

Dual Supervision Framework for Relation Extraction with Distant Supervision and Human Annotation

Woohwan Jung

Seoul National University
Seoul, Korea
whjung@kdd.snu.ac.kr

Kyuseok Shim

Seoul National University
Seoul, Korea
kshim@snu.ac.kr

Abstract

Relation extraction (RE) has been extensively studied due to its importance in real-world applications such as knowledge base construction and question answering. Most of the existing works train the models on either distantly supervised data or human-annotated data. To take advantage of the high accuracy of human annotation and the cheap cost of distant supervision, we propose the dual supervision framework which effectively utilizes both types of data. However, simply combining the two types of data to train a RE model may decrease the prediction accuracy since distant supervision has labeling bias. We employ two separate prediction networks *HA-Net* and *DS-Net* to predict the labels by human annotation and distant supervision, respectively, to prevent the degradation of accuracy by the incorrect labeling of distant supervision. Furthermore, we propose an additional loss term called *disagreement penalty* to enable *HA-Net* to learn from distantly supervised labels. In addition, we exploit additional networks to adaptively assess the labeling bias by considering contextual information. Our performance study on sentence-level and document-level REs confirms the effectiveness of the dual supervision framework.

1 Introduction

Relation extraction (RE) has been widely used in real-world applications such as knowledge base construction (Dong et al., 2014a; Dong et al., 2014b; Jung et al., 2019), question answering (Xu et al., 2016) and biomedical data mining (Ahmed et al., 2019). Given a pair of entities in a text (e.g., sentence or document), the goal of RE is to discover the relationships between the entities expressed in the text. More specifically, we aim to extract triples from the text in the form of $\langle e_h, r, e_t \rangle$ where e_h is a head entity, e_t is a tail entity and r is a relationship between the entities.

To train a model for RE, we need a large volume of fully labeled training data in the form of text-triple pairs. Although human annotation provides high-quality labels to train the relation extraction models, it is difficult to produce a large-scale training data since manual labeling is expensive and time-consuming. Thus, Mintz et al. (2009) proposed distant supervision to automatically produce a large labeled data by using an external knowledge base (KB). For a text with a head entity e_h and a tail entity e_t , when a triple $\langle e_h, r, e_t \rangle$ exists in the KB for any relation type r , distant supervision produces a label $\langle e_h, r, e_t \rangle$ even though the relationship is not expressed in the text. Thus, it suffers from the wrong labeling problem. For instance, if a triple $\langle UK, capital, London \rangle$ is in the KB, distant supervision labels the triple even for the sentence ‘London is the largest city of the UK’.

Although each of the two labeling methods has a certain weakness, most of the existing works for RE utilize either human-annotated (HA) data or distantly supervised (DS) data. To take advantage of the high accuracy of human annotation and the cheap cost of distant supervision, we propose to effectively utilize a large DS data as well as a small amount of HA data. Since DS data is likely to have *labeling bias*, simply combining the two types of data to train a RE model may decrease the prediction accuracy. To take a close look at the labeling bias, let the *inflation* of a relation type be the ratio of the average frequencies of the relation type per text in DS data and HA data, respectively. We say that a relation

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

type is *unbiased* if the average frequency of the relation type in DS data is the same as that in HA data (i.e., the inflation of the relation is 1). By examining a document-level RE dataset (DocRED) (Yao et al., 2019) with 96 relation types, we found that the inflations of the relation types are from 0.48 to 85.9. It indicates that distant supervision tends to generate a large number of false labels for some relation types.

Recently, Ye et al. (2019) introduced a domain adaptation approach to tackle the labeling bias problem for RE. It trains a RE model on DS data and adjusts the bias term of the output layer by using HA data. Although the bias adjustment achieves a meaningful accuracy improvement, it has a limitation. An underlying assumption of the method is that the labeling bias is static for every text since it adjusts the bias term only once after training and uses the same bias during the test time. However, the labeling bias varies depending on contextual information. For example, in DocRED dataset, most of the `capital` relation labeled by distant supervision are false positive. However, if the phrase ‘is the capital city of’ appears in the text, the label is likely to be a true label. Thus, we need to take account of contextual information to extract relations more accurately by considering the labeling bias.

To effectively utilize DS data and HA data for training RE models, we propose the *dual supervision framework* that can be applied to most existing RE models to achieve additional accuracy gain. Since the label distributions in HA data and DS data are quite different, we cast the task of training RE models with both data as a multi-task learning problem. Thus, we employ the two separate output modules *HA-Net* and *DS-Net* to predict the labels by human annotation and distant supervision, respectively, while previous works utilize a single output module. This allows the different predictions of the labels for human annotation and distant supervision, and thus it prevents the degradation of accuracy by incorrect labels in DS data. If we simply separate the prediction networks to apply the multi-task learning, *HA-Net* cannot learn from distantly supervised labels. To enable *HA-Net* to learn from DS data, we propose an additional loss term called *disagreement penalty*. It models the ratio of the output probabilities from the prediction networks *HA-Net* and *DS-Net* by using maximum likelihood estimation with log-normal distributions to generate the calibrated gradient to update *HA-Net* to effectively reflect distantly supervised labels. Furthermore, our framework exploits two additional networks μ -Net and σ -Net to adaptively estimate the log-normal distribution by considering contextual information. Moreover, we theoretically show that the disagreement penalty enables *HA-Net* to effectively utilize the labels generated by distant supervision. Finally, we validate the effectiveness of the dual supervision framework on two types of tasks: sentence-level and document-level REs. The experimental results confirm that our dual supervision framework significantly improves the prediction accuracy of existing RE models. In addition, the dual supervision framework substantially outperforms the state-of-the-art method (Ye et al., 2019) in both sentence-level and document-level REs with the relative F1 score improvement of up to 32%.

2 Preliminaries

We present the problems of sentence-level and document-level relation extractions and next introduce existing works for relation extraction.

2.1 Problem Statement

Following the works (Yao et al., 2019; Wang et al., 2019), we assume that each text is annotated with entity mentions. For a pair of entities, since a sentence usually describes a single relationship between them, the sentence-level relation extraction is generally regarded as a *multi-class* classification problem.

Definition 2.1 (Sentence-level relation extraction) *For a pair of the head and tail entities e_h and e_t , a relation type set R and a sentence s annotated with entity mentions, we determine the relation $r \in R$ between e_h and e_t in the sentence. Note that R includes a special relation type NA which indicates that there does not exist any relation between e_h and e_t .*

Since multiple relationships between a pair of entities can be expressed in a document, document-level relation extraction is usually defined as a *multi-label* classification problem.

Definition 2.2 (Document-level relation extraction) *For a pair of the head and tail entities e_h and e_t , a relation type set R and a document d annotated with entity mentions, we find the set of all relations*

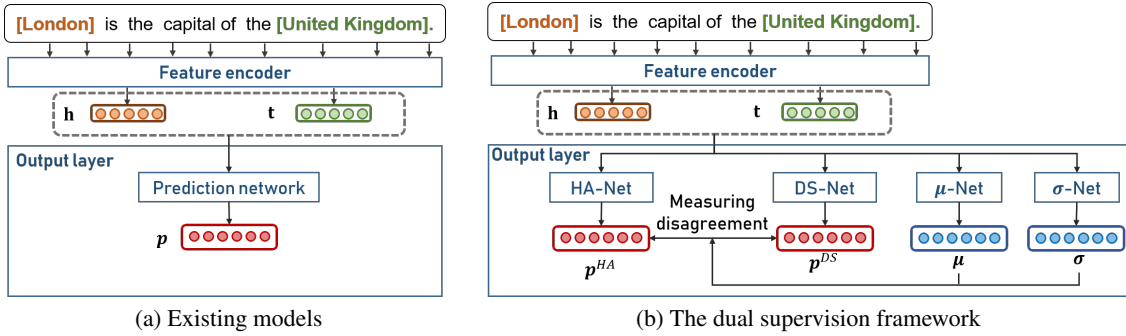


Figure 1: The overall architectures of existing models and our framework

$R^* \subset R$ between e_h and e_t appearing in document d . Note that R does not include NA in this case since it can be represented by an empty set of R^* .

In this paper, we mainly discuss sentence-level RE and extend our framework to document-level RE.

2.2 Existing Works of Relation Extraction

A typical RE model consists of a feature encoder and a prediction network, as shown in Figure 1(a). The feature encoder converts a text into the hidden representations of the head and tail entities. Cai et al. (2016) and Wang et al. (2019) exploit Bi-LSTM and BERT, respectively, to encode the text. On the other hand, Zeng et al. (2014) and Zeng et al. (2015) use CNN for the encoder. In addition, Zeng et al. (2014) propose the position embedding to consider the relative distance from each word to head and tail entities.

The prediction network outputs the probability distribution of the relations between the entities. Since sentence-level RE is a multi-class classification task, sentence-level RE models (Cai et al., 2016; Zeng et al., 2014; Zeng et al., 2015) utilize a *softmax classifier* as the prediction network and use categorical cross entropy as the loss function. On the other hand, document-level RE models (Yao et al., 2019; Wang et al., 2019) use a *sigmoid classifier* and binary cross entropy as the prediction network and the loss function, respectively. Since the labels obtained from distant supervision are noisy and biased, with a single prediction network, it is hard to make accurate predictions for DS data and HA data together.

3 Dual Supervision Framework

We first present an overview of the dual supervision framework which effectively utilizes both human-annotated (HA) data and distantly supervised (DS) data for training RE models. We next introduce the detailed structure of the output layer in our framework and propose our novel loss function with disagreement penalty that considers the labeling bias of distant supervision. Then, we describe how to train the proposed model with both types of data as well as how to extract relations from the test data. Finally, we discuss how the disagreement penalty makes each prediction network learn from the labels for the other prediction network although we use separate prediction networks.

3.1 An Overview of the Dual Supervision Framework

As shown in Figure 1(b), our framework consists of a feature encoder and an output layer with 4 sub-networks. It is general enough to accommodate a variety of existing RE models to improve their accuracy. We can apply our framework to an existing RE model by using the feature encoder of the model and building the four sub-networks which exploit the structure of the original prediction network. Since our framework uses the feature encoder of the existing models, we briefly describe only the output layer here.

Unlike the previous works, to allow the difference in the predictions for human annotated labels and distantly supervised labels, we exploit multi-task learning by employing two separate prediction networks *HA-Net* and *DS-Net* to predict the labels in HA data and DS data, respectively. We also use *HA-Net* to extract relations from the test data. The separation of the prediction networks prevents the accuracy degradation caused by incorrect labels from distant supervision. If we simply utilize two prediction networks to apply the multi-task learning, *HA-Net* cannot learn from distantly supervised labels

although the prediction networks share the feature encoder. To enable *HA-Net* to learn from distantly supervised labels, we introduce an additional loss term called *disagreement penalty*. It models the disagreement between the outputs of *HA-Net* and *DS-Net* by using maximum likelihood estimation with log-normal distributions. Furthermore, to adaptively estimate the parameters of the log-normal distribution by considering contextual information, we exploit two parameter networks μ -Net and σ -Net.

For a label $\langle e_h, r, e_t \rangle$, let I_{HA} be an indicator variable that is 1 if the label is obtained by human annotation and 0 otherwise. The proposed framework uses the following loss function for a label $\langle e_h, r, e_t \rangle$

$$L_{h,t} = I_{HA} \cdot L_{h,t}^{HA} + (1 - I_{HA}) \cdot L_{h,t}^{DS} + \lambda \cdot L_{h,t}^{DS-HA} \quad (1)$$

where $L_{h,t}^{HA}$ and $L_{h,t}^{DS}$ denote the prediction loss of *HA-Net* and *DS-Net*, respectively, and $L_{h,t}^{DS-HA}$ is the disagreement penalty to capture the distance between the predictions by *HA-Net* and *DS-Net*. The hyper parameter λ controls the relative importance of the disagreement penalty to the prediction errors. By using a separate prediction network for each type of data and introducing the disagreement penalty, *HA-Net* learns from distantly supervised labels while reducing overfitting to noisy DS data.

3.2 Separate Prediction Networks

To alleviate the accuracy degradation from the noisy labels in DS data, we utilize two prediction networks. The network *HA-Net* is used to predict the human-annotated labels from the train data and to predict relations from the test data. The other prediction network *DS-Net* predicts the labels obtained by distant supervision. We use the prediction network of an existing model for both prediction networks of our framework without sharing the model parameters. The prediction networks *HA-Net* and *DS-Net* output the $|R|$ -dimensional vectors $\mathbf{p}^{HA} = [p(r_1|e_h, e_t, HA), \dots, p(r_{|R|}|e_h, e_t, HA)]$ and $\mathbf{p}^{DS} = [p(r_1|e_h, e_t, DS), \dots, p(r_{|R|}|e_h, e_t, DS)]$, respectively, where $p(r|e_h, e_t, HA)$ and $p(r|e_h, e_t, DS)$ are the probabilities that there exists a label $\langle e_h, r, e_t \rangle$, in HA data and DS data, respectively. We simply denote $p(r|e_h, e_t, HA)$ and $p(r|e_h, e_t, DS)$ by p_r^{HA} and p_r^{DS} , respectively.

3.3 Disagreement Penalty

Distant supervision labels are biased and the size of the bias varies depending on the type of relation. Moreover, the bias can vary depending on many other features such as the types of head and tail entities as well as the contents of a text. Thus, we propose to use an effective disagreement penalty to model the labeling bias depending on the context where the head and tail entities are located.

Distribution of inflations. We measure the labeling bias by using the inflations of relations. Recall that the inflation of a relation type is the ratio of the average frequencies of the relation type per text in DS data and HA data, respectively. To investigate the distribution of inflations, we computed the inflations of 96 relation types in DocRED data. Since Kolmogorov-Smirnov (K-S) test (Massey Jr, 1951) is widely used to determine whether an observed data is drawn from a given probability distribution, we used it to find the best-fit distribution of the inflations. Since the range of the inflation is $[0, \infty)$, we evaluated p-values of the four probability distributions supported on $[0, \infty)$: Log-normal, Weibull, chi-square and exponential distributions. In addition, we include the normal distribution as a baseline. Table 1 shows the result of K-S test for DocRED data. Note that a probability distribution has a high p-value if the probability distribution fits the data well. Since the log-normal distribution has the highest p-value, it is the best-fit distribution among the five probability distributions. Based on the observation, we model the disagreement penalty between the outputs of the two prediction networks.

Distribution	p-value
Log-normal	0.008
Weibull	0.001
Chi-square	4.6×10^{-10}
Exponential	3.6×10^{-13}
Normal	1.2×10^{-15}

Table 1: The result of K-S test

Modeling the disagreement penalty. We develop the disagreement penalty based on the maximum likelihood estimation. Let X_r be the random variable which denotes the ratio of p_r^{DS} to p_r^{HA} . Since the inflation is the ratio of the number of labels in DS data and HA data, the ratio p_r^{DS}/p_r^{HA} represents the

conditional inflation of the relation type r conditioned on the text with head and tail entities. Thus, we assume that X_r follows a log-normal distribution $L(\mu_r, \sigma_r^2)$ whose probability density function is

$$f(x) = \frac{1}{x\sigma_r\sqrt{2\pi}} \exp\left(-\frac{(\log x - \mu_r)^2}{2\sigma_r^2}\right). \quad (2)$$

The disagreement penalty $L_{h,t}^{DS-HA}$ is defined as the negative log likelihood of the conditional inflation p_r^{DS}/p_r^{HA} , which is obtained by substituting p_r^{DS}/p_r^{HA} into Equation (2) as follows:

$$-\log f(p_r^{DS}/p_r^{HA}) = \frac{1}{2} \left(\frac{\log p_r^{DS} - \log p_r^{HA} - \mu_r}{\sigma_r} \right)^2 + \log p_r^{DS} - \log p_r^{HA} + \log \sigma_r + \frac{\log 2\pi}{2}. \quad (3)$$

Since $\frac{\log 2\pi}{2}$ is constant, we utilize the disagreement penalty in Equation (3) without the constant term.

If we set μ_r and σ_r to fixed values, we cannot effectively assess the conditional inflation since it can vary depending on the context. For example, although the inflation of the relation type `capital` is high, the conditional inflation should be lower if a particular phrase such as ‘is the capital city of’ appears in the text. To take account of the contextual information, we employ two additional networks μ -Net and σ -Net to estimate the μ_r and σ_r that are the parameters of log-normal distribution $L(\mu_r, \sigma_r^2)$.

3.4 Parameter Networks

The parameter networks μ -Net and σ -Net output the vectors $\boldsymbol{\mu} = [\mu_1, \dots, \mu_{|R|}]$ and $\boldsymbol{\sigma} = [\sigma_1, \dots, \sigma_{|R|}]$, respectively, which are the parameters of the log-normal distributions to represent the conditional inflation for $r \in R$. Both μ -Net and σ -Net have the same structure as those of the prediction networks except their output activation functions. For a log-normal distribution $L(\mu, \sigma)$, the parameter μ can be positive or negative, and σ is always positive. Thus, we use a hyperbolic tangent function and a softplus function (Dugas et al., 2001) as the output activation functions of μ -Net and σ -Net, respectively.

For example, if the prediction network of the original RE model consists of a bilinear layer and an output activation function, the parameter vectors $\boldsymbol{\mu} \in \mathbb{R}^{|R|}$ and $\boldsymbol{\sigma} \in \mathbb{R}^{|R|}$ are computed from the head entity vector $\mathbf{h} \in \mathbb{R}^d$ and tail entity vector $\mathbf{t} \in \mathbb{R}^d$ as

$$\boldsymbol{\mu} = \tanh(\mathbf{h}^\top \mathbf{W}^\mu \mathbf{t} + \mathbf{b}^\mu), \quad \boldsymbol{\sigma} = \text{softplus}(\mathbf{h}^\top \mathbf{W}^\sigma \mathbf{t} + \mathbf{b}^\sigma) + \varepsilon$$

where $\text{softplus}(x) = \log(1 + e^x)$ and ε is a sanity bound preventing extremely small values of σ_r from dominating the loss function, and $\mathbf{W}^\mu \in \mathbb{R}^{d \times |R| \times d}$, $\mathbf{W}^\sigma \in \mathbb{R}^{d \times |R| \times d}$, $\mathbf{b}^\mu \in \mathbb{R}^{|R|}$ and $\mathbf{b}^\sigma \in \mathbb{R}^{|R|}$ are learnable parameters. We set the sanity bound ε to 0.0001 in our experiment.

3.5 Loss Function

For sentence-level relation extraction, we use the categorical cross entropy loss as the prediction losses $L_{h,t}^{HA}$ and $L_{h,t}^{DS}$. For a label $\langle e_h, r, e_t \rangle$, we obtain the following loss function from Equations (1) and (3)

$$\begin{aligned} L_{h,t} &= I_{HA} \cdot L_{h,t}^{HA} + (1 - I_{HA}) \cdot L_{h,t}^{DS} + \lambda \cdot L_{h,t}^{DS-HA} \\ &= -I_{HA} \cdot \log p_r^{HA} - (1 - I_{HA}) \log p_r^{DS} + \lambda \left[\frac{1}{2} \left(\frac{\ell_r - \mu_r}{\sigma_r} \right)^2 + \ell_r + \log \sigma_r \right] \end{aligned} \quad (4)$$

where $\ell_r = \log p_r^{DS} - \log p_r^{HA}$, and I_{HA} is 1 if the label is from HA data and 0 otherwise.

3.6 Analysis of the Disagreement Penalty

Let \mathbf{w}_{HA} be a learnable parameter of *HA-Net* which predicts relations in the test time. We investigate the effect of the disagreement penalty by comparing the gradients of loss functions with respect to \mathbf{w}_{HA} for a human annotated label and a distantly supervised label.

For a label $\langle e_h, r, e_t \rangle$, let $\phi_r = (\log(p_r^{DS}/p_r^{HA}) - \mu_r)/\sigma_r^2$. If the label is human annotated, we obtain the following gradient of the loss $L_{h,t}$ with respect to \mathbf{w}_{HA} from Equation (4)

$$\nabla L_{h,t} = \nabla L_{h,t}^{HA} + \mathbf{0} + \lambda \nabla L_{h,t}^{DS-HA} = -(1 + \lambda(1 + \phi_r)) \frac{1}{p_r^{HA}} \nabla p_r^{HA}. \quad (5)$$

On the other hand, if the label is annotated by distant supervision, the gradient becomes

$$\nabla L_{h,t} = \mathbf{0} + \mathbf{0} + \lambda \nabla L_{h,t}^{DS-HA} = -\lambda(1 + \phi_r) \frac{1}{p_r^{HA}} \nabla p_r^{HA}. \quad (6)$$

The two gradients in Equations (5) and (6) have the same direction of $-\nabla p_r^{HA}$. It implies that a human annotated label and a distantly supervised label have similar effects on training *HA-Net* except that the magnitudes of gradients are calibrated by $1 + \lambda(1 + \phi_r)$ and $\lambda(1 + \phi_r)$, respectively. Thus, *HA-Net* can learn from not only human annotated labels but also distantly supervised labels by introducing the disagreement penalty. Recall that the log-normal distribution $L(\mu_r, \sigma_r)$ describes the conditional inflation for a given sentence with a head entity and a tail entity. If the median e^{μ_r} of $L(\mu_r, \sigma_r)$ has a high value, the distantly supervised label is likely to be a false label. Thus, we decrease the size of ϕ_r to reduce the effect of a distantly supervised label. On the other hand, as the median e^{μ_r} becomes lower, the size of ϕ_r increases to aggressively utilize the distantly supervised label.

3.7 Extension to Document-level Relation Extraction

For the document-level RE, we use the *binary* cross entropy as the prediction losses $L_{h,t}^{HA}$ and $L_{h,t}^{DS}$. For a pair of entities e_h and e_t , let $R_{h,t}$ be the set of relation types between the entities. In the train time, we use the following loss function for document relation extraction

$$L_{h,t} = -I_{HA} \left(\sum_{r \in R_{h,t}} \log p_r^{HA} + \sum_{r \in R \setminus R_{h,t}} \log (1 - p_r^{HA}) \right) - (1 - I_{HA}) \left(\sum_{r \in R_{h,t}} \log p_r^{DS} + \sum_{r \in R \setminus R_{h,t}} \log (1 - p_r^{DS}) \right) + \lambda \sum_{r \in R_{h,t}} \left[\frac{1}{2} \left(\frac{\ell_r - \mu_r}{\sigma_r} \right)^2 + \ell_r + \log \sigma_r \right].$$

where $\ell_r = \log p_r^{DS} / p_r^{HA}$, and I_{HA} is 1 if the labels are from HA data and 0 otherwise. We obtain the same property shown in Section 3.6 for the above loss function. In the test time, we regard that the model outputs the triple $\langle e_h, r, e_t \rangle$ if p_r^{HA} is greater than a threshold which is tuned on the development dataset.

4 Experiments

We conducted a performance study for sentence-level and document-level REs by following the experimental settings of (Ye et al., 2019) and (Yao et al., 2019; Wang et al., 2019), respectively. All models are implemented in PyTorch and trained on a V100 GPU. We initialized *HA-Net* and *DS-Net* to have the same initial parameters. More experimental details including implementations can be found in Appendix A.

4.1 Experimental Settings

Dataset. KBP (Ling and Weld, 2012; Ellis, 2012) and NYT (Riedel et al., 2010; Hoffmann et al., 2011) are datasets for sentence-level RE, and DocRED (Yao et al., 2019) is a dataset for document-level RE. The statistics of the datasets are summarized in Table 2. Since KBP and NYT do not have HA train data, we use 20% of the HA test data as the HA train data. In addition, we randomly split 10% of train data on KBP and NYT for the development (dev) data. Note that the ground truth of the test data in DocRED is not publicly available. However, we can get the F1 score of the result extracted from the test data by submitting the result to the DocRED competition hosted by CodaLab (available at <https://competitions.codalab.org/competitions/20717>). We report both the F1 scores computed from the dev data and the test data.

Data	Number of instances				# of rel. types
	Train-HA	Train-DS	Dev	Test	
KBP	378	132,369	14,103	1,488	7
NYT	756	323,126	34,871	3,021	25
DocRED	38,269	1,508,320	12,332	12,842	96

Table 2: Statistics of datasets

Dataset RE models	KBP						NYT					
	<i>BiGRU_S</i>	<i>PaLSTM_S</i>	<i>BiLSTM_S</i>	<i>PCNN_S</i>	<i>CNN_S</i>	<i>BERT_S</i>	<i>BiGRU_S</i>	<i>PaLSTM_S</i>	<i>BiLSTM_S</i>	<i>PCNN_S</i>	<i>CNN_S</i>	<i>BERT_S</i>
<i>HA-Only</i>	0.1984	0.1153	0.1787	0.3410	0.2586	0.1631	0.0884	0.1259	0.1504	0.4463	0.3978	0.1953
<i>DS-Only</i>	0.3909	0.3521	0.3519	0.2705	0.2810	0.3610	0.4532	0.4429	0.4297	0.4177	0.4463	0.4625
<i>BASet</i>	0.3972	0.4055	0.4053	0.2410	0.2400	0.3858	0.4966	0.4555	0.4561	0.3584	0.4358	0.5081
<i>BAFix</i>	0.4241	0.4027	0.3581	0.2931	0.2473	0.3383	0.4613	0.4507	0.4707	0.4023	0.4532	0.5145
<i>MaxThres</i>	0.4264	0.3630	0.4053	0.2815	0.2645	0.3751	0.4531	0.4462	0.4350	0.4258	0.4655	0.4952
<i>EntThres</i>	0.4470	0.4018	0.4248	0.2925	0.2826	0.3539	0.4553	0.4472	0.4210	0.4154	0.4427	0.4940
<i>DUAL</i>	0.4749	0.4420	0.4207	0.3872	0.2969	0.4013	0.5455	0.5210	0.4524	0.4986	0.4744	0.5300

Table 3: Sentence-level RE datasets (KBP and NYT)

RE models	Dev					Test				
	<i>BERT_D</i>	<i>BiLSTM_D</i>	<i>CA_D</i>	<i>LSTM_D</i>	<i>CNN_D</i>	<i>BERT_D</i>	<i>BiLSTM_D</i>	<i>CA_D</i>	<i>LSTM_D</i>	<i>CNN_D</i>
<i>HA-Only</i>	0.5513	0.4992	0.4986	0.4817	0.4788	0.5478	0.4982	0.4992	0.4815	0.4681
<i>DS-Only</i>	0.4683	0.4951	0.4890	0.4877	0.4166	0.4587	0.4809	0.4772	0.4713	0.4160
<i>BASet</i>	0.4807	0.5123	0.5024	0.5012	0.4349	0.4716	0.4949	0.4905	0.4905	0.4320
<i>BAFix</i>	0.4802	0.5136	0.5070	0.5166	0.4365	0.4730	0.5061	0.4989	0.4977	0.4354
<i>DUAL</i>	0.5880	0.5510	0.5372	0.5392	0.4967	0.5774	0.5379	0.5306	0.5277	0.4909

Table 4: Document-level RE dataset (DocRED)

Compared methods. We compare our dual supervision framework, denoted by *DUAL*, with the state-of-the-art methods *BASet* and *BAFix* in (Ye et al., 2019). For sentence-level RE, we compare *DUAL* with two additional baselines *MaxThres* (Ren et al., 2017) and *EntThres* (Liu et al., 2017) which are only applicable to multi-class classification and cannot be used in document-level RE. *MaxThres* outputs NA if the maximum output probability is less than a threshold. Similarly, *EntThres* outputs NA if the entropy of the output probability distribution is greater than a threshold.

Used relation extraction models. For *sentence-level RE*, we used the six models: *BiGRU_S* (Zhang et al., 2017), *PaLSTM_S* (Zhang et al., 2017), *BiLSTM_S* (Zhang et al., 2017), *CNN_S* (Zeng et al., 2014), *PCNN_S* (Zeng et al., 2015) and *BERT_S* (Wang et al., 2019). On the other hand, for *document-level RE*, we used the five models: *BERT_D* (Wang et al., 2019), *CNN_D* (Zeng et al., 2014), *LSTM_D* (Yao et al., 2019), *BiLSTM_D* (Cai et al., 2016) and *CA_D* (Sorokin and Gurevych, 2017). Note that *CNN_D*, *BiLSTM_D*, and *CA_D* are originally proposed for sentence-level RE and we used the adaptation of them to document-level RE by Yao et al. (2019). In addition, we adapt *BERT_D* to the sentence-level RE by changing the output activation function from sigmoid to softmax and denote it by *BERT_S*.

4.2 Comparison with Existing Methods

We compare the dual supervision framework with the existing methods.

Sentence-level RE. Table 3 shows F1 scores for relation extraction on KBP and NYT. Note that *DS-Only* and *HA-Only* represent the original RE models trained only on distantly supervised and human-annotated labels, respectively. *DUAL* shows the highest F1 scores with all RE models except *BiLSTM_S*. Since KBP and NYT have a small number of human-annotated labels in train data, *HA-Only* shows worse F1 scores than *DS-Only*. Furthermore, *DUAL* achieves improvements of F1 score from 5% to 40% over *DS-Only* by additionally using the small amount of human annotated labels. On the other hand, the compared methods *BAFix*, *BASet*, *MaxThres* and *EntThres* often perform worse than *DS-Only* and *HA-Only*.

Document-level RE. We present F1 scores on DocRED in Table 4. *DUAL* outperforms *BASet* and *BAFix* with all RE models. Especially, the F1 score of dual framework with *BERT_D* shows more than 22% of improvement over *BASet* and *BAFix*. Since DocRED has a large human-annotated train data, *HA-Only* shows better performance than *DS-Only*. For *BERT_D* and *CNN_D*, the existing methods show lower F1 scores compared to *HA-Only*. It shows that the accuracy can be degraded although we use additional DA data in addition to HA data due to the labeling bias. Meanwhile, we achieve a consistent and significant improvement by applying *DUAL*. In the rest of this paper, we will provide a detailed evaluation of performance on DocRED data which is the largest dataset in this experiment. For the test

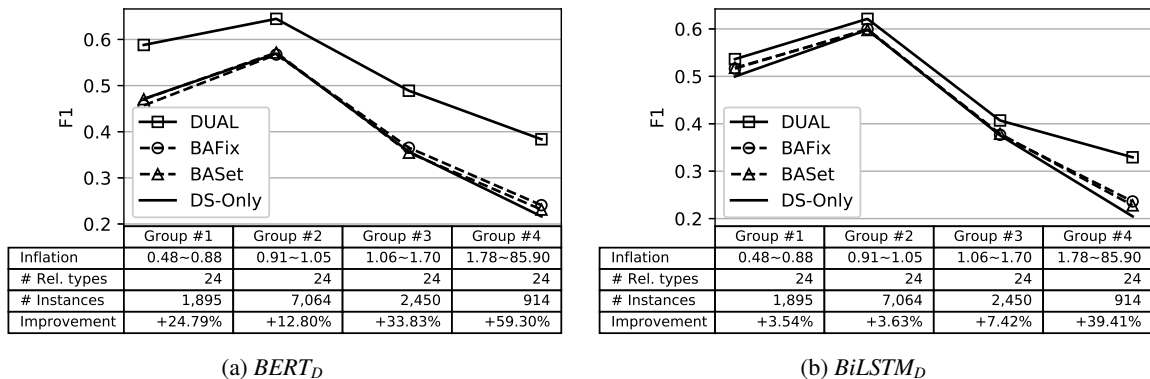


Figure 2: F1 scores of different groups

data of DocRED, the ground truth is not publicly available and only a F1 score can be obtained from the DocRED competition. Thus, we provide detailed evaluations of performance on the dev data only.

Inflation vs. accuracy. To investigate the effect of the inflation to the accuracy of relation extraction, we split the relation types into 4 groups based on the inflation of the relation types. In Figure 2, we present the characteristics of each group and plot the F1 scores by groups for $BERT_D$ model and $BiLSTM_D$ model. All methods have the highest F1 scores when the inflation is close to 1 (at the 2nd group). Furthermore, the improvement of F1 score by $DUAL$ compared to the second best performer increases as the inflation moves away from 1. Thus, it confirms that our dual supervision framework effectively utilizes both human annotation and distant supervision by modeling the bias of the distant supervision. Since the other models CA_D , $LSTM_D$ and CNN_D show similar results with $BiLSTM_D$, we omit the result.

4.3 Ablation Study

We conducted an ablation study with the existing model $BERT_D$ on DocRED to validate the effectiveness of individual components of our framework. We compared $DUAL$ (separate prediction networks + disagreement penalty) and two variations of our framework $Multitask$ (separate prediction networks only) and $Single$. $Multitask$ denotes a variation of $DUAL$ which does not utilize the disagreement penalty, while $BERT_D$ without applying the dual supervision framework is referred to as $Single$. Note that $Single$ is also trained on both HA data and DS data together.

To show the effectiveness of the components depending on the size of HA data, we plotted the F1 scores with varying the number of human-annotated document from 152 to 3,053 (i.e., from 5% to 100% of the documents with HA). As we expected, $DUAL$ outperforms both variations all the time. Furthermore, separation of the prediction networks significantly improves the accuracy when we have enough number of human-annotated labels. However, when we use less than 10% of the human annotated documents, $Multitask$ suffers from the sparsity problem. By utilizing the disagreement penalty additionally, $DUAL$ outperforms $Single$ even when we use only 5% of the human-annotated documents for training the model. It implies that the disagreement penalty enables $HA-Net$ to effectively learn from DS data as well as HA data.




Figure 3: Varying the size of HA data

4.4 Quality Comparison

To give an idea of what false relations are found by existing methods, we provide two example documents in the dev data of DocRED and the relations extracted by $DUAL$, $BAFix$ and $DS-Only$ with $BERT_D$ in Table 5. The relation $\langle Sweden, capital, Stockholm \rangle$ is expressed in the document titled ‘Kungliga Hovkapellet’ and all methods find the relation correctly. In the document titled ‘Loopline Bride’, the relation $\langle Ireland, capital, Dublin \rangle$ does not exist. However, $BAFix$ and $DS-Only$ output the incorrect

Document	Title: Kungliga Hovkapellet [1] Kungliga Hovkapellet is a Swedish orchestra, originally part of the Royal Court in [Sweden]’s capital [Stockholm]. [2] Its existence ...	Title: Loopline Bridge [1] The Loopline Bridge (or the Liffey Viaduct) is a railway bridge spanning the River Liffey and several streets in [Dublin], [Ireland]. [2] It joins ...
Relations	True label: ⟨Sweden, capital, Stockholm⟩ <i>DUAL</i> : ⟨Sweden, capital, Stockholm⟩ <i>BAFix</i> : ⟨Sweden, capital, Stockholm⟩ <i>DS-Only</i> : ⟨Sweden, capital, Stockholm⟩	True label: NA <i>DUAL</i> : NA <i>BAFix</i> : ⟨Ireland, capital, Dublin⟩ <i>DS-Only</i> : ⟨Ireland, capital, Dublin⟩

Table 5: Examples of documents and extracted relations

relation. Since *DUAL* adaptively assess the labeling bias with μ -Net and σ -Net, *DUAL* does not output the false relation. In addition, since the RE models trained with *BAFix* and *DS-Only* fail to learn the text pattern corresponding to the relation type due to the labeling bias, they output many false labels such as ⟨Vietnam, capital, Taipei⟩ in many documents. It shows that the dual supervision framework effectively deal with the labeling bias of distant supervision by considering contextual information.

4.5 Topic-aware RE

Topic-aware RE is a special case of document-level RE to extract the relations between the topic entity of a document and the other entities. Jung and Shim (2020) proposed a topic-aware relation extraction (T-REX) model which is robust to the omitted mentions of topic entities in documents. We apply our dual supervision framework to the T-REX on DocRED dataset and report the result in Table 6. The result shows that our dual supervision framework is also effective in the topic-aware RE task.

	F1	AUC
<i>HA-Only</i>	0.6569	0.6456
<i>DS-Only</i>	0.6624	0.6978
<i>DUAL</i>	0.6930	0.7125

Table 6: Topic-aware RE

5 Related Works

We briefly survey the existing works for RE. Mintz et al. (2009) propose distant supervision to overcome the limitation of the quantity of human-annotated labels. They utilize lexical, syntactic and named entity tag features obtained by existing NLP tools to extract relations. Other early works in (Riedel et al., 2010; Hoffmann et al., 2011) also utilized hand-crafted features to find the relations in text. However, since such RE models take the input features from NLP tools, the errors generated by the NLP tools are propagated to the RE models. In order to deal with the error propagation, as we discussed in Section 2.2, the works (Lin et al., 2016; Zeng et al., 2014; Zeng et al., 2015; Sorokin and Gurevych, 2017; Wang et al., 2019) use deep neural networks such as CNN, LSTM and BERT instead of handcrafted features to encode the text for finding the relations. Since many relational facts are expressed across multiple sentences, the recent works (Yao et al., 2019; Wang et al., 2019) studied document-level RE. Yao et al. (2019) provide a document-level RE dataset (DocRED) as well as compare the models adapted from the sentence-level RE models (Zeng et al., 2014; Hochreiter and Schmidhuber, 1997; Cai et al., 2016; Sorokin and Gurevych, 2017). Moreover, a fine tuned model (Wang et al., 2019) of BERT (Devlin et al., 2018) for document-level RE achieved a higher F1 score than the baselines on DocRED.

The wrong labeling problem in distant supervision has been addressed in many previous works (Zeng et al., 2015; Lin et al., 2016; Ye and Ling, 2019; Beltagy et al., 2018). Among them, Zeng et al. (2015), Lin et al. (2016) and Ye and Ling (2019) build a bag-of-sentences for a pair of entities and extract relational facts from the bag-of-sentences with attention over the sentences. Beltagy et al. (2018) propose a bag-of-sentences-level model which utilizes human annotation. However, they use the human annotated labels only to determine whether there exists a relationship or not since the labels are obtained from a different domain. The goal of these works is different from ours which is to find the relations *appearing in a given text* (e.g., a document). Thus, the bag-of-sentences-level models have a limitation to be used for some applications such as question answering.

The most relevant work to ours is (Ye et al., 2019). This paper proposes the bias adjustment methods to utilize a small amount of HA data to improve RE models trained on DS data by considering the different distribution of human annotated labels and distantly supervised labels. However, they do not use HA data to train the models and use the HA data only to obtain a statistic to be used to determine the size of the bias adjustment. Thus, the bias adjustment methods cannot consider contextual information.

6 Conclusion

We proposed the dual supervision framework to utilize human annotation and distant supervision based on the analysis of labeling bias in distant supervision. We devised a new structure for the output layer of RE models that consists of 4 sub networks. The new structure is robust to the noisy labeling of distant supervision since the labels obtained by human annotation and distant supervision are predicted by separate prediction networks *HA-Net* and *DS-Net*, respectively. In addition, we introduced an additional loss term called *disagreement penalty* which enables *HA-Net* to learn from distantly supervised labels. The parameter networks μ -Net and σ -Net adaptively assess the labeling bias by considering contextual information. Moreover, we theoretically analyzed the effect of the disagreement penalty. Our experiments showed that the dual supervision framework significantly outperforms the existing methods.

Acknowledgements

This work was supported by Next-Generation Information Computing Development Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Science, ICT (No. NRF-2017M3C4A7063570) and was also supported by Institute of Information & communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No. 2020-0-00857, Development of cloud robot intelligence augmentation, sharing and framework technology to integrate and enhance the intelligence of multiple robots). This research was results of a study on the "HPC Support" Project, supported by the Ministry of Science and ICT and NIPA.

References

- M. Ahmed, J. Islam, M. R. Samee, and R. E. Mercer. 2019. Identifying protein-protein interaction using tree lstm and structured attention. In *2019 IEEE 13th International Conference on Semantic Computing (ICSC)*, pages 224–231, Jan.
- Iz Beltagy, Kyle Lo, and Waleed Ammar. 2018. Combining distant and direct supervision for neural relation extraction. *arXiv preprint arXiv:1810.12956*.
- Rui Cai, Xiaodong Zhang, and Houfeng Wang. 2016. Bidirectional recurrent convolutional neural network for relation classification. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 756–765.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Xin Dong, Evgeniy Gabilovich, Jeremy Heitz, Wilko Horn, Ni Lao, Kevin Murphy, Thomas Strohmann, Shaohua Sun, and Wei Zhang. 2014a. Knowledge vault: A web-scale approach to probabilistic knowledge fusion. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 601–610.
- Xin Luna Dong, Evgeniy Gabilovich, Jeremy Heitz, Wilko Horn, Kevin Murphy, Shaohua Sun, and Wei Zhang. 2014b. From data fusion to knowledge fusion. *PVLDB*, 7(10):881–892.
- Charles Dugas, Yoshua Bengio, François Bélisle, Claude Nadeau, and René Garcia. 2001. Incorporating second-order functional knowledge for better option pricing. In *Advances in neural information processing systems*, pages 472–478.
- Joe Ellis. 2012. Linguistic resources for 2013 knowledge base population evaluations.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.

- Raphael Hoffmann, Congle Zhang, Xiao Ling, Luke Zettlemoyer, and Daniel S Weld. 2011. Knowledge-based weak supervision for information extraction of overlapping relations. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 541–550. Association for Computational Linguistics.
- WooHwan Jung and Kyuseok Shim. 2020. T-rer: A topic-aware relation extraction model. In *Proceedings of the 29th ACM International Conference on Information Knowledge Management, CIKM '20*, page 20732076, New York, NY, USA. Association for Computing Machinery.
- WooHwan Jung, Younghoon Kim, and Kyuseok Shim. 2019. Crowdsourced truth discovery in the presence of hierarchies for knowledge fusion. In *Advances in Database Technology - 22nd International Conference on Extending Database Technology, EDBT 2019, Lisbon, Portugal, March 26-29, 2019*, pages 205–216.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Yankai Lin, Shiqi Shen, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2016. Neural relation extraction with selective attention over instances. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2124–2133.
- Xiao Ling and Daniel S Weld. 2012. Fine-grained entity recognition. In *Twenty-Sixth AAAI Conference on Artificial Intelligence*.
- Liyuan Liu, Xiang Ren, Qi Zhu, Shi Zhi, Huan Gui, Heng Ji, and Jiawei Han. 2017. Heterogeneous supervision for relation extraction: A representation learning approach. *arXiv preprint arXiv:1707.00166*.
- Frank J Massey Jr. 1951. The kolmogorov-smirnov test for goodness of fit. *Journal of the American statistical Association*, 46(253):68–78.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 2-Volume 2*, pages 1003–1011. Association for Computational Linguistics.
- Xiang Ren, Zeqiu Wu, Wenqi He, Meng Qu, Clare R Voss, Heng Ji, Tarek F Abdelzaher, and Jiawei Han. 2017. Cotype: Joint extraction of typed entities and relations with knowledge bases. In *Proceedings of the 26th International Conference on World Wide Web*, pages 1015–1024.
- Sebastian Riedel, Limin Yao, and Andrew McCallum. 2010. Modeling relations and their mentions without labeled text. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 148–163. Springer.
- Daniil Sorokin and Iryna Gurevych. 2017. Context-aware representations for knowledge base relation extraction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1784–1789.
- Hong Wang, Christfried Focke, Rob Sylvester, Nilesh Mishra, and William Wang. 2019. Fine-tune bert for docred with two-step process. *arXiv preprint arXiv:1909.11898*.
- Kun Xu, Siva Reddy, Yansong Feng, Songfang Huang, and Dongyan Zhao. 2016. Question answering on freebase via relation extraction and textual evidence. *arXiv preprint arXiv:1603.00957*.
- Yuan Yao, Deming Ye, Peng Li, Xu Han, Yankai Lin, Zhenghao Liu, Zhiyuan Liu, Lixin Huang, Jie Zhou, and Maosong Sun. 2019. Docred: A large-scale document-level relation extraction dataset. *arXiv preprint arXiv:1906.06127*.
- Zhi-Xiu Ye and Zhen-Hua Ling. 2019. Distant supervision relation extraction with intra-bag and inter-bag attentions. *arXiv preprint arXiv:1904.00143*.
- Qinyuan Ye, Liyuan Liu, Maosen Zhang, and Xiang Ren. 2019. Looking beyond label noise: Shifted label distribution matters in distantly supervised relation extraction. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3832–3841.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. Relation classification via convolutional deep neural network. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2335–2344.

- Daojian Zeng, Kang Liu, Yubo Chen, and Jun Zhao. 2015. Distant supervision for relation extraction via piecewise convolutional neural networks. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1753–1762.
- Yuhao Zhang, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D Manning. 2017. Position-aware attention and supervised data improve slot filling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 35–45.

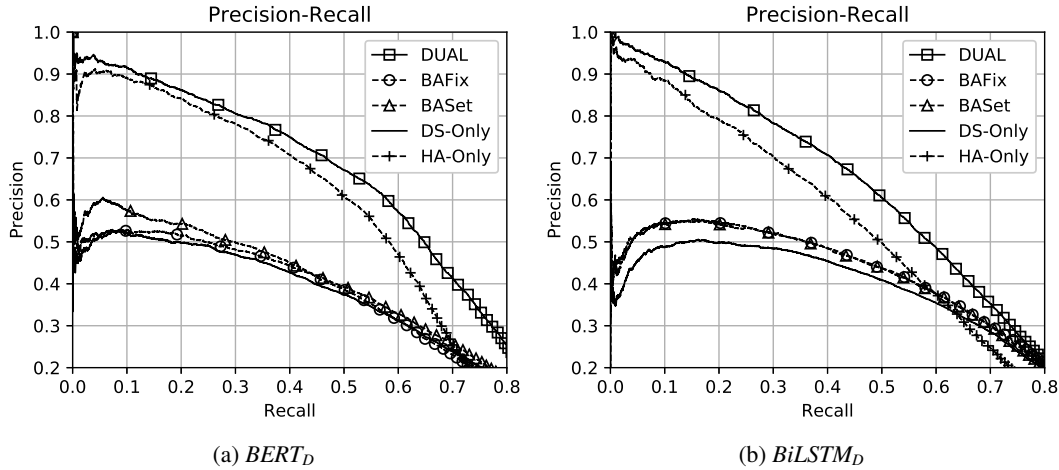


Figure 4: Precision-recall curves

Appendix

A Experimental Details

Our implementation is available at <https://github.com/woohwanjung/dual>.

Document-level RE. For $BiLSTM_D$, $LSTM_D$, CA_D and CNN_D , we utilized the code which is available at <https://github.com/thunlp/DocRED> and implemented by Yao et al. (2019). In addition, we used the implementation of $BERT_D$ that is available at <https://github.com/hongwang600/DocRed> and provided by Wang et al. (2019). We used Adam optimizer (Kingma and Ba, 2014) to optimize the RE models. For the $BERT_D$ model, we set the batch size to 12 and learning rate to 10^{-5} . For the other models, we followed the setting provided in (Yao et al., 2019): batch size is 40, learning rate is 10^{-3} . We set the hyperparameters λ and d to 10^{-5} and 128, respectively. Each training batch has half of the instances with human-annotated labels and the other half of instances with distantly supervised labels.

Sentence-level RE. We use the code which is made publicly available by Ye et al. (2019) at <https://github.com/INK-USC/shifted-label-distribution>. All models except $BERT_S$ are trained by stochastic gradient descent. Learning rate is initially set to 1.0, and decreased to 10% if there is no improvement on the dev data for 3 consecutive epochs. For the models, we set the hyperparameters λ and d to 10^{-3} and 200, respectively. To train $BERT_S$ model, we used Adam optimizer with learning rate 10^{-5} . Moreover, the hyperparameters λ and d are set to 10^{-4} and 128, respectively. We alternately used an HA batch and a DS batch for dual supervision where an HA batch consists of training instances with human annotated labels and a DS batch consists of training instances with distantly supervised labels.

B Additional Experiments

The precision-recall curves of the compared methods are shown in Figure 4. As expected, $DUAL$ consistently outperforms all compared methods. $BAFix$ and $BASet$ have similar precision-recall curves with $DS-Only$. Although $HA-Only$ shows comparable precisions with $DUAL$ when recall is low, the precision of $HA-Only$ drops faster than that of $DUAL$ with increasing recall. It implies that human annotated labels are not enough for training a model to extract a large number of relations. Meanwhile, $DUAL$ extracts more relations from the document compared to existing models at the same precision level.