

Heterogeneous Supervision for Relation Extraction: A Representation Learning Approach

Liyuan Liu^{†*} Xiang Ren^{†*} Qi Zhu[†] Huan Gui[†] Shi Zhi[†] Heng Ji[‡] Jiawei Han[†]

[†] University of Illinois at Urbana-Champaign, Urbana, IL, USA

[‡] Computer Science Department, Rensselaer Polytechnic Institute, USA

[†]{ll2, xren7, qiz3, huangui2, shizhi2, hanj}@illinois.edu [‡]{jih}@rpi.edu

Abstract

Relation extraction is a fundamental task in information extraction. Most existing methods have heavy reliance on annotations labeled by human experts, which are costly and time-consuming. To overcome this drawback, we propose a novel framework, REHESSION, to conduct relation extractor learning using annotations from heterogeneous information source, e.g., knowledge base and domain heuristics. These annotations, referred as *heterogeneous supervision*, often conflict with each other, which brings a new challenge to the original relation extraction task: how to infer the true label from noisy labels for a given instance. Identifying context information as the backbone of both relation extraction and true label discovery, we adopt embedding techniques to learn the distributed representations of context, which bridges all components with mutual enhancement in an iterative fashion. Extensive experimental results demonstrate the superiority of REHESSION over the state-of-the-art.

1 Introduction

One of the most important tasks towards text understanding is to detect and categorize semantic relations between two entities in a given context. For example, in Fig. 1, with regard to the sentence of c_1 , relation between *Jesse James* and *Missouri* should be categorized as *died_in*. With accurate identification, relation extraction systems can provide essential support for many applications. One

example is question answering, regarding a specific question, relation among entities can provide valuable information, which helps to seek better answers (Bao et al., 2014). Similarly, for medical science literature, relations like protein-protein interactions (Fundel et al., 2007) and gene disease associations (Chun et al., 2006) can be extracted and used in knowledge base population. Additionally, relation extractors can be used in ontology construction (Schutz and Buitelaar, 2005).

Typically, existing methods follow the supervised learning paradigm, and require extensive annotations from domain experts, which are costly and time-consuming. To alleviate such drawback, attempts have been made to build relation extractors with a small set of seed instances or human-crafted patterns (Nakashole et al., 2011; Carlson et al., 2010), based on which more patterns and instances will be iteratively generated by bootstrap learning. However, these methods often suffer from semantic drift (Mintz et al., 2009). Besides, knowledge bases like Freebase have been leveraged to automatically generate training data and provide *distant supervision* (Mintz et al., 2009). Nevertheless, for many domain-specific applications, distant supervision is either non-existent or insufficient (usually less than 25% of relation mentions are covered (Ren et al., 2015; Ling and Weld, 2012)).

Only recently have preliminary studies been developed to unite different supervisions, including knowledge bases and domain specific patterns, which are referred as *heterogeneous supervision*. As shown in Fig. 1, these supervisions often conflict with each other (Ratner et al., 2016). To address these conflicts, data programming (Ratner et al., 2016) employs a generative model, which encodes supervisions as *labeling functions*, and adopts the source consistency assumption: *a source is likely to provide true information with*

*Equal contribution.

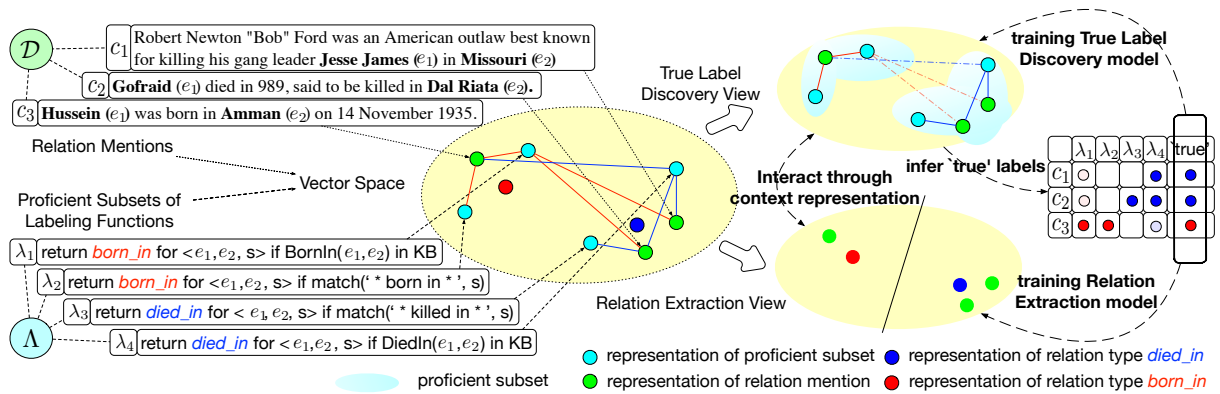


Figure 1: REHESSION Framework except Extraction and Representation of Text Features

the same probability for all instances. This assumption is widely used in true label discovery literature (Li et al., 2016) to model reliabilities of information sources like crowdsourcing and infer the true label from noisy labels. Accordingly, most true label discovery methods would trust a human annotator on all instances to the same level.

However, labeling functions, unlike human annotators, do not make casual mistakes but follow certain “error routine”. Thus, the reliability of a labeling function is not consistent among different pieces of instances. In particular, a labeling function could be more reliable for a certain subset (Varma et al., 2016) (also known as its *proficient subset*) comparing to the rest. We identify these proficient subsets based on context information, only trust labeling functions on these subsets and avoid assuming global source consistency.

Meanwhile, embedding methods have demonstrated great potential in capturing semantic meanings, which also reduce the dimension of overwhelming text features. Here, we present REHESSION, a novel framework capturing context’s semantic meaning through representation learning, and conduct both relation extraction and true label discovery in a context-aware manner. Specifically, as depicted in Fig. 1, we embed relation mentions in a low-dimension vector space, where similar relation mentions tend to have similar relation types and annotations. ‘True’ labels are further inferred based on reliabilities of labeling functions, which are calculated with their proficient subsets’ representations. Then, these inferred true labels would serve as supervision for all components, including context representation, true label discovery and relation extraction. Besides, the context representation bridges relation extraction with true label dis-

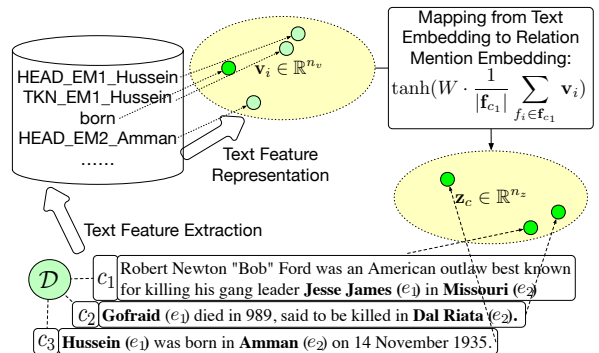


Figure 2: Relation Mention Representation

covery, and allows them to enhance each other.

To the best of our knowledge, the framework proposed here is the first method that utilizes representation learning to provide heterogeneous supervision for relation extraction. The high-quality context representations serve as the backbone of true label discovery and relation extraction. Extensive experiments on benchmark datasets demonstrate significant improvements over the state-of-the-art.

The remaining of this paper is organized as follows. Section 2 gives the definition of relation extraction with heterogeneous supervision. We then present the REHESSION model and the learning algorithm in Section 3, and report our experimental evaluation in Section 4. Finally, we briefly survey related work in Section 5 and conclude this study in Section 6.

2 Preliminaries

In this section, we would formally define relation extraction and heterogeneous supervision, including the format of labeling functions.

2.1 Relation Extraction

Here we conduct relation extraction in *sentence-level* (Bao et al., 2014). For a sentence d , an entity mention is a token span in d which represents an entity, and a relation mention is a triple (e_1, e_2, d) which consists of an ordered entity pair (e_1, e_2) and d . And the relation extraction task is to categorize relation mentions into a given set of relation types \mathcal{R} , or Not-Target-Type (None) which means the type of the relation mention does not belong to \mathcal{R} .

2.2 Heterogeneous Supervision

Similar to (Ratner et al., 2016), we employ labeling functions as basic units to encode supervision information and generate annotations. Since different supervision information may have different proficient subsets, we require each labeling function to encode only one elementary supervision information. Specifically, in the relation extraction scenario, we require each labeling function to only annotate one relation type based on one elementary piece of information, e.g., four examples are listed in Fig. 1.

Notice that knowledge-based labeling functions are also considered to be noisy because relation extraction is conducted in sentence-level, e.g. although `president_of` (*Obama, USA*) exists in KB, it should not be assigned with “*Obama was born in Honolulu, Hawaii, USA*”, since `president_of` is irrelevant to the context.

2.3 Problem Definition

For a POS-tagged corpus \mathcal{D} with detected entities, we refer its relation mentions as $\mathcal{C} = \{c_i = (e_{i,1}, e_{i,2}, d), \forall d \in \mathcal{D}\}$. Our goal is to annotate entity mentions with relation types of interest ($\mathcal{R} = \{r_1, \dots, r_K\}$) or None. We require users to provide heterogeneous supervision in the form of labeling function $\Lambda = \{\lambda_1, \dots, \lambda_M\}$, and mark the annotations generated by Λ as $\mathcal{O} = \{o_{c,i} | \lambda_i \text{ generate annotation } o_{c,i} \text{ for } c \in \mathcal{C}\}$. We record relation mentions annotated by Λ as \mathcal{C}_l , and refer relation mentions without annotation as \mathcal{C}_u . Then, our task is to train a relation extractor based on \mathcal{C}_l and categorize relation mentions in \mathcal{C}_u .

3 The REHESSION Framework

Here, we present REHESSION, a novel framework to infer true labels from automatically generated noisy labels, and categorize unlabeled instances

\mathbf{f}_c	c 's text features set, where $c \in \mathcal{C}$
\mathbf{v}_i	text feature embedding for $f_i \in \mathcal{F}$
\mathbf{z}_c	relation mention embedding for $c \in \mathcal{C}$
\mathbf{l}_i	embedding for λ_i 's proficient subset, $\lambda_i \in \Lambda$
$o_{c,i}$	annotation for c , generated by labeling function λ_i
o_c^*	underlying true label for c
$\rho_{c,i}$	identify whether $o_{c,i}$ is correct
\mathcal{S}_i	the proficient subset of labeling function λ_i
$s_{c,i}$	identify whether c belongs to λ_i 's proficient subset
\mathbf{t}_i	relation type embedding for $r_i \in \mathcal{R}$

Table 1: Notation Table.

into a set of relation types. Intuitively, errors of annotations (\mathcal{O}) come from mismatch of contexts, e.g., in Fig. 1, λ_1 annotates c_1 and c_2 with ‘true’ labels but for mismatched contexts ‘killing’ and ‘killed’. Accordingly, we should only trust labeling functions on *matched* context, e.g., trust λ_1 on c_3 due to its context ‘was born in’, but not on c_1 and c_2 . On the other hand, relation extraction can be viewed as *matching* appropriate relation type to a certain context. These two *matching* processes are closely related and can enhance each other, while context representation plays an important role in both of them.

Framework Overview. We propose a general framework to learn the relation extractor from automatically generated noisy labels. As plotted in Fig. 1, distributed representation of context bridges relation extraction with true label discovery, and allows them to enhance each other. Specifically, it follows the steps below:

1. After being extracted from context, text features are embedded in a low dimension space by representation learning (see Fig. 2);
2. Text feature embeddings are utilized to calculate relation mention embeddings (see Fig. 2);
3. With relation mention embeddings, true labels are inferred by calculating labeling functions’ reliabilities in a context-aware manner (see Fig. 1);
4. Inferred true labels would ‘supervise’ all components to learn model parameters (see Fig. 1).

We now proceed by introducing these components of the model in further details.

3.1 Modeling Relation Mention

As shown in Table 2, we extract abundant lexical features (Ren et al., 2016; Mintz et al., 2009) to characterize relation mentions. However, this abundance also results in the gigantic dimension of original text features ($\sim 10^7$ in our case). In

Feature	Description	Example
Entity mention (EM) head	Syntactic head token of each entity mention	"HEAD_EM1_Hussein", ...
Entity Mention Token	Tokens in each entity mention	"TKN_EM1_Hussein", ...
Tokens between two EMs	Tokens between two EMs	"was", "born", "in"
Part-of-speech (POS) tag	POS tags of tokens between two EMs	"VBD", "VBN", "IN"
Collocations	Bigrams in left/right 3-word window of each EM	"Hussein was", "in Amman"
Entity mention order	Whether EM 1 is before EM 2	"EM1_BEFORE_EM2"
Entity mention distance	Number of tokens between the two EMs	"EM_DISTANCE_3"
Body entity mentions numbers	Number of EMs between the two EMs	"EM_NUMBER_0"
Entity mention context	Unigrams before and after each EM	"EM_AFTER_was", ...
Brown cluster (learned on \mathcal{D})	Brown cluster ID for each token	"BROWN_010011001", ...

Table 2: Text features \mathcal{F} used in this paper. ("Hussein", "Amman", "Hussein was born in Amman") is used as an example.

order to achieve better generalization ability, we represent relation mentions with low dimensional ($\sim 10^2$) vectors. In Fig. 2, for example, relation mention c_3 is first represented as bag-of-features. After learning text feature embeddings, we use the average of feature embedding vectors to derive the embedding vector for c_3 .

Text Feature Representation. Similar to other principles of embedding learning, we assume text features occurring in the same contexts tend to have similar meanings (also known as distributional hypothesis (Harris, 1954)). Furthermore, we let each text feature’s embedding vector to predict other text features occurred in the same relation mentions or context. Thus, text features with similar meaning should have similar embedding vectors. Formally, we mark text features as $\mathcal{F} = \{f_1, \dots, f_{|\mathcal{F}|}\}$, record the feature set for $\forall c \in \mathcal{C}$ as \mathbf{f}_c , and represent the embedding vector for f_i as $\mathbf{v}_i \in \mathbb{R}^{n_v}$, and we aim to maximize the following log likelihood: $\sum_{c \in \mathcal{C}_l} \sum_{f_i, f_j \in \mathbf{f}_c} \log p(f_i | f_j)$, where $p(f_i | f_j) = \exp(\mathbf{v}_i^T \mathbf{v}_j^*) / \sum_{f_k \in \mathcal{F}} \exp(\mathbf{v}_i^T \mathbf{v}_k^*)$.

However, the optimization of this likelihood is impractical because the calculation of $\nabla p(f_i | f_j)$ requires summation over all text features, whose size exceeds 10^7 in our case. In order to perform efficient optimization, we adopt the negative sampling technique (Mikolov et al., 2013) to avoid this summation. Accordingly, we replace the log likelihood with Eq. 1 as below:

$$\mathcal{J}_E = \sum_{\substack{c \in \mathcal{C}_l \\ f_i, f_j \in \mathbf{f}_c}} (\log \sigma(\mathbf{v}_i^T \mathbf{v}_j^*) - \sum_{k=1}^V \mathbb{E}_{f_{k'} \sim \hat{P}} [\log \sigma(-\mathbf{v}_i^T \mathbf{v}_{k'}^*)]) \quad (1)$$

where \hat{P} is noise distribution used in (Mikolov et al., 2013), σ is the sigmoid function and V is number of negative samples.

Relation Mention Representation. With text feature embeddings learned by Eq. 1, a naive method to




Figure 3: Graphical model of $o_{c,i}$ ’s correctness

represent relation mentions is to concatenate or average its text feature embeddings. However, text features embedding may be in a different semantic space with relation types. Thus, we directly learn a mapping g from text feature representations to relation mention representations (Van Gysel et al., 2016a,b) instead of simple heuristic rules like concatenate or average (see Fig. 2):

$$\mathbf{z}_c = g(\mathbf{f}_c) = \tanh(W \cdot \frac{1}{|\mathbf{f}_c|} \sum_{f_i \in \mathbf{f}_c} \mathbf{v}_i) \quad (2)$$

where \mathbf{z}_c is the representation of $c \in \mathcal{C}_l$, W is a $n_z \times n_v$ matrix, n_z is the dimension of relation mention embeddings and \tanh is the element-wise hyperbolic tangent function.

In other words, we represent bag of text features with their average embedding, then apply linear map and hyperbolic tangent to transform the embedding from text feature semantic space to relation mention semantic space. The non-linear \tanh function allows non-linear class boundaries in other components, and also regularize relation mention representation to range $[-1, 1]$ which avoids numerical instability issues.

3.2 True Label Discovery

Because heterogeneous supervision generates labels in a discriminative way, we suppose its errors follow certain underlying principles, i.e., if a

Datasets	NYT	Wiki-KBP
% of None in Training	0.6717	0.5552
% of None in Test	0.8972	0.8532

Table 3: Proportion of None in Training/Test Set

labeling function annotates a instance correctly / wrongly, it would annotate other similar instances correctly / wrongly. For example, λ_1 in Fig. 1 generates wrong annotations for two similar instances c_1, c_2 and would make the same errors on other similar instances. Since context representation captures the semantic meaning of relation mention and would be used to identify relation types, we also use it to identify the mismatch of context and labeling functions. Thus, we suppose for each labeling function λ_i , there exists an proficient subset \mathcal{S}_i on \mathcal{R}^{nz} , containing instances that λ_i can precisely annotate. In Fig. 1, for instance, c_3 is in the proficient subset of λ_1 , while c_1 and c_2 are not. Moreover, the generation of annotations are not really random, and we propose a probabilistic model to describe the level of mismatch from labeling functions to real relation types instead of annotations' generation.

As shown in Fig. 3, we assume the indicator of whether c belongs to \mathcal{S}_i , $s_{c,i} = \delta(c \in \mathcal{S}_i)$, would first be generated based on context representation

$$p(s_{c,i} = 1 | \mathbf{z}_c, \mathbf{l}_i) = p(c \in \mathcal{S}_i) = \sigma(\mathbf{z}_c^T \mathbf{l}_i) \quad (3)$$

Then the correctness of annotation $o_{c,i}, \rho_{c,i} = \delta(o_{c,i} = o_c^*)$, would be generated. Furthermore, we assume $p(\rho_{c,i} = 1 | s_{c,i} = 1) = \phi_1$ and $p(\rho_{c,i} = 1 | s_{c,i} = 0) = \phi_0$ to be constant for all relation mentions and labeling functions.

Because $s_{c,i}$ would not be used in other components of our framework, we integrate out $s_{c,i}$ and write the log likelihood as

$$\begin{aligned} \mathcal{J}_T = & \sum_{o_{c,i} \in \mathcal{O}} \log(\sigma(\mathbf{z}_c^T \mathbf{l}_i) \phi_1^{\delta(o_{c,i}=o_c^*)} (1 - \phi_1)^{\delta(o_{c,i} \neq o_c^*)} \\ & + (1 - \sigma(\mathbf{z}_c^T \mathbf{l}_i)) \phi_0^{\delta(o_{c,i}=o_c^*)} (1 - \phi_0)^{\delta(o_{c,i} \neq o_c^*)}) \end{aligned} \quad (4)$$

Note that o_c^* is a hidden variable but not a model parameter, and \mathcal{J}_T is the likelihood of $\rho_{c,i} = \delta(o_{c,i} = o_c^*)$. Thus, we would first infer $o_c^* = \operatorname{argmax}_{o_c^*} \mathcal{J}_T$, then train the true label discovery model by maximizing \mathcal{J}_T .

3.3 Modeling Relation Type

We now discuss the model for identifying relation types based on context representation. For each relation mention c , its representation \mathbf{z}_c implies its relation type, and the distribution of relation type can be described by the soft-max function:

$$p(r_i | \mathbf{z}_c) = \frac{\exp(\mathbf{z}_c^T \mathbf{t}_i)}{\sum_{r_j \in \mathcal{R} \cup \{\text{None}\}} \exp(\mathbf{z}_c^T \mathbf{t}_j)} \quad (5)$$

where $\mathbf{t}_i \in \mathbb{R}^{v_z}$ is the representation for relation type r_i . Moreover, with the inferred true label o_c^* , the relation extraction model can be trained as a multi-class classifier. Specifically, we use Eq. 5 to approach the distribution

$$p(r_i | o_c^*) = \begin{cases} 1 & r_i = o_c^* \\ 0 & r_i \neq o_c^* \end{cases} \quad (6)$$

Moreover, we use KL-divergence to measure the dissimilarity between two distributions, and formulate model learning as maximizing \mathcal{J}_R :

$$\mathcal{J}_R = - \sum_{c \in \mathcal{C}_l} KL(p(\cdot | \mathbf{z}_c) || p(\cdot | o_c^*)) \quad (7)$$

where $KL(p(\cdot | \mathbf{z}_c) || p(\cdot | o_c^*))$ is the KL-divergence from $p(r_i | o_c^*)$ to $p(r_i | \mathbf{z}_c)$, $p(r_i | \mathbf{z}_c)$ and $p(r_i | o_c^*)$ has the form of Eq. 5 and Eq. 6.

3.4 Model Learning

Based on Eq. 1, Eq. 4 and Eq. 7, we form the joint optimization problem for model parameters as

$$\begin{aligned} \min_{W, \mathbf{v}, \mathbf{v}^*, \mathbf{l}, \mathbf{t}, o_c^*} \mathcal{J} = & -\mathcal{J}_R - \lambda_1 \mathcal{J}_E - \lambda_2 \mathcal{J}_T \\ \text{s.t. } \forall c \in \mathcal{C}_l, o_c^* = & \operatorname{argmax}_{o_c^*} \mathcal{J}_T, \mathbf{z}_c = g(\mathbf{f}_c) \end{aligned} \quad (8)$$

Collectively optimizing Eq. 8 allows heterogeneous supervision guiding all three components, while these components would refine the context representation, and enhance each other.

In order to solve the joint optimization problem in Eq. 8 efficiently, we adopt the stochastic gradient descent algorithm to update $\{W, \mathbf{v}, \mathbf{v}^*, \mathbf{l}, \mathbf{t}\}$ iteratively, and o_c^* is estimated by maximizing \mathcal{J}_T after calculating \mathbf{z}_c . Additionally, we apply the widely used dropout techniques (Srivastava et al., 2014) to prevent overfitting and improve generalization performance.

The learning process of REHESSION is summarized as below. In each iteration, we would sample a relation mention c from \mathcal{C}_l , then sample c 's text

features and conduct the text features’ representation learning. After calculating the representation of c , we would infer its true label o_c^* based on our true label discovery model, and finally update model parameters based on o_c^* .

3.5 Relation Type Inference

We now discuss the strategy of performing type inference for C_u . As shown in Table 3, the proportion of `None` in C_u is usually much larger than in C_l . Additionally, not like other relation types in \mathcal{R} , `None` does not have a coherent semantic meaning. Similar to (Ren et al., 2016), we introduce a heuristic rule: identifying a relation mention as `None` when (1) our relation extractor predict it as `None`, or (2) the entropy of $p(\cdot|z_c)$ over \mathcal{R} exceeds a pre-defined threshold η . The entropy is calculated as $H(p(\cdot|z_c)) = -\sum_{r_i \in \mathcal{R}} p(r_i|z_c) \log(p(r_i|z_c))$. And the second situation means based on relation extractor this relation mention is not likely belonging to any relation types in \mathcal{R} .

4 Experiments

In this section, we empirically validate our method by comparing to the state-of-the-art relation extraction methods on news and Wikipedia articles.

4.1 Datasets and settings

In the experiments, we conduct investigations on two benchmark datasets from different domains:¹ `NYT` (Riedel et al., 2010) is a news corpus sampled from $\sim 294k$ 1989-2007 New York Times news articles. It consists of 1.18M sentences, while 395 of them are annotated by authors of (Hoffmann et al., 2011) and used as test data;

`Wiki-KBP` utilizes 1.5M sentences sampled from 780k Wikipedia articles (Ling and Weld, 2012) as training corpus, while test set consists of the 2k sentences manually annotated in 2013 KBP slot filling assessment results (Ellis et al., 2012).

For both datasets, the training and test sets partitions are maintained in our experiments. Furthermore, we create validation sets by randomly sampling 10% mentions from each test set and used the remaining part as evaluation sets.

Feature Generation. As summarized in Table 2, we use a 6-word window to extract context features for each entity mention, apply the Stanford

¹ Codes and datasets used in this paper can be downloaded at: <https://github.com/LiyuanLucasLiu/ReHession>.

Kind	Wiki-KBP		NYT	
	#Types	#LF	#Types	#LF
Pattern	13	147	16	115
KB	7	7	25	26

Table 4: Number of labeling functions and the relation types they can annotated w.r.t. two kinds of information

CoreNLP tool (Manning et al., 2014) to generate entity mentions and get POS tags for both datasets. Brown clusters (Brown et al., 1992) are derived for each corpus using public implementation². All these features are shared with all compared methods in our experiments.

Labeling Functions. In our experiments, labeling functions are employed to encode two kinds of supervision information. One is knowledge base, the other is handcrafted domain-specific patterns. For domain-specific patterns, we manually design a number of labeling functions³; for knowledge base, annotations are generated following the procedure in (Ren et al., 2016; Riedel et al., 2010).

Regarding two kinds of supervision information, the statistics of the labeling functions are summarized in Table 4. We can observe that heuristic patterns can identify more relation types for KBP datasets, while for NYT datasets, knowledge base can provide supervision for more relation types. This observation aligns with our intuition that single kind of information might be insufficient while different kinds of information can complement each other.

We further summarize the statistics of annotations in Table 6. It can be observed that a large portion of instances is only annotated as `None`, while lots of conflicts exist among other instances. This phenomenon justifies the motivation to employ true label discovery model to resolve the conflicts among supervision. Also, we can observe most conflicts involve `None` type, accordingly, our proposed method should have more advantages over traditional true label discovery methods on the relation extraction task comparing to the relation classification task that excludes `None` type.

4.2 Compared Methods

We compare REHESSION with below methods: `FIGER` (Ling and Weld, 2012) adopts multi-label

²<https://github.com/percyliang/brown-cluster>

³pattern-based labeling functions can be accessed at: <https://github.com/LiyuanLucasLiu/ReHession>

Method	Relation Extraction						Relation Classification	
	NYT			Wiki-KBP			NYT	Wiki-KBP
	Prec	Rec	F1	Prec	Rec	F1	Accuracy	Accuracy
NL+FIGER	0.2364	0.2914	0.2606	0.2048	0.4489	0.2810	0.6598	0.6226
NL+BFK	0.1520	0.0508	0.0749	0.1504	0.3543	0.2101	0.6905	0.5000
NL+DSL	0.4150	0.5414	0.4690	0.3301	0.5446	0.4067	0.7954	0.6355
NL+MultiR	0.5196	0.2755	0.3594	0.3012	0.5296	0.3804	0.7059	0.6484
NL+FCM	0.4170	0.2890	0.3414	0.2523	0.5258	0.3410	0.7033	0.5419
NL+CoType-RM	0.3967	0.4049	0.3977	0.3701	0.4767	0.4122	0.6485	0.6935
TD+FIGER	0.3664	0.3350	0.3495	0.2650	0.5666	0.3582	0.7059	0.6355
TD+BFK	0.1011	0.0504	0.0670	0.1432	0.1935	0.1646	0.6292	0.5032
TD+DSL	0.3704	0.5025	0.4257	0.2950	0.5757	0.3849	0.7570	0.6452
TD+MultiR	0.5232	0.2736	0.3586	0.3045	0.5277	0.3810	0.6061	0.6613
TD+FCM	0.3394	0.3325	0.3360	0.1964	0.5645	0.2914	0.6803	0.5645
TD+CoType-RM	0.4516	0.3499	0.3923	0.3107	0.5368	0.3879	0.6409	0.6890
REHESSION	0.4122	0.5726	0.4792	0.3677	0.4933	0.4208	0.8381	0.7277

Table 5: Performance comparison of relation extraction and relation classification

Dataset	Wiki-KBP	NYT
Total Number of RM	225977	530767
RM annotated as None	100521	356497
RM with conflicts	32008	58198
Conflicts involving None	30559	38756

Table 6: Number of relation mentions (RM), relation mentions annotated as None, relation mentions with conflicting annotations and conflicts involving None

learning with Perceptron algorithm.

BFK (Bunescu and Mooney, 2005) applies bag-of-feature kernel to train a support vector machine;

DSL (Mintz et al., 2009) trains a multi-class logistic classifier⁴ on the training data;

MultiR (Hoffmann et al., 2011) models training label noise by multi-instance multi-label learning;

FCM (Gormley et al., 2015) performs compositional embedding by neural language model.

CoType-RM (Ren et al., 2016) adopts partial-label loss to handle label noise and train the extractor.

Moreover, two different strategies are adopted to feed heterogeneous supervision to these methods. The first is to keep all noisy labels, marked as ‘NL’. Alternatively, a true label discovery method, Investment (Pasternack and Roth, 2010), is applied to resolve conflicts, which is based on the source consistency assumption and iteratively updates inferred true labels and label functions’ reliabilities. Then, the second strategy is to only feed the inferred true labels, referred as ‘TD’.

⁴We use liblinear package from <https://github.com/cjlin1/liblinear>

Universal Schemas (Riedel et al., 2013) is proposed to unify different information by calculating a low-rank approximation of the annotations \mathcal{O} . It can serve as an alternative of the Investment method, i.e., selecting the relation type with highest score in the low-rank approximation as the true type. But it doesn’t explicitly model noise and not fit our scenario very well. Due to the constraint of space, we only compared our method to Investment in most experiments, and Universal Schemas is listed as a baseline in Sec. 4.4. Indeed, it performs similarly to the Investment method.

Evaluation Metrics. For relation classification task, which excludes None type from training / testing, we use the classification accuracy (Acc) for evaluation, and for relation extraction task, precision (Prec), recall (Rec) and F1 score (Bunescu and Mooney, 2005; Bach and Badaskar, 2007) are employed. Notice that both relation extraction and relation classification are conducted and evaluated in *sentence-level* (Bao et al., 2014).

Parameter Settings. Based on the semantic meaning of proficient subset, we set ϕ_2 to $1/|\mathcal{R} \cup \{\text{None}\}|$, i.e., the probability of generating right label with random guess. Then we set ϕ_1 to $1 - \phi_2$, $\lambda_1 = \lambda_2 = 1$, and the learning rate $\alpha = 0.025$. As for other parameters, they are tuned on the validation sets for each dataset. Similarly, all parameters of compared methods are tuned on validation set, and the parameters achieving highest F1 score are chosen for relation extraction.

Relation Mention	REHESSION	Investment & Universal Schemas
<i>Ann Demeulemeester</i> (born 1959 , Waregem , Belgium) is ...	born-in	None
<i>Raila Odinga</i> was born at ..., in Maseno, Kisumu District, ...	born-in	None
<i>Ann Demeulemeester</i> (elected 1959 , Waregem , Belgium) is ...	None	None
<i>Raila Odinga</i> was examined at ..., in Maseno, Kisumu District, ...	None	None

Table 7: Example output of true label discovery. The first two relation mentions come from Wiki-KBP, and their annotations are {born-in, None}. The last two are created by replacing key words of the first two. Key words are marked as bold and entity mentions are marked as Italics.

4.3 Performance Comparison

Given the experimental setup described above, the averaged evaluation scores in 10 runs of relation classification and relation extraction on two datasets are summarized in Table 5.

From the comparison, it shows that NL strategy yields better performance than TD strategy, since the true labels inferred by Investment are actually wrong for many instances. On the other hand, as discussed in Sec. 4.4, our method introduces context-awareness to true label discovery, while the inferred true label guides the relation extractor achieving the best performance. This observation justifies the motivation of avoiding the source consistency assumption and the effectiveness of proposed true label discovery model.

One could also observe the difference between REHESSION and the compared methods is more significant on the NYT dataset than on the Wiki-KBP dataset. This observation accords with the fact that the NYT dataset contains more conflicts than KBP dataset (see Table 6), and the intuition is that our method would have more advantages on more conflicting labels.

Among four tasks, the relation classification of Wiki-KBP dataset has highest label quality, i.e. conflicting label ratio, but with least number of training instances. And CoType-RM and DSL reach relatively better performance among all compared methods. CoType-RM performs much better than DSL on Wiki-KBP relation classification task, while DSL gets better or similar performance with CoType-RM on other tasks. This may be because the representation learning method is able to generalize better, thus performs better when the training set size is small. However, it is rather vulnerable to the noisy labels compared to DSL. Our method employs embedding techniques, and also integrates context-aware true label dis-

Dataset & Method	Prec	Rec	F1	Acc	
Wiki-KBP	Ori	0.3677	0.4933	0.4208	0.7277
	TD	0.3032	0.5279	0.3850	0.7271
	US	0.3380	0.4779	0.3960	0.7268
NYT	Ori	0.4122	0.5726	0.4792	0.8381
	TD	0.3758	0.4887	0.4239	0.7387
	US	0.3573	0.5145	0.4223	0.7362

Table 8: Comparison among REHESSION (Ori), REHESSION-US (US) and REHESSION-TD (TD) on relation extraction and relation classification

covery to de-noise labels, making the embedding method rather robust, thus achieves the best performance on all tasks.

4.4 Case Study

Context Awareness of True Label Discovery.

Although Universal Schemas does not adopted the source consistency assumption, but it’s conducted in document-level, and is context-agnostic in our sentence-level setting. Similarly, most true label discovery methods adopt the source consistency assumption, which means if they trust a labeling function, they would trust it on all annotations. And our method infers true labels in a context-aware manner, which means we only trust labeling functions on matched contexts.

For example, Investment and Universal Schemas refer None as true type for all four instances in Table 7. And our method infers born-in as the true label for the first two relation mentions; after replacing the matched contexts (**born**) with other words (**elected** and **examined**), our method no longer trusts born-in since the modified contexts are no longer matched, then infers None as the true label. In other words, our proposed method infer the true label in a context aware manner.

Effectiveness of True Label Discovery. We explore the effectiveness of the proposed context-aware true label discovery component by comparing REHESSION to its variants REHESSION-TD and REHESSION-US, which uses Investment or Universal Schemas to resolve conflicts. The averaged evaluation scores are summarized in Table 8. We can observe that REHESSION significantly outperforms its variants. Since the only difference between REHESSION and its variants is the model employed to resolve conflicts, this gap verifies the effectiveness of the proposed context-aware true label discovery method.

5 Related Work

5.1 Relation Extraction

Relation extraction aims to detect and categorize semantic relations between a pair of entities. To alleviate the dependency of annotations given by human experts, weak supervision (Bunescu and Mooney, 2007; Etzioni et al., 2004) and distant supervision (Ren et al., 2016) have been employed to automatically generate annotations based on knowledge base (or seed patterns/instances). Universal Schemas (Riedel et al., 2013; Verga et al., 2015; Toutanova et al., 2015) has been proposed to unify patterns and knowledge base, but it’s designed for document-level relation extraction, i.e., not to categorize relation types based on a specific context, but based on the whole corpus. Thus, it allows one relation mention to have multiple true relation types; and does not fit our scenario very well, which is sentence-level relation extraction and assumes one instance has only one relation type. Here we propose a more general framework to consolidate heterogeneous information and further refine the true label from noisy labels, which gives the relation extractor potential to detect more types of relations in a more precise way.

Word embedding has demonstrated great potential in capturing semantic meaning (Mikolov et al., 2013), and achieved great success in a wide range of NLP tasks like relation extraction (Zeng et al., 2014; Takase and Inui, 2016; Nguyen and Grishman, 2015). In our model, we employed the embedding techniques to represent context information, and reduce the dimension of text features, which allows our model to generalize better.

5.2 Truth Label Discovery

True label discovery methods have been developed to resolve conflicts among multi-source information under the assumption of source consistency (Li et al., 2016; Zhi et al., 2015). Specifically, in the *spammer-hammer* model (Karger et al., 2011), each source could either be a spammer, which annotates instances randomly; or a hammer, which annotates instances precisely. In this paper, we assume each labeling function would be a hammer on its proficient subset, and would be a spammer otherwise, while the proficient subsets are identified in the embedding space.

Besides data programming, socratic learning (Varma et al., 2016) has been developed to conduct binary classification under heterogeneous supervi-

sion. Its true label discovery module supervises the discriminative module in label level, while the discriminative module influences the true label discovery module by selecting a feature subset. Although delicately designed, it fails to make full use of the connection between these modules, i.e., not refine the context representation for classifier. Thus, its discriminative module might suffer from the overwhelming size of text features.

6 Conclusion and Future Work

In this paper, we propose REHENSION, an embedding framework to extract relation under heterogeneous supervision. When dealing with heterogeneous supervisions, one unique challenge is how to resolve conflicts generated by different labeling functions. Accordingly, we go beyond the “source consistency assumption” in prior works and leverage context-aware embeddings to induce proficient subsets. The resulting framework bridges true label discovery and relation extraction with context representation, and allows them to mutually enhance each other. Experimental evaluation justifies the necessity of involving context-awareness, the quality of inferred true label, and the effectiveness of the proposed framework on two real-world datasets.

There exist several directions for future work. One is to apply transfer learning techniques to handle label distributions’ difference between training set and test set. Another is to incorporate OpenIE methods to automatically find domain-specific patterns and generate pattern-based labeling functions.

7 Acknowledgments

Research was sponsored in part by the U.S. Army Research Lab. under Cooperative Agreement No. W911NF-09-2-0053 (NSCTA), National Science Foundation IIS-1320617, IIS 16-18481, and NSF IIS 17-04532, and grant 1U54GM114838 awarded by NIGMS through funds provided by the trans-NIH Big Data to Knowledge (BD2K) initiative (www.bd2k.nih.gov). The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies of the U.S. Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

- Nguyen Bach and Sameer Badaskar. 2007. A review of relation extraction. *Literature review for Language and Statistics II*.
- Junwei Bao, Nan Duan, Ming Zhou, and Tiejun Zhao. 2014. Knowledge-based question answering as machine translation. *Cell*, 2(6).
- Peter F Brown, Peter V Desouza, Robert L Mercer, Vincent J Della Pietra, and Jenifer C Lai. 1992. Class-based n-gram models of natural language. *Computational linguistics*, 18(4):467–479.
- Razvan Bunescu and Raymond Mooney. 2007. Learning to extract relations from the web using minimal supervision. In *ACL*.
- Razvan Bunescu and Raymond J Mooney. 2005. Subsequence kernels for relation extraction. In *NIPS*, pages 171–178.
- Andrew Carlson, Justin Betteridge, Richard C Wang, Estevam R Hruschka Jr, and Tom M Mitchell. 2010. Coupled semi-supervised learning for information extraction. In *Proceedings of the third ACM international conference on Web search and data mining*, pages 101–110. ACM.
- Hong-Woo Chun, Yoshimasa Tsuruoka, Jin-Dong Kim, Rie Shiba, Naoki Nagata, Teruyoshi Hishiki, and Jun’ichi Tsujii. 2006. Extraction of gene-disease relations from medline using domain dictionaries and machine learning. In *Pacific Symposium on Biocomputing*, volume 11, pages 4–15.
- Joe Ellis, Xuansong Li, Kira Griffitt, Stephanie Strassel, and Jonathan Wright. 2012. Linguistic resources for 2013 knowledge base population evaluations. In *TAC*.
- Oren Etzioni, Michael Cafarella, Doug Downey, Stanley Kok, Ana-Maria Popescu, Tal Shaked, Stephen Soderland, Daniel S Weld, and Alexander Yates. 2004. Web-scale information extraction in know-it-all:(preliminary results). In *Proceedings of the 13th international conference on World Wide Web*, pages 100–110. ACM.
- Katrin Fundel, Robert Küffner, and Ralf Zimmer. 2007. Relexrelation extraction using dependency parse trees. *Bioinformatics*, 23(3):365–371.
- Matthew R Gormley, Mo Yu, and Mark Dredze. 2015. Improved relation extraction with feature-rich compositional embedding models. *arXiv preprint arXiv:1505.02419*.
- Zellig S Harris. 1954. Distributional structure. *Word*, 10(2-3):146–162.
- 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.
- David R Karger, Sewoong Oh, and Devavrat Shah. 2011. Iterative learning for reliable crowdsourcing systems. In *Advances in neural information processing systems*, pages 1953–1961.
- Yaliang Li, Jing Gao, Chuishi Meng, Qi Li, Lu Su, Bo Zhao, Wei Fan, and Jiawei Han. 2016. A survey on truth discovery. *SIGKDD Explor. Newsl.*, 17(2):1–16.
- Xiao Ling and Daniel S Weld. 2012. Fine-grained entity recognition. In *AAAI*. Citeseer.
- Christopher D Manning, Mihai Surdeanu, John Bauer, Jenny Rose Finkel, Steven Bethard, and David McClosky. 2014. The stanford corenlp natural language processing toolkit. In *ACL (System Demonstrations)*, pages 55–60.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- 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.
- Ndapandula Nakashole, Martin Theobald, and Gerhard Weikum. 2011. Scalable knowledge harvesting with high precision and high recall. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 227–236. ACM.
- Thien Huu Nguyen and Ralph Grishman. 2015. Combining neural networks and log-linear models to improve relation extraction. *arXiv preprint arXiv:1511.05926*.
- Jeff Pasternack and Dan Roth. 2010. Knowing what to believe (when you already know something). In *Proceedings of the 23rd International Conference on Computational Linguistics*, pages 877–885. Association for Computational Linguistics.
- Alexander J Ratner, Christopher M De Sa, Sen Wu, Daniel Selsam, and Christopher Ré. 2016. Data programming: Creating large training sets, quickly. In *Advances in Neural Information Processing Systems*, pages 3567–3575.
- Xiang Ren, Ahmed El-Kishky, Chi Wang, Fangbo Tao, Clare R Voss, and Jiawei Han. 2015. Clustype: Effective entity recognition and typing by relation phrase-based clustering. In *Proceedings of the 21th*

- ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 995–1004. ACM.
- Xiang Ren, Zeqiu Wu, Wenqi He, Meng Qu, Clare R Voss, Heng Ji, Tarek F Abdelzaher, and Jiawei Han. 2016. Cotype: Joint extraction of typed entities and relations with knowledge bases. *arXiv preprint arXiv:1610.08763*.
- 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.
- Sebastian Riedel, Limin Yao, Andrew McCallum, and Benjamin M Marlin. 2013. Relation extraction with matrix factorization and universal schemas. In *HLT-NAACL*, pages 74–84.
- Alexander Schutz and Paul Buitelaar. 2005. Relext: A tool for relation extraction from text in ontology extension. In *International semantic web conference*, volume 2005, pages 593–606. Springer.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958.
- Sho Takase and Naoaki Okazaki Kentaro Inui. 2016. Composing distributed representations of relational patterns. In *Proceedings of ACL*.
- Kristina Toutanova, Danqi Chen, Patrick Pantel, Hoi-fung Poon, Pallavi Choudhury, and Michael Gamon. 2015. Representing text for joint embedding of text and knowledge bases. In *EMNLP*, volume 15, pages 1499–1509.
- Christophe Van Gysel, Maarten de Rijke, and Evangelos Kanoulas. 2016a. Learning latent vector spaces for product search. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*, pages 165–174. ACM.
- Christophe Van Gysel, Maarten de Rijke, and Marcel Worring. 2016b. Unsupervised, efficient and semantic expertise retrieval. In *Proceedings of the 25th International Conference on World Wide Web*, pages 1069–1079. International World Wide Web Conferences Steering Committee.
- Paroma Varma, Bryan He, Dan Iter, Peng Xu, Rose Yu, Christopher De Sa, and Christopher Ré. 2016. Socratic learning: Correcting misspecified generative models using discriminative models. *arXiv preprint arXiv:1610.08123*.
- Patrick Verga, David Belanger, Emma Strubell, Benjamin Roth, and Andrew McCallum. 2015. Multilingual relation extraction using compositional universal schema. *arXiv preprint arXiv:1511.06396*.
- Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, Jun Zhao, et al. 2014. Relation classification via convolutional deep neural network. In *COLING*, pages 2335–2344.
- Shi Zhi, Bo Zhao, Wenzhu Tong, Jing Gao, Dian Yu, Heng Ji, and Jiawei Han. 2015. Modeling truth existence in truth discovery. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1543–1552. ACM.