

Finding Influential Instances for Distantly Supervised Relation Extraction

Zifeng Wang^{1,2}, Rui Wen³, Xi Chen^{3*}, Shao-Lun Huang², Ningyu Zhang⁴, Yefeng Zheng³

¹University of Illinois Urbana Champaign ²TBSI, Tsinghua University

³Tencent ⁴Zhejiang University

Abstract

Distant supervision (DS) is a strong way to expand the datasets for enhancing relation extraction (RE) models but often suffers from high label noise. Current works based on attention, reinforcement learning, or GAN are black-box models so they neither provide meaningful interpretation of sample selection in DS nor stability on different domains. On the contrary, this work proposes a novel model-agnostic instance sampling method for DS by influence function (IF), namely REIF. Our method identifies favorable/unfavorable instances in the bag based on IF, then does dynamic instance sampling. We design a fast influence sampling algorithm that reduces the computational complexity from $\mathcal{O}(mn)$ to $\mathcal{O}(1)$, with analyzing its robustness on the selected sampling function. Experiments show that by simply sampling the favorable instances during training, REIF is able to win over a series of baselines which have complicated architectures. We also demonstrate that REIF can support interpretable instance selection.

1 Introduction

To expand the training data for relation extraction (RE), distant supervision (DS) was proposed by Mintz et al. (2009) who assumed that if two entities are related in existing KBs, then all sentences contain both of them express this relation. However, this heuristic inevitably suffers from wrong labels (Takamatsu et al., 2012) and undermines model performance. For example, the sentence “*Bill Gates redefined the software industry, ... said Rob Glaser, a former Microsoft executive*” does not mention the relation *founder* but is still treated as a positive training sample in DS. Dealing with noisy instances in DS has been a focus in RE. There are three main genres in the literature: (1) incorporating an attention module (Lin et al., 2016) to allocate confidence level among instances in the same bag;

* Corresponding author: jasonxchen@tencent.com

Bag Label	Instances	Influences ϕ	Probability π
/business /company /founders (Bill Gates, Microsoft)	<i>Bill Gates</i> , the founder of <i>Microsoft</i> , has indicates ...	-0.8	0.7
	... as an investor, Allen, who founded <i>Microsoft</i> with <i>Bill Gates</i> , has ...	-0.2	0.5
	<i>Bill Gates</i> redefined the software industry ... said Rob Glaser, a former <i>Microsoft</i> executive...	0.1	0.2

Figure 1: Finding influential instances within a bag via subsampling based on the calculated probability π . Note that here *negative* ϕ means a beneficial sample.

(2) using reinforcement learning (Qin et al., 2018b) for instance selection; and (3) leveraging adversarial training (Wu et al., 2017) to enhance the RE model’s robustness against noise. However, they are either black-box models thus unable to provide meaningful interpretation of sample selection or sensitive to datasets. More importantly, none of them is theoretically guaranteed to truly reduce the “noise” from the dataset.

In this work, we propose to leverage influence function (IF) to evaluate instance quality then do instance selection for DS. Influence function is a powerful tool drawn from robust statistics (Huber, 2004). It is able to approximate the influence of a single data point on the whole model learned on the dataset. Creating to this merit, it has been successfully utilized for inspecting outliers (Boente et al., 2002) and denoising datasets (Wang et al., 2020) based on shallow machine learning models, e.g., logistic regression. Although Koh & Liang (2017) extends IF to interpreting deep networks, it is still elusive if it works for denoising datasets for deep networks. In this work, we develop the Relation Extraction by InFLuence subsampling (REIF) framework, which aims for denoising DS for deep learning RE models.

The high-level idea of REIF is shown by Fig. 1. Each instance is assigned a quality measure ϕ , from which its sampling probability is obtained via the sampling function π . Accordingly, the better

an instance’s quality is, the more likely it is picked during training. We will explain the operational meaning of ϕ in Section 3.2. In a nutshell, the main contributions of this paper are

- We develop a novel IF-based denoising framework for DS RE, namely REIF, for denoising RE by sampling favorable training instances.
- An efficient implementation of REIF enables subsampling in $\mathcal{O}(1)$ complexity, instead of the $\mathcal{O}(mn)$ complexity without our implementation.
- Empirical experiments show REIF’s superiority over other baselines, and we identify its capability to support interpretable instance selection for RE by a case study.¹

2 Related Work

There are a series of works trying to address the noisy label difficulty in DS by multi-instance learning (MIL) (Hoffmann et al., 2011; Riedel et al., 2010; Surdeanu et al., 2012). MIL considers the training labels in *bag* level instead of instance level. Each bag contains at least one instance with the labeled relation while the exact label of each instance is unknown. As MIL being proved effective in relation extraction, it was firstly introduced to neural relation extraction by Zeng et al. (2015), where the piece-wise convolutional neural network (PCNN) was developed, and only one instance with the largest predicted probability was selected in each bag.

Later, attention (Lin et al., 2016; Zhou et al., 2018; Jia et al., 2019; Yuan et al., 2019; Ye and Ling, 2020; Zhou et al., 2021), reinforcement learning (Feng et al., 2018; Yang et al., 2018; Qin et al., 2018b; Chen et al., 2021), and adversarial training (Wu et al., 2017; Qin et al., 2018a; Han et al., 2018; Shi et al., 2018) have been proposed for further improvement. However, above works usually require intense trials in fine-tuning of the hyper-parameters in practice, or are not interpretable to human-beings. In this work, we propose a model-agnostic and interpretable instance selection method via IF, which is easy-to-use for most DL models without many hyperparameters to choose.

¹Code is available in the supplementary materials.

3 Methodology

In this section, we elaborate on the major steps of REIF associated with the technical details and the theoretic foundation of measuring data quality by influences. Also, an analysis supporting our choice of sampling function is given.

3.1 Relation Extraction by Influence Subsampling

Our REIF is model-agnostic thus amenable to most DL models. Without loss of generality, we pick PCNN (Zeng et al., 2015) as the encoder for the input texts. The flowchart of our framework is shown in Fig. 2. It includes three main parts: 1) backbone model and 2) instance selection.

Backbone Model. Inputs of the encoder are raw sentences represented by indices of words, e.g., a sentence x_* with l words $x_* = \{x_{*,1}, \dots, x_{*,l}\}$. We transform them into dense real-valued representation vectors as $\mathbf{w}_* = \{\mathbf{w}_{*,1}, \dots, \mathbf{w}_{*,l}\}$, by concatenating the word embedding from $\mathbf{V} \in \mathbb{R}^{d^a \times |V|}$ (where $|V|$ denotes the size of the vocabulary and d^a is the dimension of word embedding) and position embedding with dimension d^p together. As there are two position embeddings, each word vector in \mathbf{w} has dimension $d^a + 2 \times d^p$. Convolution layer processes the word representations as

$$\mathbf{x}_* = \text{CNN}(\mathbf{w}_*). \quad (1)$$

The CNN model receives representation vectors \mathbf{w}_* and outputs the processed feature vectors $\mathbf{x}_* \in \mathbb{R}^{d \times l}$. The probability for relation prediction, taking \mathbf{x}_* as input, is given by

$$P(y = k | \mathbf{x}_*) = \frac{\exp(\boldsymbol{\beta}^{(k)\top} \mathbf{x}_*)}{\sum_{k'} \exp(\boldsymbol{\beta}^{(k')\top} \mathbf{x}_*)}, \quad (2)$$

where $\boldsymbol{\beta} = \{\boldsymbol{\beta}^{(1)} \dots \boldsymbol{\beta}^{(K)}\} \in \mathbb{R}^{d \times K}$ is the weight matrix of the last fully-connected layer; K is the total number of relations.

Dynamic Instance Sampling. One possible way to do sample selection by IF is *post-hoc*, i.e., it first samples from the full training set, then retrains the model on the subsamples. However, we argue it is unsuitable for DS. In post-hoc sampling, all instances are gathered together, hence the subsamples are dominated by majority relations with lots of training instances, resulting in severe class imbalance. In an extreme case, minority relations may completely disappear after subsampling.

On contrast, we propose dynamic instance sampling (DIS) which is executed within bags during

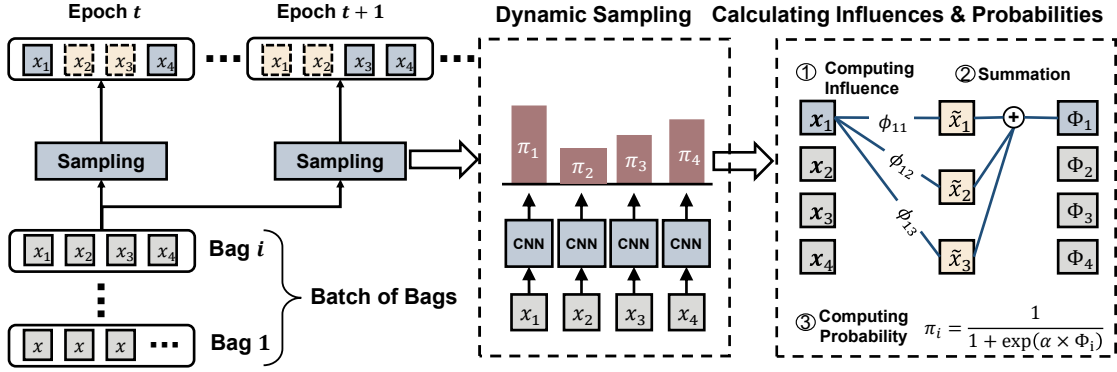


Figure 2: The flowchart of the instance-level subsampling method, where x is training sentence; \tilde{x} is the validation sample; ϕ is the computed influence; and a dotted box means the instance is dropped after subsampling.

training. Given a bag $X = \{x_1, \dots, x_n\}$ containing n sentences, we try to sample a subset X_{sub} with $|X_{sub}| < n$ from X . To this end, we calculate the influences $\Phi_i, \forall i = 1, \dots, n$, and sampling probabilities π_i are

$$\pi_i = \pi(\Phi_i) := \frac{1}{1 + \exp(\alpha \times \Phi_i)}, \quad (3)$$

where π_i is the probability of x_i being selected and α is a hyper-parameter. Consequently, the training objective function $J(\theta)$ is

$$J(\theta) = \frac{1}{|X_{sub}|} \sum_{x_i \in X_{sub}} \ell_i(\theta), \quad (4)$$

where $\ell(\theta)$ is the abbreviation of loss function $\ell(x, y; \theta)$ for notation simplicity.

3.2 Theoretic Foundation of Influence-based Sample Quality Measure

The core step of REIF is to measure the instance influence Φ . Intuitively, adverse instances, which cause model validation loss increasing, should be assigned low probability being sampled, and vice versa. We next present the property of Φ and substantiate this intuition in a rigorous way.

Consider a classification problem where we attempt to obtain a model $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$, which is parametrized by θ , that can make prediction from an input space \mathcal{X} (e.g., sentences) to an output space \mathcal{Y} (e.g., relations). Given a set of training data $\{x_i\}_{i=1}^n$ and the corresponding labels $\{y_i\}_{i=1}^n$, the optimal $\hat{\theta}$ defined by

$$\hat{\theta} := \arg \min_{\theta \in \Theta} \frac{1}{n} \sum_{i=1}^n \ell_i(\theta). \quad (5)$$

We evaluate the learned $f_{\hat{\theta}}$ on an additional valida-

tion set $\{(x_j^v, y_j^v)\}_{j=1}^m$ such as

$$L(\hat{\theta}) := \frac{1}{m} \sum_{j=1}^m \ell_j^v(\hat{\theta}) \quad (6)$$

where $\ell_j^v(\hat{\theta})$ is the validation loss on x_j^v .

In order to quantitatively measure the i -th training sample's influence over model's validation loss, we can perturb the training loss $\ell_i(\theta)$ by a small ϵ , then retrain a perturbed risk minimizer $\tilde{\theta}$ as

$$\tilde{\theta} := \arg \min_{\theta \in \Theta} \frac{1}{n} \sum_{i'=1}^n \ell_{i'}(\theta) + \epsilon \times \ell_i(\theta). \quad (7)$$

As a result, we are able to compute the validation loss change of the validation sample x_j^v by

$$\delta_j(\epsilon) := \ell_j^v(\tilde{\theta}) - \ell_j^v(\hat{\theta}). \quad (8)$$

It indicates to what extent x_i influences the prediction on x_j^v . If $\epsilon = -1/n$, according to Eq. (7), x_i 's loss $\ell_i(\theta)$ is actually removed from the objective function. In this situation, $\delta_j(\epsilon) > 0$, i.e., $\ell_j^v(\tilde{\theta}) - \ell_j^v(\hat{\theta}) > 0$, implies that removing x_i causes the validation loss on x_j^v increasing, i.e.,

$$\delta_j \left(-\frac{1}{n} \right) > 0 \rightarrow x_i \text{ is good for } x_j^v. \quad (9)$$

The influence function $\phi_{i,j} := \phi(x_i, x_j^v; \hat{\theta})$ linearly approximate $\delta_j(\epsilon)$ by

$$\delta_j(\epsilon) = \ell_j^v(\tilde{\theta}) - \ell_j^v(\hat{\theta}) \simeq \epsilon \times \phi_{i,j}, \quad (10)$$

where the closed-form expression of ϕ is given in (Koh and Liang, 2017) as

$$\phi_{i,j} := -\nabla_{\theta} \ell_j^v(\hat{\theta})^\top H_{\hat{\theta}}^{-1} \nabla_{\theta} \ell_i(\hat{\theta}) \quad (11)$$

and $H_{\hat{\theta}} := \frac{1}{n} \sum_{i=1}^n \nabla_{\theta}^2 \ell_i(\hat{\theta})$ is the Hessian matrix.

In short, by Eq. (10), $\delta_j(-1/n) > 0$ is equivalent to $\phi_{i,j} < 0$. We can compute x_i 's influence over the whole validation set by summation

$$\Phi_i = \sum_{j=1}^m \phi_{i,j} = - \sum_{j=1}^m \nabla_{\theta} \ell_j^{v\top}(\hat{\theta}) H_{\hat{\theta}}^{-1} \nabla_{\theta} \ell_i(\hat{\theta}). \quad (12)$$

Now, $\Phi_i < 0$ implies that x_i is *good* for the whole validation set. Also, if Φ_i is smaller, then x_i is more likely to be a favorable sample, and vice versa.

3.3 On Robustness of Sampling Functions

With the influence measure Φ , it seems that we can simply drop all *unfavorable* samples that have $\Phi > 0$. However, we argue that using 0 as the threshold usually results in failure to the out-of-sample test, due to its sensitivity to distribution shift. Instead, we take the measure of *probabilistic* sampling by designing a sampling function $\pi(\Phi) \in [0, 1]$. We give the reason of this choice based on the deviation of the induced validation loss by inaccurate estimate of influence. Let's denote the validation loss with inaccurate influence by $\ell^v(\hat{\theta}; \hat{\Phi})$, thus

$$\Delta^2(L) := \frac{1}{m} \sum_{j=1}^m (\ell_j^v(\hat{\theta}; \hat{\Phi}) - \ell_j^v(\hat{\theta}))^2 \quad (13)$$

indicates the robustness of the model under $\hat{\Phi}$. We then give the following proposition on $\Delta^2(L)$ with respect to sampling function π . Proof can be found in Appendix A.

Proposition 1 (Robustness of Probabilistic Sampling under Inaccurate Influence). *Let $\pi'(\Phi_i)$ be the derivative of $\pi(\cdot)$ function when taking Φ_i as its input, we have*

$$\sup_{\Phi, \hat{\Phi}} \Delta^2(L) = \gamma \sum_{i=1}^n (\pi(\hat{\Phi}_i) - \pi(\Phi_i))^2 \sum_{j=1}^m \phi_{i,j}^2 \quad (14)$$

$$\simeq \gamma \sum_{i=1}^n \left((\hat{\Phi}_i - \Phi_i) \pi'(\Phi_i) \right)^2 \sum_{j=1}^m \phi_{i,j}^2 \quad (15)$$

where γ is a constant.

It can be viewed that $\Delta^2(L)$ is controlled by the derivative of sampling function $\pi'(\Phi)$. For the sigmoid sampling in Eq. (3), it is easy to derive that

$$\pi'(\Phi) = -\alpha \pi(\Phi)(1 - \pi(\Phi)), \quad (16)$$

which means $\max |\pi'(\Phi)| = \frac{1}{4}\alpha$ when $\Phi = 0$. $\Delta^2(L)$ is hence controlled by the hyper-parameter α . When $|\Phi|$ increases, $|\pi'(\Phi)|$ reduces sharply, which ensures the variance's upper bound being tight all the time. By contrast, in deterministic sampling, $\Delta^2(L)$ is sensitive to inaccurate $\hat{\Phi}$ because it is "hard", or more rigorously, because $\Delta^2(L)$ is probably large due to large $|\pi(\Phi) - \pi(\hat{\Phi})|$ caused by an improper dropout threshold.

4 Efficient Implementation

Recap Eq. (12), computing Φ_i requires $\phi_{i,j}$ in Eq. (11) for $j = 1, \dots, m$ on all validation samples. As a result, the computation of all $\{\Phi_i\}_{i=1}^n$ has $\mathcal{O}(mn)$ time complexity. Moreover, for DNNs with massive parameters, computing the layer-wise gradients $\nabla_{\theta} \ell(\theta)$ is intractable. These limitations prevent the use of IF from DL RE models. To address it, we here propose a rather efficient implementation of REIF. We demonstrate how to reduce the complexity of calculating influences from $\mathcal{O}(mn)$ to $\mathcal{O}(n)$, then to $\mathcal{O}(1)$. In addition, we show how to compute the influence function by stochastic estimation.

4.1 Computing Influences in Linear Time

We argue that in Eq. (12), it is unnecessary to calculate $\phi_{i,j}$ separately, since here we only care about their summations. Specifically, since the summation is only related to the subscript j , we can cast it to

$$\Phi_i = -\nabla_{\theta} \ell_i^{\top}(\hat{\theta}) H_{\hat{\theta}}^{-1} \sum_{j=1}^m \nabla_{\theta} \ell_j^v(\hat{\theta}) \quad (17)$$

$$= -\nabla_{\theta} \ell_i^{\top}(\hat{\theta}) H_{\hat{\theta}}^{-1} \nabla_{\theta} \sum_{j=1}^m \ell_j^v(\hat{\theta}) \quad (18)$$

$$= -m \nabla_{\theta} \ell_i^{\top}(\hat{\theta}) H_{\hat{\theta}}^{-1} \nabla_{\theta} L(\hat{\theta}), \quad (19)$$

where $L(\hat{\theta})$ comes from Eq. (6). By this derivation, we can calculate $L(\hat{\theta})$ rather than all $l_j(\hat{\theta})$, then take derivative of $L(\hat{\theta})$. Since $L(\hat{\theta})$ only needs to be calculated once and it is shared in calculating all Φ_i s, this process only requires $\mathcal{O}(n)$ time, without loss of accuracy.

4.2 Linear Approximation for $\mathcal{O}(1)$ Complexity

$\nabla_{\theta} \ell(\hat{\theta})$ in Eq. (17) usually has complicated expression when $f_{\theta}(\cdot)$ is a neural network, hence the previous works implemented it by the auto-grad

systems like TensorFlow (Abadi et al., 2016) and PyTorch (Paszke et al., 2019). However, when the number of alternative training instances is large, even $\mathcal{O}(n)$ is not satisfactory enough, because additional differential operations need to be done on each $\ell_i(\hat{\theta})$ sequentially. Moreover, when faced with complex neural networks with massive parameters, computing the Hessian matrix $H_{\hat{\theta}}$ and its inversion is intractable. Considering these issues, we propose a linear approximation approach to reduce the complexity to $\mathcal{O}(1)$, and avoid operating on all parameters of the neural network.

Suppose the cross entropy loss function is used:

$$\ell(\theta) = - \sum_{k=1}^K \mathbb{I}\{y = k\} \log P(y = k|x; \theta) \quad (20)$$

where $\mathbb{I}(\cdot)$ is an indicator function. Let $\mathbf{y}, \hat{\mathbf{y}} \in \mathbb{R}^K$ be the one-hot label vector, e.g., $(1, 0, 0)^\top$, and prediction vector, e.g., $(0.8, 0.1, 0.1)^\top$, respectively. We replace $\nabla_{\theta} \ell(\theta)$ in Eq. (11) with the derivatives on β (the weight of the last fully-connected layer):

$$\nabla_{\theta} \ell(\theta) \Rightarrow \nabla_{\beta} \ell(\theta) = (\hat{\mathbf{y}} - \mathbf{y}) \mathbf{x}^\top \in \mathbb{R}^{d \times K} \quad (21)$$

where \mathbf{x} is the input of the last fully-connected layer. This closed-form expression allows computing batch gradients in $\mathcal{O}(1)$ time. Although the calculated influence might be inaccurate, it is still reliable for measuring instances' *relative quality* in general. We will validate this claim in our experiments.

4.3 Algorithm

Algorithm 1 shows the details of REIF, please refer to Appendix B. It has two hyper-parameters: the sampling ratio r and the sigmoid sampling parameter α . The optimal value of r depends on quality of the dataset, since the higher quality it is, the more favorable instances it might have. Keeping $\alpha = 1$ is satisfactory in most scenarios.

In particular, on the line #14 of Algorithm 1, we compute the product between the inverse Hessian matrix and a gradient vector via the stochastic estimation procedure by Koh & Liang (Koh and Liang, 2017). Denoting the vector $\nabla_{\theta} L(\hat{\theta})$ by v , it first initializes the approximate inverse Hessian-Vector-Product (HVP) by $\tilde{H}_0^{-1} v \leftarrow v$, then repeatedly samples n_b training instances and updates as

$$\tilde{H}_t^{-1} v \leftarrow v + \left(I - \frac{1}{n_b} \sum \nabla_{\theta}^2 \ell(\hat{\theta}) \right) \tilde{H}_{t-1}^{-1} v \quad (22)$$

until $\tilde{H}_t^{-1} v$ converges. In our algorithm, we only need to do this once after each epoch, to get the pre-computed inverse HVP $s = H_{\hat{\theta}}^{-1} \nabla_{\theta} L(\hat{\theta})$. Therefore, during training, we directly compute $\nabla_{\theta} \ell_i(\hat{\theta})$ for each instance according to Eq. (21), then multiply it with the precomputed s .

5 Experiments

We concentrate on the following research questions:

RQ1. How does our REIF perform as compared with classical baselines?

RQ2. How does the sampling ratio r influence the performance of the REIF?

RQ3. Does the sigmoid function lead to more robust sampling than the deterministic sampling?

RQ4. How does the proposed dynamic instance sampling perform compared with the post-hoc sampling using IF?

5.1 Datasets

In our experiments, we use two versions of widely used NYT datasets, the NYT-SMALL and NYT-LARGE. The small version is released in (Riedel et al., 2010), by aligning Freebase with the New York Times corpus. In particular, we use the filtered version of the NYT-SMALL released by (Zeng et al., 2015). The large version was released by (Lin et al., 2016). Data statistics can be found in Appendix C.

5.2 Experimental Setups

We pick PCNN (PCNN+ONE) (Zeng et al., 2015) as the backbone in our experiments, and include several baselines for comparison: the attention-based PCNN (PCNN+ATT) and the naive average method (PCNN+AVE) (Lin et al., 2016). Note that our REIF method is model-agnostic, hence it is applicable for other deep learning based backbones as well, e.g., CNN and RNN. Setups of models can be found in Appendix D.

We sample a clean validation set from training set by a rule-based approach used in (Jia et al., 2019), in order to obtain the inverse HVP required for calculating influences. The details of its establishment and discussions of this validation set can be found in Appendix E. During subsampling, we set $\alpha = 1$ and $r \in \{5\%, 10\%, 20\%, 30\%\}^2$ for our REIF.

²The ceiling function is used for rounding.

Table 1: P@N for relation extraction results, on NYT-SMALL and NYT-LARGE, where the best ones are in bold.

Dataset	NYT-SMALL				NYT-LARGE			
	100	200	300	Mean	100	200	300	Mean
PCNN + ONE	54.0	52.7	52.2	53.0	70.4	66.4	63.6	66.8
PCNN + AVE	52.7	50.8	47.3	50.3	73.0	71.2	67.8	70.6
PCNN + ATT	52.7	50.7	49.5	50.9	79.7	76.0	71.6	75.8
PCNN + REIF (Proposed)	75.2	65.1	60.8	67.0	86.4	82.5	80.3	83.1

Table 2: Precision (%) of various DS methods using PCNN as backbones / other DS methods for different recalls (0.1, 0.2, 0.3) on NYT-LARGE. The results of cited methods are drawn from their papers, and the best are in bold.

PCNN	0.1	0.2	0.3	Mean
+ONE	64.7	57.1	48.9	56.9
+ATT	74.3	63.3	56.5	64.7
+ONE+soft-label (Liu et al., 2017)	71.6	62.5	54.1	62.7
+ATT+soft-label (Liu et al., 2017)	75.1	67.5	55.8	66.1
+ONE+DSGAN (Qin et al., 2018a)	65.5	57.2	50.0	57.6
+ATT+DSGAN (Qin et al., 2018a)	70.5	62.2	53.3	62.0
+PE+REINF (Zeng et al., 2018)	70.1	66.2	56.1	64.1
+ONE+RL (Qin et al., 2018b)	66.7	56.1	48.3	64.1
+ATT+RL (Qin et al., 2018b)	68.3	60.0	52.2	60.2
+ONE+ADV (Wu et al., 2017)	71.7	58.9	51.1	60.6
+ONE+AN (Han et al., 2018)	80.3	70.2	60.3	70.3
+ATT-RA+BAG-ATT (Ye and Ling, 2020)	78.8	68.9	62.1	69.9
+SATT (Zhou et al., 2021)	78.2	69.1	59.5	68.9
DISTRE (Alt et al., 2019)	65.2	64.4	60.9	63.5
RedSandT (Christou and Tsoumakas, 2021)	73.1	67.3	58.0	66.1
Trans-SA (Xiao et al., 2022)	74.1	67.2	57.9	66.4
PCNN+REIF (Ours)	82.6	73.9	60.9	72.5

5.3 Effects of Influence Subsampling (RQ1)

Fig. 3 shows the precision-recall curve in held-out evaluation of ONE, AVE, ATT, and our REIF, and Table 1 illustrates the corresponding P@N of all methods. Our REIF performs the best among all methods. In details, on NYT-SMALL, our REIF improves 14% over ONE, and 16.1% over ATT; on NYT-LARGE, the improvements are 14.1% and 5.1%, respectively, in terms of the mean P@N. Specifically, REIF only leverages part of instances during training, while ATT involves all instances but performs badly on NYT-SMALL, and ONE only picks one instance per bag. It means that neither picking too many nor too few instances gains satisfactory performance in distant supervision. On contrast, our REIF can detect and pick those favorable ones from the noisy dataset, thus achieving a better model. In distant supervision, our method is effective for achieving nice trade-off between efficiency and effectiveness. Moreover, we compare our method with many DS baselines, including ad-

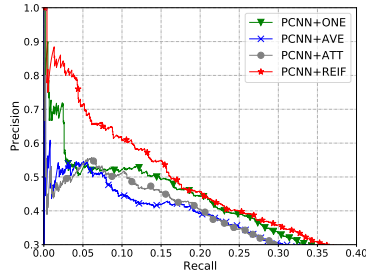
versarial training, reinforcement learning, attention, and GAN based methods, using the reported results. As shown in Table 2, REIF still is superior.

5.4 Effects of Sampling Ratio (RQ2)

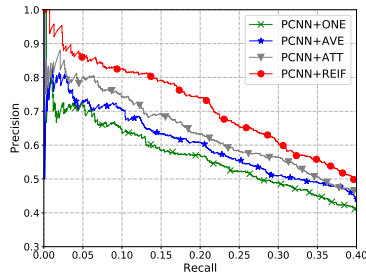
We evaluate the performance of REIF with respect to different r by repeat experiments. Results are reported in Fig. 5. REIF keeps stable when sampling ratio ranges from 5% to 30%, such that adding more instances does not make much difference, which might be due to high noise in the NYT dataset, i.e., focusing on those favorable instances is enough for training a satisfactory RE model.

5.5 Effects of Sigmoid Sampling & Dynamic Sampling (RQ3, RQ4)

Our REIF is engaged with the proposed probabilistic sigmoid sampling and DIS, namely REIF+P+DIS. We would like to validate these two techniques compared with the deterministic sampling (REIF+D+DIS), and the post-hoc sampling (REIF+P+PH). Our main observations from Fig. 4



(a) NYT-SMALL



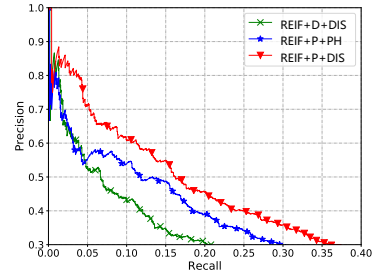
(b) NYT-LARGE

Figure 3: Aggregated precision-recall (P-R) curves obtained by PCNN+ONE, PCNN+AVE, PCNN+ATT, and the proposed PCNN+REIF on NYT-SMALL and NYT-LARGE datasets.

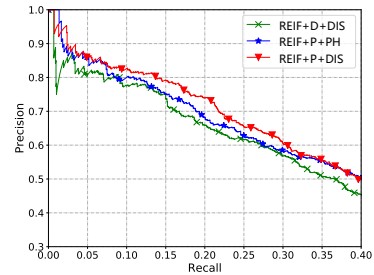
are as follows:

(1) The probabilistic sigmoid sampling is crucial for robust subsampling, as the REIF+D+DIS performs the worst in both datasets. As mentioned in Proposition 1, drawbacks of REIF+D mainly come from the inaccurate estimate of influence $\hat{\Phi}$, due to the non-convexity of neural networks and the use of linear approximations. That is, we could not determine the instances that have $\hat{\Phi}$ around the threshold with very high confidence, e.g., deterministic ranking and selecting, since this causes high variance of the resulting test loss, as indicated by Eq. (14). By contrast, we should assign them similar probabilities to be sampled, as done in REIF+P, to avoid sharp variation of the test loss caused by inaccurate influences in deterministic selection.

(2) Our dynamic sampling method generally performs better than post-hoc sampling in DS, especially on the tail instances. When recall is high, REIF+DIS performs better on the minor relations, thus has higher precision than REIF+PH. In DIS, more minor relation instances are maintained, which facilitates the model’s capacity of mining minor relation instances. Considering efficiency and the overall effectiveness, we shall prefer DIS in practice.



(a) NYT-SMALL



(b) NYT-LARGE

Figure 4: Precision-recall curve of compared REIF variants, where the REIF+P+DIS is the REIF with probabilistic sigmoid sampling and dynamic sampling, +D means deterministic sampling and +PH means post-hoc sampling.

6 Manual Evaluation & Case Study

Held-out evaluation usually suffers from false negative examples in Freebase (Zeng et al., 2015). To further check our method, we perform manual evaluation by choosing the entity pairs which are labeled as “NA” but predicted a relation (not “NA”) with high confidence. The top- k precisions are reported in Table 3, where the results of Mintz (Mintz et al., 2009), MultiR (Hoffmann et al., 2011), MIML (Surdeanu et al., 2012), PCNN+ONE (Zeng et al., 2015) and APCNN (Ji et al., 2017) are drawn from their papers. It could be seen our method outperforms baselines in extracting new facts from the false negative examples.

Fig. 6 reports an example of calculating influences that support instance selection. Picking a relation *children* as the example, influences and attention scores (Lin et al., 2016) are computed, from which we can identify that the influences quantitatively measure their individual quality. Recall in Section 3.2 that the smaller influences indicate better data quality. The first and the last instances are clearly right and wrong, respectively, in terms of indicating the relation *children* between their entities. By contrast, the second one tends to be right

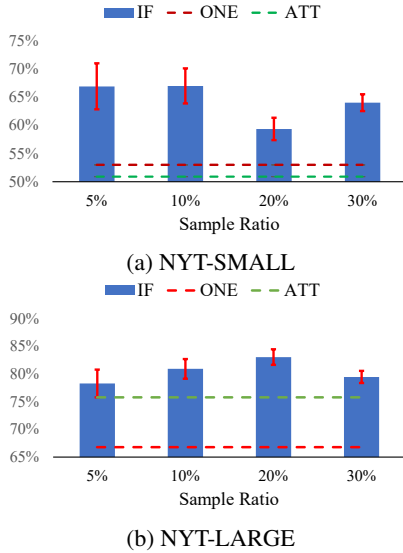


Figure 5: Mean P@N (average of P@100/200/300) varies with sampling ratio of REIF (IF) method. Red bar represents standard error by 5 times repeat experiments.

Instances	Influences	Att Scores
... because of art rooney , the legendary steelers ' owner ... and they have continued to be a family under his oldest son , dan rooney .	-2.23E-02	1.11E-04
mother of joseph paula and walter eva, grandmother of david, lauren, jacob , miriam and Leah .	-1.07E-04	2.61E-09
... the suspense novelists mary higgins clark and carol higgins clark signed books and posed for photographs for five hours ...	1.50E-05	1.44E-07
... daughter jamie baldinger and her husband, joseph ; son david goldring and his wife rachel ...	7.81E-04	1.39E-09

Figure 6: Examples of influences calculated with the relation *children*, on NYT-LARGE. The words in bold are entities. The *Att Scores* (Lin et al., 2016) are standardized into $[0, 1]$ by softmax, and *Influence* is the smaller the better.

because it implies that *Joseph* is the parent of *Jacob*. Although two entities in the third instance are very similar, no evidence shows they are relatives. Therefore, sampling probabilities can be obtained via these influences for the further subsampling process.

7 Conclusion & Discussion

In this work, we proposed an efficient subsampling scheme to find the influential instances for DS, namely REIF. Our method is model-agnostic, therefore it can be engaged in the majority of RE models. REIF can be generalized to other tasks which also confront noisy data. For instance, in other weak supervision scenarios such as active learning, our method can be an effective approach

Table 3: Precision values for the top 100, 200 and 500 via manual evaluation. Avg denotes the average of the former three columns. Best ones are in bold.

Accuracy (%)	Top 100	Top 200	Top 500	Avg
Mintz	77	71	55	67.7
MultiR	83	74	49	68.7
MIML	85	75	61	73.7
PCNN+ONE	86	80	69	78.3
APCNN	87	82	72	80.3
PCNN+ATT	86	81	70	79.0
PCNN+REIF	88	84	76	82.7

to build data pipeline from data quality measure to data selection. We leave this as our future work.

References

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. 2016. TensorFlow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*.
- Christoph Alt, Marc Hübner, and Leonhard Hennig. 2019. Fine-tuning pre-trained transformer language models to distantly supervised relation extraction. In *Annual Meeting of the Association for Computational Linguistics*, pages 1388–1398.
- Graciela Boente, Ana M Pires, and Isabel M Rodrigues. 2002. Influence functions and outlier detection under the common principal components model: A robust approach. *Biometrika*, 89(4):861–875.
- Jing Chen, Zhiqiang Guo, and Jie Yang. 2021. Distant supervision for relation extraction via noise filtering. In *International Conference on Machine Learning and Computing*, pages 361–367.
- Despina Christou and Grigorios Tsoumakas. 2021. Improving distantly-supervised relation extraction through bert-based label and instance embeddings. *IEEE Access*, 9:62574–62582.
- Jun Feng, Minlie Huang, Li Zhao, Yang Yang, and Xiaoyan Zhu. 2018. Reinforcement learning for relation classification from noisy data. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Xu Han, Zhiyuan Liu, and Maosong Sun. 2018. Denoising distant supervision for relation extraction via instance-level adversarial training. *arXiv preprint arXiv:1805.10959*.
- Raphael Hoffmann, Congle Zhang, Xiao Ling, Luke Zettlemoyer, and Daniel S. Weld. 2011. Knowledge-based weak supervision for information extraction of overlapping relations. In *Annual Meeting of the Association for Computational Linguistics*, pages 541–550.

- Peter J. Huber. 2004. *Robust Statistics*, volume 523. John Wiley & Sons.
- Guoliang Ji, Kang Liu, Shizhu He, Jun Zhao, et al. 2017. Distant supervision for relation extraction with sentence-level attention and entity descriptions. In *AAAI Conference on Artificial Intelligence*, volume 3060.
- Wei Jia, Dai Dai, Xinyan Xiao, and Hua Wu. 2019. ARNOR: Attention regularization based noise reduction for distant supervision relation classification. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1399–1408.
- Pang Wei Koh and Percy Liang. 2017. Understanding black-box predictions via influence functions. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1885–1894.
- 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*, pages 2124–2133.
- Tianyu Liu, Kexiang Wang, Baobao Chang, and Zhifang Sui. 2017. A soft-label method for noise-tolerant distantly supervised relation extraction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1790–1795.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics*, pages 1003–1011. Association for Computational Linguistics.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035.
- Pengda Qin, Weiran Xu, and William Yang Wang. 2018a. DSGAN: Generative adversarial training for distant supervision relation extraction. *arXiv preprint arXiv:1805.09929*.
- Pengda Qin, Weiran Xu, and William Yang Wang. 2018b. Robust distant supervision relation extraction via deep reinforcement learning. *arXiv preprint arXiv:1805.09927*.
- 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.
- Ge Shi, Chong Feng, Lifu Huang, Boliang Zhang, Heng Ji, Lejian Liao, and He-Yan Huang. 2018. Genre separation network with adversarial training for cross-genre relation extraction. In *Conference on Empirical Methods in Natural Language Processing*, pages 1018–1023.
- Mihai Surdeanu, Julie Tibshirani, Ramesh Nallapati, and Christopher D. Manning. 2012. Multi-instance multi-label learning for relation extraction. In *Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 455–465.
- Shingo Takamatsu, Issei Sato, and Hiroshi Nakagawa. 2012. Reducing wrong labels in distant supervision for relation extraction. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics*, pages 721–729.
- Zifeng Wang, Hong Zhu, Zhenhua Dong, Xiuqiang He, and Shao-Lun Huang. 2020. Less is better: Unweighted data subsampling via influence function. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Yi Wu, David Bamman, and Stuart Russell. 2017. Adversarial training for relation extraction. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1778–1783.
- Yan Xiao, Yaochu Jin, Ran Cheng, and Kuangrong Hao. 2022. Hybrid attention-based transformer block model for distant supervision relation extraction. *Neurocomputing*, 470:29–39.
- Yaosheng Yang, Wenliang Chen, Zhenghua Li, Zhengqiu He, and Min Zhang. 2018. Distantly supervised NER with partial annotation learning and reinforcement learning. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2159–2169.
- Zhi-Xiu Ye and Zhen-Hua Ling. 2020. Distant supervision relation extraction with intra-bag and inter-bag attentions. In *Conference of the North American Chapter of the Association for Computational Linguistics*, pages 2810–2819.
- Yujin Yuan, Liyuan Liu, Siliang Tang, Zhongfei Zhang, Yueting Zhuang, Shiliang Pu, Fei Wu, and Xiang Ren. 2019. Cross-relation cross-bag attention for distantly-supervised relation extraction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 419–426.
- Matthew D Zeiler. 2012. ADADELTA: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.
- Daojian Zeng, Kang Liu, Yubo Chen, and Jun Zhao. 2015. Distant supervision for relation extraction via piecewise convolutional neural networks. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1753–1762.

Xiangrong Zeng, Shizhu He, Kang Liu, and Jun Zhao. 2018. Large scaled relation extraction with reinforcement learning. In *AAAI Conference on Artificial Intelligence*, pages 5658–5665.

Peng Zhou, Jiaming Xu, Zhenyu Qi, Hongyun Bao, Zhineng Chen, and Bo Xu. 2018. Distant supervision for relation extraction with hierarchical selective attention. *Neural Networks*, 108:240 – 247.

Yanru Zhou, Limin Pan, Chongyou Bai, Senlin Luo, and Zhouting Wu. 2021. Self-selective attention using correlation between instances for distant supervision relation extraction. *Neural Networks*, 142:213–220.

A Proof of Proposition 1

Proposition 1 (Robustness of Probabilistic Sampling under Inaccurate Influence). *Let $\pi'(\Phi_i)$ be the derivative of $\pi(\cdot)$ function when taking Φ_i as its input, we have*

$$\begin{aligned} \sup_{\Phi, \hat{\Phi}} \Delta^2(L) &= \gamma \sum_{i=1}^n (\pi(\hat{\Phi}_i) - \pi(\Phi_i))^2 \sum_{j=1}^m \phi_{i,j}^2 \\ &\simeq \gamma \sum_{i=1}^n \left((\hat{\Phi}_i - \Phi_i) \pi'(\Phi_i) \right)^2 \sum_{j=1}^m \phi_{i,j}^2 \end{aligned} \quad (\text{A.1})$$

where γ is a constant.

Proof.

$$\Delta^2(L) \propto \sum_{j=1}^m (\ell_j^v(\tilde{\theta}; \hat{\Phi}) - \ell_j^v(\tilde{\theta}))^2 \quad (\text{A.2})$$

$$= \sum_{j=1}^m (\ell_j^v(\tilde{\theta}; \hat{\Phi}) - \ell_j^v(\hat{\theta}) + \ell_j(\hat{\theta}) - \ell_j^v(\tilde{\theta}))^2 \quad (\text{A.3})$$

$$\propto \sum_{j=1}^m \left(\sum_{i=1}^n \pi(\hat{\Phi}_i) \phi_{i,j} - \pi(\Phi_i) \phi_{i,j} \right)^2 \quad (\text{A.4})$$

$$\leq \sum_{i=1}^n (\pi(\hat{\Phi}_i) - \pi(\Phi_i))^2 \sum_{j=1}^m \phi_{i,j}^2 \quad (\text{A.5})$$

Eq. (A.4) is obtained by definition of probabilistic subsampling because

$$\begin{aligned} \ell_j^v(\tilde{\theta}) - \ell_j^v(\hat{\theta}) &\simeq \sum_{i=1}^n \epsilon_i \phi_{i,j} \\ &\propto \sum_{i=1}^n \pi(\Phi_i) \phi_{i,j}. \end{aligned} \quad (\text{A.6})$$

Details can be referred to (Wang et al., 2020). Taking linear Taylor expansion of the $\pi(\hat{\Phi}_i) - \pi(\Phi_i)$ at the last line yields the final result. \square

B Algorithm

C Dataset Statistics

D General Setups for Training PCNN

Following the configurations of previous works, we employ word2vec³ to extract the word embeddings, to process the raw data. Parameters of PCNN are set according to (Zeng et al., 2015): window size $d^w = 3$, sentence embedding size $d^s = 230$, word

³<https://code.google.com/p/word2vec/>

Algorithm 1 Finding Influential Instances for DS on RE by Influence Subsampling.

Input: Training and validation data $\mathcal{D}_{tr}, \mathcal{D}_{va}$;

Hyper-parameters: r and α ;

- 1: **for** epoch $t = 1 \rightarrow T$ **do**
 - 2: **repeat**
 - 3: Initialize the selected instances set $X_{sub} = \emptyset$;
 - 4: Sequentially sample a batch of bags $\{X_1, \dots, X_B\}$ from \mathcal{D}_{tr} ;
 - 5: **for** bag $b = 1 \rightarrow B$ **do**
 - 6: Obtain instance-level loss as $\vec{\ell} \leftarrow (\ell_1(\hat{\theta}_t), \dots, \ell_{|X_b|}(\hat{\theta}_t))^\top$;
 - 7: Compute influences $\Phi_i \leftarrow s_t^\top \nabla_{\theta} \ell_i(\hat{\theta}_t) \forall i = 1, \dots, |X_b|$;
 - 8: Compute sampling probability $\pi_i \leftarrow 1/(1 + \exp(\alpha \times \Phi_i)) \forall i$;
 - 9: Sample $r \times |X_b|$ instances from X_b to get \tilde{X}_b , and $X_{sub} \leftarrow X_{sub} \cup \tilde{X}_b$;
 - 10: **end for**
 - 11: Update $\hat{\theta}_t$ using the selected subset X_{sub} by gradient descent;
 - 12: **until** going through all bags in \mathcal{D}_{tr} .
 - 13: Get validation loss by $L(\hat{\theta}_t) \leftarrow \frac{1}{m} \sum_{j=1}^m \ell_j^v(\hat{\theta}_t)$ on \mathcal{D}_{va} ;
 - 14: Obtain $s_t \leftarrow H_t^{-1} \nabla_{\theta} L(\hat{\theta}_t)$ by stochastic estimation as done in Eq. (22);
 - 15: **end for**
-

dimension $d^a = 50$ and position dimension $d^p = 5$ for fair comparison. During training, we fix the batch size $B = 128$, dropout ratio $p = 0.5$, and use the ADADELTA (Zeiler, 2012) with parameters $\rho = 0.95$ and $\varepsilon = 10^{-6}$ for optimization. Since we find the default hyperparameters already lead superior performance of REIF, we did not make further tuning.

E Establishing the Validation Set

Due to lacking clean validation set, we utilize automatic selection similar to ARNOR (Jia et al., 2019). It takes top 10% high-frequency patterns of each relation as initial pattern, then takes max 5 new patterns in one loop for each relation in bootstrap procedure. We stop bootstrap until 10% training samples are involved. Our experiments demonstrate REIF can gain significantly from this automatically built validation set, although it is collected by heuristics and not absolutely clean.

Table 1: Data statistics of used two NYT datasets. “# Pos”, “# Ins”, “# Rel”: number of positive bags, instances and relations, respectively.

	NYT-SMALL		NYT-LARGE	
	Train	Test	Train	Test
# Bags	65,726	93,574	281,270	96,678
# Pos	4,266	1,732	18,252	1,950
# Ins	112,941	152,416	522,611	172,448
# Rel	26	26	53	53