pass, backward pass and combination of both. The optimization objective is to minimize the cross-entropy error between the ground-truth label and softmax output. The parameters are learned using backpropagation (Werbos 1990).

Key Features: The features focus on characteristics of the full sentence, dependency path or individual entities. The various features used in our experiments are: (1) *Position-Indicator* (PI): A one-hot vector for SVM which indicates the position of the entity in the vocabulary. Four additional words ($\langle e_1 \rangle$, $\langle /e_1 \rangle$, $\langle e_2 \rangle$, $\langle /e_2 \rangle$) to mark start and end of entity mentions e_1 and e_2 , used in NNs. See details about PI in Gupta (2015). (2) *Entity Types* (ET): A one-hot vector to represent the entity type in SVM and embedding vectors in NNs. (3) *Part-of-speech* (POS): A bag-of-words (BoW) in SVM and embedding vector for each POS type in NNs. (4) *Dependency*: In SVM, the specific edge types in the dependency path are captured with a BoW vector, similar to

Relation	Intra	Inter	Relation	Intra	Inter
BioNLP ST 2011 (<i>Medical</i>)			BioNLP ST 2013 (<i>Medical</i>)		
PartOf	99	103	PartOf	104	83
Localization	261	732	Localization	246	677
Total	360	835 (70%)	Total	350	760 (69%)
BioNLP ST 2016 (<i>Medical</i>)			MUC6 (<i>News</i>)		
Lives_In	363	135	Per-Org	245	112
			Per-Post	407	66
			Org-Post	268	113
Total	363	135 (27%)	Total	920	291 (24%)

Table 1: Count of intra- and inter-sentential relationships in datasets (train+dev) from two domains

Grouin (2016). In NNs, it refers to *iDepNN-ADP*. (5) [*inter-sentential*]-*Shortest-Dependency-Path* ([i-]SDP): Sequence of Words on the [i-]SDP.

Evaluation and Analysis

Dataset. We evaluate our proposed methods on four datasets from medical and news domain. Table 1 shows counts of intra- and inter-sentential relationships. The three medical domain datasets are taken from the BioNLP shared task (ST) of relation/event extraction (Bossy et al. 2011; Nédellec et al. 2013; Deléger et al. 2016). We compare our proposed techniques with the systems published at these venues. The Bacteria Biotope task (Bossy et al. 2011) of the BioNLP ST 2011 focuses on extraction of habitats of bacteria, which is extended by the BioNLP ST 2013 (Nédellec et al. 2013), while the BioNLP ST 2016 focuses on extraction of *Lives_in* events. We have standard train/dev/test splits for the BioNLP ST 2016 dataset, while we perform 3-fold cross-validation¹ on BioNLP ST 2011 and 2013. For BioNLP ST 2016, we generate negative examples by randomly sampling co-occurring entities without known interactions. Then we sample the same number as positives to obtain a balanced dataset during training and validation for different sentence range. See supplementary for further details.

The MUC6 (Grishman and Sundheim 1996) dataset contains information about management succession events from newswire. The task organizers provided a training corpus and a set of templates that contain the management succession events, the names of people who are starting or leaving management posts, the names of their respective posts and organizations and whether the named person is currently in the job. **Entity Tagging:** We tag entities *Person* (*Per*), *Organization* (*Org*) using Stanford NER tagger (Finkel, Grenager, and Manning 2005). The entity type *Position* (*Post*) is annotated based on the templates. **Relation Tagging:** We have three types of relations: *Per-Org*, *Per-Post* and *Post-Org*. We follow Swampillai and Stevenson (2010) and annotate binary relations (within and across sentence boundaries) using management succession events between two entity pairs. We randomly split the collection 60/20/20 into train/dev/test.

Experimental Setup. For MUC6, we use the pretrained GloVe (Pennington, Socher, and Manning 2014) embed-

¹the official evaluation is not accessible any more and therefore, the annotations for their test sets are not available

Dataset: BioNLP ST 2016						
Features	SVM			iDepNN		
	P	R	F_1	P	R	F_1
iSDP	.217	.816	.344	.352	.574	.436
+ PI + ET	.218	.819	.344	.340	.593	.432
+ POS	.269	.749	.396	.348	.568	.431
+ Dependency	.284	.746	.411	.402	.509	.449
Dataset: MUC6						
Features	SVM			iDepNN		
	P	R	F_1	P	R	F_1
iSDP	.689	.630	.627	.916	.912	.913
+ PI	.799	.741	.725	.912	.909	.909
+ POS	.794	.765	.761	.928	.926	.926
+ Dependency	.808	.768	.764	.937	.934	.935

Table 2: SVM vs iDepNN: Features in inter-sentential ($k \leq 1$) training and inter-sentential ($k \leq 1$) evaluation. iSDP+Dependency refers to iDepNN-ADP structure.

dings (200-dimension). For the BioNLP datasets, we use 200-dimensional embedding² vectors from six billion words of biomedical text (Moen and Ananiadou 2013). We randomly initialize a 5-dimensional vectors for PI and POS. We initialize the recurrent weight matrix to identity and biases to zero. We use the macro-averaged F_1 score (the official evaluation script by SemEval-2010 Task 8 (Hendrickx et al. 2010)) on the development set to choose hyperparameters (see supplementary). To report results on BioNLP ST 2016 test set, we use the official web service³.

Baselines. Swampillai and Stevenson’s (2010) annotation of MUC6 intra- and inter-sentential relationships is not available. They investigated SVM with dependency and linguistic features for relationships spanning sentence boundaries. In BioNLP shared tasks, the top performing systems are SVM-based and limited to relationships within single sentences. As an NN baseline, we also develop Connectionist biRNN (Vu et al. 2016a) that spans sentence boundaries; we refer to it as i-biRNN (architecture in supplementary). Similarly, we also investigate using a bidirectional LSTM (i-biLSTM). As a competitive baseline in the inter-sentential relationship extraction, we run⁴ graphLSTM (Peng et al. 2017). This work compares SVM and graphLSTM with i-biRNN, i-biLSTM, iDepNN-SDP and iDepNN-ADP for different values of the sentence range parameter k (the distance in terms of the number of sentences between the entity pairs for a relation), i.e., k ($= 0, \leq 1, \leq 2$ and ≤ 3).

Contribution of different components. Table 2 shows the contribution of each feature, where both training and evaluation is performed over relationships within and across sentence boundaries for sentence range parameter $k \leq 1$. Note: iSDP+Dependency refers to iDepNN-ADP structure that exhibits a better precision, F_1 and balance in precision

²<http://bio.nplab.org/>

³<http://bibliome.jouy.inra.fr/demo/BioNLP-ST-2016-Evaluation/index.html>

⁴hyperparameters in supplementary

train param	Model	Evaluation for different values of sentence range k															
		$k = 0$				$k \leq 1$				$k \leq 2$				$k \leq 3$			
		pr	P	R	F_1	pr	P	R	F_1	pr	P	R	F_1	pr	P	R	F_1
$k = 0$	SVM	363	.474	.512	.492	821	.249	.606	.354	1212	.199	.678	.296	1517	.153	.684	.250
	graphLSTM	473	.472	.668	.554	993	.213	.632	.319	1345	.166	.660	.266	2191	.121	.814	.218
	i-biLSTM	480	.475	.674	.556	998	.220	.652	.328	1376	.165	.668	.265	1637	.132	.640	.219
	i-biRNN	286	.517	.437	.474	425	.301	.378	.335	540	.249	.398	.307	570	.239	.401	.299
	iDepNN-SDP	297	.519	.457	.486	553	.313	.510	.388	729	.240	.518	.328	832	.209	.516	.298
	iDepNN-ADP	266	<u>.526</u>	.414	.467	476	<u>.311</u>	.438	.364	607	<u>.251</u>	.447	.320	669	.226	.447	.300
$k \leq 1$	SVM	471	.464	.645	.540	888	.284	.746	.411	1109	.238	.779	.365	1196	.221	.779	.344
	graphLSTM	406	.502	.607	.548	974	.226	.657	.336	1503	.165	.732	.268	2177	.126	.813	.218
	i-biLSTM	417	.505	.628	.556	1101	.224	.730	.343	1690	.162	.818	.273	1969	.132	.772	.226
	i-biRNN	376	.489	.544	.515	405	.393	.469	.427	406	.391	.469	.426	433	.369	.472	.414
	iDepNN-SDP	303	.561	.503	.531	525	.358	.555	.435	660	.292	.569	.387	724	.265	.568	.362
	iDepNN-ADP	292	.570	.491	.527	428	.402	.509	.449	497	.356	.522	.423	517	.341	.521	.412
$k \leq 2$	SVM	495	.461	.675	.547	1016	.259	.780	.389	1296	.218	.834	.345	1418	.199	.834	.321
	graphLSTM	442	.485	.637	.551	1016	.232	.702	.347	1334	.182	.723	.292	1758	.136	.717	.230
	i-biLSTM	404	.487	.582	.531	940	.245	.682	.360	1205	.185	.661	.289	2146	.128	.816	.222
	i-biRNN	288	.566	.482	.521	462	.376	.515	.435	556	.318	.524	.396	601	.296	.525	.378
	iDepNN-SDP	335	.537	.531	.534	633	.319	.598	.416	832	.258	.634	.367	941	.228	.633	.335
	iDepNN-ADP	309	<u>.538</u>	.493	.514	485	<u>.365</u>	.525	.431	572	<u>.320</u>	.542	.402	603	<u>.302</u>	.540	.387
$k \leq 3$	SVM	507	.458	.686	.549	1172	.234	.811	.363	1629	.186	.894	.308	1874	.162	.897	.275
	graphLSTM	429	.491	.624	.550	1082	.230	.740	.351	1673	.167	.833	.280	2126	.124	.787	.214
	i-biLSTM	417	.478	.582	.526	1142	.224	.758	.345	1218	.162	.833	.273	2091	.128	.800	.223
	i-biRNN	405	.464	.559	.507	622	.324	.601	.422	654	.310	.604	.410	655	.311	.607	.410
	iDepNN-SDP	351	.533	.552	.542	651	.315	.605	.414	842	.251	.622	.357	928	.227	.622	.333
	iDepNN-ADP	313	<u>.553</u>	.512	.532	541	<u>.355</u>	.568	.437	654	<u>.315</u>	.601	.415	687	<u>.300</u>	.601	.401
$k \leq 1$	<i>ensemble</i>	480	.478	.680	.561	837	.311	.769	.443	1003	.268	.794	.401	1074	.252	.797	.382

Table 3: BioNLP ST 2016 Dataset: Performance of the intra-and-inter-sentential training/evaluation for different k . **Underline:** Better precision by *iDepNN-ADP* over *iDepNN-SDP*, graphLSTM and SVM. **Bold:** Best in column. **pr:** Count of predictions

and recall, compared to SVM. See supplementary for feature analysis on BioNLP ST 2011 / 2013.

State-of-the-Art Comparisons

BioNLP ST 2016 dataset: Table 3 shows the performance of {SVM, graphLSTM} vs {i-biRNN, iDepNN-SDP, iDepNN-ADP} for relationships within and across sentence boundaries. *Moving left to right for each training parameter, the recall increases while precision decreases due to increasing noise with larger k . In the inter-sentential evaluations ($k \leq 1, \leq 2, \leq 3$ columns) for all the training parameters, the iDepNN variants outperform both SVM and graphLSTM in terms of F_1 and maintain a better precision as well as balance in precision and recall with increasing k ; e.g., at $k \leq 1$ (train/eval), precision and F_1 are (.402 vs .226) and (.449 vs .336), respectively for (iDepNN-ADP vs graphLSTM). We find that SVM mostly leads in recall.*

In comparison to graphLSTM, i-biRNN and i-biLSTM, we observe that iDepNN-ADP offers precise structure in relation extraction within and across sentence boundaries. For instance, at training $k \leq 1$ and evaluation $k = 0$, iDepNN-ADP reports precision of .570 vs .489 and .561 in i-biRNN and iDepNN-SDP, respectively. During training at $k \leq 1$, iDepNN-SDP and iDepNN-ADP report better F_1 than i-biRNN for evaluations at all k , suggesting that the shortest and augmented paths provide useful dependency features via

threshold	<i>ensemble</i> (train $k \leq 1$ and evaluation $k = 0$)							
	<i>Dev</i> (official scores)				<i>Test</i> (official scores)			
	pr	P	R	F_1	pr	P	R	F_1
$p \geq 0.85$	160	.694	.514	.591	419	.530	.657	.587
$p \geq 0.90$	151	.709	.496	.583	395	.539	.630	.581
$p \geq 0.95$	123	.740	.419	.535	293	.573	.497	.533

Table 4: Ensemble scores at various thresholds for BioNLP ST 2016 dataset. p : output probability

recurrent and recursive compositions, respectively. Between the proposed architectures iDepNN-SDP and iDepNN-ADP, the former achieves higher recall for all k . We find that the training at $k \leq 1$ is optimal for intra- and inter-sentential relations over development set (see supplementary). We also observe that i-biRNN establishes a strong NN baseline for relationships spanning sentences. The proposed architectures consistently outperform graphLSTM in both precision and F_1 across sentence boundaries.

Ensemble: We exploit the precision and recall bias of the different models via an ensemble approach, similar to TurkuNLP (Mehryary et al. 2016) and UMS (Deléger et al. 2016) systems that combined predictions from SVM and NNs. We aggregate the prediction outputs of the neural (i-biRNN, iDepNN-SDP and iDepNN-ADP) and non-neural




Figure 3: Left: SVM, graphLSTM & iDepNN-ADP on BioNLP ST 2016: Performance analysis on relations that span sentence boundaries, with different sentence range parameters Right: BioNLP 2016 ST dataset (official results on test set): Comparison with the published systems in the BioNLP ST, where pr is the count of predictions. This work demonstrates a better balance in precision and recall, and achieves the highest F_1 and *recall*. We extract 419 predictions within and across sentence boundaries, which is closer to the count of gold predictions, i.e., 340 (Deléger et al. 2016).

(SVM) classifiers, i.e., a relation to hold if any classifier has predicted it. We perform the *ensemble* scheme on the development and official test sets for intra- and inter-sentential (optimal at $k \leq 1$) relations. Table 3 shows the ensemble scores on the official test set for relations within and across sentence boundaries, where *ensemble* achieves the highest F_1 (.561) over individual models.

Confident Extractions: We consider the high confidence prediction outputs by the different models participating in the *ensemble*, since it lacks precision (.478). Following Peng et al. (2017), we examine three values of the output probability p , i.e., ($\geq 0.85, 0.90$ and 0.95) of each model in the *ensemble*. Table 4 shows the *ensemble* performance on the development and official test sets, where the predictions with $p \geq 0.85$ achieve the state-of-the-art performance and rank us at the top out of 11 systems (Figure 3, right).

This Work vs Competing Systems in BioNLP ST 2016: As shown in Figure 3 (right), we rank at the top and achieve a gain of 5.2% (.587 vs .558) in F_1 compared to VERSE. We also show a better balance in precision and recall, and report the highest recall compared to all other systems. Most systems do not attempt to predict relations spanning sentences. The most popular algorithms are SVM (VERSE, HK, UTS, LIMSI) and NNs (TurkuNLP, WhuNlpRE, DUTIR). UMS combined predictions from an SVM and an NN. Most systems rely on syntactic parsing, POS, word embeddings and entity recognition features (VERSE, TurkuNLP, UMS, HK, DUTIR, UTS). VERSE and TurkuNLP obtained top scores on intra-sentential relations relying on the dependency path features between entities; however they are limited to intra-sentential relations. TurkuNLP employed an ensemble of 15

different LSTM based classifiers. DUTIR is based on CNN for intra-sentential relationships. LIMSI is the only system that considers inter-sentential relationships during training; however it is SVM-based and used additional manually annotated training data, linguistic features using biomedical resources (PubMed, Cocoa web API, OntoBiotope ontology, etc.) and post-processing to annotate biomedical abbreviations. We report a noticeable gain of 21% (.587 vs .485) in F_1 with an improved precision and recall over LIMSI.

BioNLP ST 2011 and 2013 datasets: Following the BioNLP ST 2016 evaluation, we also examine two additional datasets from the same domain. iDepNN-ADP (Table 5) is the leading performer in terms of precision and F_1 within and across boundaries for BioNLP ST 2013. Examining BioNLP ST 2011, the iDepNN variants lead both SVM and i-biRNN for $k \leq 1$ and $k \leq 2$.

MUC6 dataset: Similar to BioNLP ST 2016, we perform training and evaluation of SVM, i-biRNN, iDepNN-SDP and iDepNN-ADP for different sentence range with best feature combination (Table 2) using MUC6 dataset. Table 6 shows that both iDepNN variants consistently outperform graphLSTM and SVM for relationships within and across sentences. For within ($k=0$) sentence evaluation, iDepNN-ADP reports .963 F_1 , compared to .779 and .783 by SVM and graphLSTM, respectively. iDepNN-ADP is observed more precise than iDepNN-SDP and graphLSTM with increasing k , e.g., at $k \leq 3$. Training at sentence range $k \leq 1$ is found optimal in extracting inter-sentential relationships.

Model	Dataset: BioNLP ST 2013												Dataset: BioNLP ST 2011											
	$k = 0$			$k \leq 1$			$k \leq 2$			$k \leq 3$			$k = 0$			$k \leq 1$			$k \leq 2$			$k \leq 3$		
	P	R	F ₁	P	R	F ₁	P	R	F ₁	P	R	F ₁	P	R	F ₁	P	R	F ₁	P	R	F ₁	P	R	F ₁
SVM	.95	.90	.92	.87	.83	.85	.95	.90	.92	.95	.90	.92	.98	.96	.97	.87	.87	.87	.95	.94	.94	.91	.88	.90
graphLSTM	.98	.97	.97	.94	.95	.94	.95	.89	.92	.90	.97	.93	.99	.99	.99	.95	.98	.96	.95	.97	.96	.96	.92	.93
i-biLSTM	.98	.97	.97	.96	.95	.95	.93	.96	.94	.91	.93	.92	.99	.99	.99	.95	.98	.96	.96	.97	.96	.95	.92	.93
i-biRNN	.95	.94	.94	.93	.90	.91	.94	.92	.93	.94	.84	.89	.97	.99	.98	.88	.94	.90	.92	.94	.93	.96	.96	.96
iDepNN-SDP	.94	.96	.95	.94	.95	.94	.87	.92	.89	.91	.94	.92	.97	.99	.98	.96	.92	.93	.97	.97	.97	.94	.91	.92
iDepNN-ADP	.99	.98	.99	.97	.94	.95	.98	.95	.96	.96	.91	.93	.97	.97	.97	.93	.96	.94	.92	.98	.95	.93	.94	.93

Table 5: BioNLP ST 2011 and 2013 datasets: Performance for training ($k \leq 1$) and evaluation for different k . Underline: Better precision in iDepNN-ADP than iDepNN-SDP, graphLSTM, i-biLSTM, i-biRNN and SVM. **Bold**: best in column.

train param	Model	Evaluation for different k								
		$k = 0$			$k \leq 1$			$k \leq 3$		
		P	R	F ₁	P	R	F ₁	P	R	F ₁
$k = 0$	SVM	.796	.765	.760	.775	.762	.759	.791	.779	.776
	graphLSTM	.910	.857	.880	.867	.897	.870	.870	.867	.870
	i-biLSTM	.917	.837	.873	.833	.896	.863	.853	.870	.863
	i-biRNN	.875	.859	.864	.828	.822	.824	.830	.827	.827
	iDepNN-SDP	.958	.948	.952	.934	.928	.930	.935	.930	.932
	iDepNN-ADP	.933	.927	.929	.924	.920	.921	.930	.927	.927
$k \leq 1$	SVM	.815	.772	.769	.808	.768	.764	.802	.775	.770
	graphLSTM	.730	.900	.783	.727	.907	.773	.730	.913	.770
	i-biLSTM	.760	.880	.780	.670	.950	.767	.697	.937	.770
	i-biRNN	.925	.934	.927	.870	.872	.860	.868	.866	.858
	iDepNN-SDP	.949	.945	.946	.928	.926	.926	.934	.932	.932
	iDepNN-ADP	.961	.955	.957	.937	.934	.935	.942	.940	.940
$k \leq 3$	SVM	.840	.785	.779	.816	.779	.774	.822	.788	.781
	graphLSTM	.737	.907	.783	.703	.927	.773	.710	.927	.767
	i-biLSTM	.720	.920	.780	.680	.943	.770	.700	.932	.770
	i-biRNN	.944	.934	.938	.902	.890	.895	.926	.923	.924
	iDepNN-SDP	.956	.947	.951	.920	.916	.917	.939	.938	.936
	iDepNN-ADP	.965	.963	.963	.933	.933	.931	.939	.938	.936

Table 6: MUC6 Dataset: Performance over the intra- and inter-sentential training and evaluation for different k . Underline signifies better precision by iDepNN-ADP over iDepNN-SDP, graphLSTM, i-biLSTM, i-biRNN and SVM. **Bold** indicates the best score column-wise.

Error Analysis and Discussion

In Figure 3 (left), we analyze predictions using different values of sentence range k ($=0, \leq 1, \leq 2$ and ≤ 3) during both training and evaluation of SVM, graphLSTM and iDepNN-ADP for BioNLP ST 2016 dataset. For instance, an SVM (top-left) is trained for intra-sentential (*same sentence*) relations, while iDepNN-ADP (bottom-right) for both intra- and inter-sentential spanning three sentences (*three sentences apart*). We show how the count of true positives (TP), false negatives (FN) and false positives (FP) varies with k .

Observe that as the distance of the relation increases, the classifiers predict larger ratios of false positives to true positives. However, as the sentence range increases, iDepNN-ADP outperforms both SVM and graphLSTM due to fewer false positives (red colored bars). On top, the ratio of FP to TP is better in iDepNN-ADP than graphLSTM and SVM for all values of k . Correspondingly in Table 3, iDepNN-ADP reports better precision and balance between precision and recall, signifying its robustness to noise in handling inter-sentential relationships.

iDepNN vs graphLSTM: Peng et al. (2017) focuses on general relation extraction framework using graphLSTM

with challenges such as potential cycles in the document graph leading to expensive model training and difficulties in convergence due to loopy gradient backpropagation. Therefore, they further investigated different strategies to back-propagate gradients. The graphLSTM introduces a number of parameters with a number of edge types and thus, requires abundant supervision/training data. On other hand, our work introduces simple and robust neural architectures (iDepNN-SDP and iDepNN-ADP), where the iDepNN-ADP is a special case of document graph in form of a parse tree spanning sentence boundaries. We offer a smooth gradient backpropagation in the complete structure (e.g., in iDepNN-ADP via recurrent and recursive hidden vectors) that is more efficient than graphLSTM due to non-cyclic (i.e., tree) architecture. We have also shown that iDepNN-ADP is robust to false positives and maintains a better balance in precision and recall than graphLSTM for inter-sentential relationships (Figure 3).

Conclusion

We have proposed to classify relations between entities within and across sentence boundaries by modeling the inter-sentential shortest and augmented dependency paths within a novel neural network, named as inter-sentential Dependency-based Neural Network (iDepNN) that takes advantage of both recurrent and recursive neural networks to model the structures in the intra- and inter-sentential relationships. Experimental results on four datasets from newswire and medical domains have demonstrated that iDepNN is robust to false positives, shows a better balance in precision and recall and achieves the state-of-the-art performance in extracting relationships within and across sentence boundaries. We also perform better than 11 teams participating in the BioNLP shared task 2016.

References

- Bahcall, O. 2015. Precision medicine. *Nature* 526:335.
- Bossy, R.; Jourde, J.; Bessieres, P.; Van De Guchte, M.; and Nédellec, C. 2011. Bionlp shared task 2011: bacteria biotope. In *Proceedings of the BioNLP Shared Task 2011 Workshop*, 56–64. Association for Computational Linguistics.
- Bunescu, R. C., and Mooney, R. J. 2005. A shortest path dependency kernel for relation extraction. In *Proceedings of the conference on human language technology and empirical methods in natural language processing*, 724–731. Association for Computational Linguistics.

- Deléger, L.; Bossy, R.; Chaix, E.; Ba, M.; Ferré, A.; Bessieres, P.; and Nédellec, C. 2016. Overview of the bacteria biotope task at bionlp shared task 2016. In *Proceedings of the 4th BioNLP Shared Task Workshop*, 12–22.
- Finkel, J. R.; Grenager, T.; and Manning, C. 2005. Incorporating non-local information into information extraction systems by gibbs sampling. In *Proceedings of the 43rd annual meeting on association for computational linguistics*, 363–370. Association for Computational Linguistics.
- Gerber, M., and Chai, J. Y. 2010. Beyond nombank: A study of implicit arguments for nominal predicates. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, 1583–1592. Association for Computational Linguistics.
- Grishman, R., and Sundheim, B. 1996. Message understanding conference-6: A brief history. In *COLING 1996 Volume 1: The 16th International Conference on Computational Linguistics*, volume 1.
- Grouin, C. 2016. Identification of mentions and relations between bacteria and biotope from pubmed abstracts. *ACL 2016* 64.
- Gupta, P.; Schütze, H.; and Andrassy, B. 2016. Table filling multi-task recurrent neural network for joint entity and relation extraction. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, 2537–2547.
- Gupta, P. 2015. Deep Learning Methods for the Extraction of Relations in Natural Language Text. Master’s thesis, Technical University of Munich, Germany.
- Hagberg, A.; Swart, P.; and S Chult, D. 2008. Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Laboratory (LANL).
- Hendrickx, I.; Kim, S. N.; Kozareva, Z.; Nakov, P.; Séaghdha, D. O.; Padó, S.; Pennacchiotti, M.; Romano, L.; and Szpakowicz, S. 2010. Semeval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals. *ACL 2010* 33.
- Kambhatla, N. 2004. Combining lexical, syntactic, and semantic features with maximum entropy models for extracting relations. In *Proceedings of the ACL 2004 on Interactive poster and demonstration sessions*, 22. Association for Computational Linguistics.
- Liu, Y.; Wei, F.; Li, S.; Ji, H.; Zhou, M.; and Wang, H. 2015. A dependency-based neural network for relation classification. *arXiv preprint arXiv:1507.04646*.
- Manning, C. D.; Surdeanu, M.; Bauer, J.; Finkel, J. R.; Bethard, S.; and McClosky, D. 2014. The stanford corenlp natural language processing toolkit. In *ACL (System Demonstrations)*, 55–60.
- Mehryary, F.; Björne, J.; Pyysalo, S.; Salakoski, T.; and Ginter, F. 2016. Deep learning with minimal training data: Turkunlp entry in the bionlp shared task 2016. In *Proceedings of the 4th BioNLP Shared Task Workshop*, 73–81.
- Mintz, M.; Bills, S.; Snow, R.; and Jurafsky, D. 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*, 1003–1011. Association for Computational Linguistics.
- Moen, S., and Ananiadou, T. S. S. 2013. Distributional semantics resources for biomedical text processing.
- Nédellec, C.; Bossy, R.; Kim, J.-D.; Kim, J.-J.; Ohta, T.; Pyysalo, S.; and Zweigenbaum, P. 2013. Overview of bionlp shared task 2013. In *Proceedings of the BioNLP Shared Task 2013 Workshop*, 1–7. Association for Computational Linguistics Sofia, Bulgaria.
- Nguyen, D. P.; Matsuo, Y.; and Ishizuka, M. 2007. Relation extraction from wikipedia using subtree mining. In *Proceedings of the National Conference on Artificial Intelligence*, volume 22, 1414. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999.
- Peng, N.; Poon, H.; Quirk, C.; Toutanova, K.; and Yih, W.-t. 2017. Cross-sentence n-ary relation extraction with graph lstms. *Transactions of the Association for Computational Linguistics* 5:101–115.
- Pennington, J.; Socher, R.; and Manning, C. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 1532–1543.
- Quirk, C., and Poon, H. 2016. Distant supervision for relation extraction beyond the sentence boundary. *arXiv preprint arXiv:1609.04873*.
- Schuster, M., and Paliwal, K. K. 1997. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing* 45(11):2673–2681.
- Socher, R.; Huval, B.; Manning, C. D.; and Ng, A. Y. 2012. Semantic compositionality through recursive matrix-vector spaces. In *Proceedings of the 2012 joint conference on empirical methods in natural language processing and computational natural language learning*, 1201–1211. Association for Computational Linguistics.
- Socher, R.; Karpathy, A.; Le, Q. V.; Manning, C. D.; and Ng, A. Y. 2014. Grounded compositional semantics for finding and describing images with sentences. *Transactions of the Association for Computational Linguistics* 2:207–218.
- Swampillai, K., and Stevenson, M. 2010. Inter-sentential relations in information extraction corpora. In *LREC*.
- Swampillai, K., and Stevenson, M. 2011. Extracting relations within and across sentences. In *RANLP*, 25–32.
- Vu, N. T.; Adel, H.; Gupta, P.; and Schütze, H. 2016a. Combining recurrent and convolutional neural networks for relation classification. In *Proceedings of the NAACL-HLT*, 534–539. San Diego, California USA: Association for Computational Linguistics.
- Vu, N. T.; Gupta, P.; Adel, H.; and Schütze, H. 2016b. Bi-directional recurrent neural network with ranking loss for spoken language understanding. In *Proceedings of the Acoustics, Speech and Signal Processing (ICASSP)*, 6060–6064. Shanghai, China: IEEE.
- Werbos, P. J. 1990. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE* 78(10):1550–1560.
- Yoshikawa, K.; Riedel, S.; Hirao, T.; Asahara, M.; and Matsumoto, Y. 2011. Coreference based event-argument relation extraction on biomedical text. *Journal of Biomedical Semantics* 2(5):S6.
- Zhang, M.; Zhang, J.; Su, J.; and Zhou, G. 2006. A composite kernel to extract relations between entities with both flat and structured features. In *Proceedings of the 21st International Conference on Computational Linguistics and the 44th annual meeting of the Association for Computational Linguistics*, 825–832. Association for Computational Linguistics.