

On the use of Silver Standard Data for Zero-shot Classification Tasks in Information Extraction

Jianwei Wang¹, Tianyin Wang², Ziqian Zeng^{1*}

¹Shien-Ming Wu School of Intelligent Engineering, South China University of Technology, China,

²Information Network Institute, Carnegie Mellon University

wiwjwilliam@mail.scut.edu.cn, tianyinw@andrew.cmu.edu, zqzeng@scut.edu.cn

Abstract

The superior performance of supervised classification methods in the information extraction (IE) area heavily relies on a large amount of gold standard data. Recent zero-shot classification methods converted the task to other NLP tasks (e.g., textual entailment) and used off-the-shelf models of these NLP tasks to directly perform inference on the test data without using a large amount of IE annotation data. A potentially valuable by-product of these methods is the large-scale silver standard data, i.e., pseudo-labeled data by the off-the-shelf models of other NLP tasks. However, there is no further investigation into the use of these data. In this paper, we propose a new framework, Clean-LaVe, which aims to utilize silver standard data to enhance the zero-shot performance. Clean-LaVe includes four phases: (1) Obtaining silver data; (2) Identifying relatively clean data from silver data; (3) Finetuning the off-the-shelf model using clean data; (4) Inference on the test data. The experimental results show that Clean-LaVe can outperform the baseline by 5% and 6% on TACRED and Wiki80 dataset in the zero-shot relation classification task, and by 3% ~7% on Smile (Korean and Polish) in the zero-shot cross-lingual relation classification task, and by 8% on ACE05-E+ in the zero-shot event argument classification task.

Keywords: Zero-shot Classification, Information Extraction, Noise Label Learning, Cross-lingual NLP

1. Introduction

Information Extraction (IE) is a fundamental problem in natural language processing. The predominant approaches to solve IE tasks are supervised methods (Shen et al., 2022; Zhu and Li, 2022; Zhong and Chen, 2021; Lyu and Chen, 2021; Yang et al., 2019; Lu et al., 2023a). Supervised methods require a large amount of gold standard data, which restricts their applications to real-world scenarios where large-scale annotated data are not available. Zero-shot methods (Lyu et al., 2021; Sainz et al., 2021) have been proposed to alleviate this issue. We focus on zero-shot classification tasks in IE such as relation extraction (RE), cross-lingual relation extraction, and event argument classification (EAC). Zero-shot (cross-lingual) RE aims to identify the semantic relation between two entities in unstructured texts without using any annotated RE data (in the target language). Zero-shot EAC aims to assign roles to argument spans using any annotated EAC data.

Recent works (Sainz et al., 2021, 2022a,b; Lu et al., 2022) attempt to convert the zero-shot RE task and EAC task to other NLP tasks and used off-the-shelf models of these tasks to infer the relation types without using a large amount of RE or EAC annotated data. Sainz et al. (2021) used a well-trained textual entailment (TE) model to directly infer relation types on the RE test data by convert-

ing a RE task to a TE task. Their subsequent work (Sainz et al., 2022a) also used a TE model to infer argument roles on the test data by converting an EAC task to a TE task. This series of work is named LaVeEntail. SURE (Lu et al., 2022) formulated a RE task to a summarization task, and used a small amount of RE annotated data to finetune a well-trained summarization model thus it can perform inference on the RE test data. We term the TE model and summarization model in the above methods as pre-trained models. The concept of pretraining derives from transfer learning (Pan and Yang, 2009). A model is first pre-trained on the source task, i.e., textual entailment or summarization, and then finetuned on the target task, i.e., relation extraction and event argument classification.

Since pre-trained models can directly infer the categories of unlabeled data, they can serve as low-cost annotators, producing large-scale silver standard data. However, in the above works, silver standard data are not well-exploited. The straightforward way to utilize them is to directly train a supervised classifier on silver standard data. However, the performance is usually unsatisfactory due to the noisy nature of silver standard data. Learning with noisy labels has been well studied in the literature (Frénay and Verleysen, 2013; Algan and Ulusoy, 2021; Han et al., 2020). One direction is to develop noise-robust losses that can mitigate the effect of noisy labels (Ghosh et al., 2017; Zhang and Sabuncu, 2018; Charoenphakdee et al., 2019;

* Corresponding author

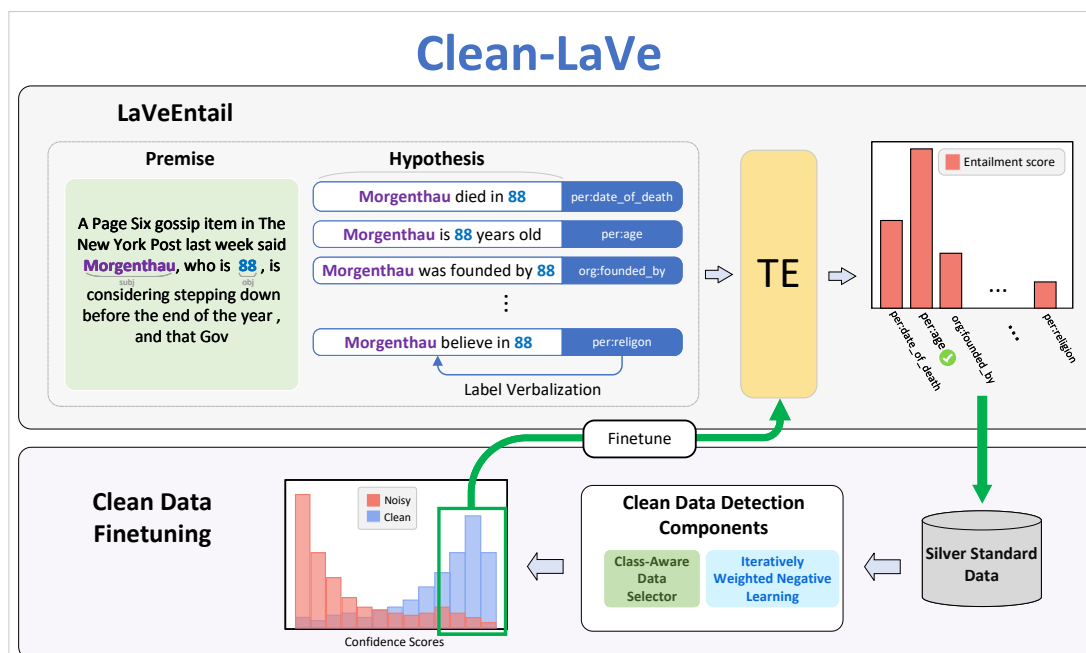


Figure 1: The diagram shows the procedure of Clean-LaVe in the zero-shot relation extraction task. First, we apply LaVeEntail on an unlabeled dataset, obtaining standard silver data. The clean data detection module uses confidence scores to distinguish the clean and noisy samples. The selected clean data are used to finetune the TE model. Finally, we use this TE model to infer relation types on the test set.

Kim et al., 2019; Lyu and Tsang, 2019; Menon et al., 2020; Thulasidasan et al., 2019). Another direction is to identify noisy data or clean data and deal with them separately either by re-weighting or converting to a semi-supervised learning task. (Han et al., 2018; Jiang et al., 2018; Arazo et al., 2019; Kim et al., 2019; Shu et al., 2019; Yao et al., 2019; Li et al., 2020). The setting of traditional noisy labels learning does not consider the existence of a pre-trained model. Is there a better way to utilize silver standard data when a pre-trained model is available? According to our best knowledge, there is no further investigation on the use of potentially valuable silver standard data when there exists a pre-trained model.

In this paper, we propose a novel framework called **Clean-LaVe**. The framework involves two main steps: firstly, detecting a small subset of clean data from the silver standard data using the clean data detection module, and secondly, utilizing the selected clean data to finetune the pre-trained model. The overall procedure is illustrated in Figure 1.

Within the clean data detection module, we introduce a **iteratively weighted negative learning** algorithm to obtain confidence scores that allow us to distinguish clean data from noisy data. The original negative learning algorithm (Kim et al., 2019) only performs well when the dataset is balanced (Huang et al., 2022b; Lu et al., 2023b). However, in real-world scenarios, this assumption may not

hold. To address this issue, we introduce an iterative weighting strategy to allow the algorithm to handle an imbalance dataset.

To select clean data, confidence scores serve as a straightforward metric (Kim et al., 2019). However, data from certain classes possibly have high confidence scores while others yield low scores. In such cases, selecting data solely based on confidence scores may lead to a narrow range of classes being selected, potentially harming overall performance. To mitigate this issue, we develop a **class-aware data selector** that enables the selection of data from a broader range of classes.

Clean-LaVe is a general framework that can be used in scenarios where a pre-trained model serves as an annotator. In our experiments, we demonstrate the usability of Clean-LaVe in various zero-shot classification tasks, such as zero-shot RE, zero-shot cross-lingual RE, and zero-shot EAC. For these aforementioned zero-shot tasks, we utilize a TE model as the pre-trained model to acquire silver standard data.

Our contributions are summarized as follows,

- We propose Clean-LaVe to first detect a small amount of clean data which are later used to finetune the pre-trained model. We then use the finetuned model to infer the categories on the test data.
- We propose a clean data detection module that enhances the selection process through Iteratively Weighted Negative Learning and Class-

Aware Data Selector.

- The experimental results demonstrate that our method can outperform the baseline by a large margin on various zero-shot classification tasks. The code is shared in https://github.com/ZeroNLP/Clean_LaVe.

2. Related Work

Zero-shot Relation Extraction. In classical zero-shot learning settings, the classes in the training and test phases are disjoint. In the training phase, it requires a large amount of annotated RE data from seen classes. In the test phase, zero example for each unseen relation type during the test phase is needed. Recent works (Levy et al., 2017; Obamuyide and Vlachos, 2018; Zhao et al., 2023) formulated the zero-shot RE to other NLP tasks such as reading comprehension (Levy et al., 2017; Zhao et al., 2023) and textual entailment (Obamuyide and Vlachos, 2018).

However, the classical zero-shot setting still requires a large amount of annotated data in the training phase. Recent works (Goswami et al., 2020; Sainz et al., 2021; Tran et al., 2021; Lu et al., 2022; Rahimi and Surdeanu, 2023) push the zero-shot setting to an extreme case where annotated data is not available in the training phase. They obtained supervision from other available resources such as language models (Tran et al., 2021; Zhang et al., 2023), relation descriptions, and off-the-shelf models from other NLP tasks. QA4RE (Zhang et al., 2023) aligns RE with question answering. QA4IE (Zhang et al., 2023) is the state-of-the-art method in the zero-shot RE task, primarily owing to the powerful capacity of Large Language Models. In this paper, we focus on this extreme zero-shot setting.

Zero-shot Cross-lingual Information Extraction. Existing approaches to zero-shot cross-lingual Information Extraction (IE) can be categorized into three main types: translation-based (Lou et al., 2022), feature-based (Huang et al., 2022a; Ma et al., 2023), and distillation-based methods (Wu et al., 2020; Ma et al., 2022). However, all of these methods require significant manual effort to obtain labeled data for the source languages, which could potentially be replaced by readily available resources in the target language domain, such as off-the-shelf TE models.

Zero-shot Event Argument Classification. Existing zero-shot event argument classification tasks are based on label representations (Huang et al., 2018; Zhang et al., 2021b), reading comprehension (Liu et al., 2020; Lyu et al., 2021; Mehta et al., 2022), and pre-trained language models (Huang

et al., 2022a; Lin et al., 2023). Lin et al. (2023) is the state-of-the-art zero-shot EAC method, which prompts the pre-trained language models and regularizes the prediction by global constraints.

Learning with Noisy Labels. One direction of learning with noisy labels is to develop noise-robust loss. The widely-used cross entropy (CE) loss in classification tasks has been shown to be not robust against label noise (Ghosh et al., 2017). Several noise-robust losses have been proposed for training models with noisy labels (Reed et al., 2015; Zhang and Sabuncu, 2018; Wang et al., 2019; Ma et al., 2020; Menon et al., 2020; Jin et al., 2021; Zhou and Chen, 2021), which were shown to be more robust than CE. However, since current deep networks have a large number of parameters, these methods can still memorize the noisy labels given sufficient training time (Zhang et al., 2017).

Another direction is to identify noisy data or clean data and cope with them separately either by re-weighting them or converting the problem to a semi-supervised learning task. (Arpit et al., 2017; Charoenphakdee et al., 2019) found out the memorization effect which is stated as although deep networks can memorize noise data, they tend to learn simple patterns first. Based on the memorization effect (Arpit et al., 2017; Zhang et al., 2021a), many methods separate clean and noisy samples by using loss value (Han et al., 2018; Jiang et al., 2018; Arazo et al., 2019; Shu et al., 2019; Yao et al., 2019; Li et al., 2020) or forgetting events (Malach and Shalev-Shwartz, 2017; Yu et al., 2019). The setting of traditional noisy labels learning does not consider the existence of a pre-trained model.

3. Method

Due to the versatility of LaVeEntail (Sainz et al., 2021, 2022a) across multiple tasks such as RE and EAC, we employ LaVeEntail as the backbone to obtain silver standard data. We will introduce the LaVeEntail method in §3.1; the clean data detection module of Clean-LaVe in §3.2; the finetuning and inference stage in Clean-LaVe in §3.3.

3.1. LaVeEntail

LaVeEntail (Sainz et al., 2021, 2022a) includes two processes for relation extraction and event argument extraction, i.e., label verbalization and textual entailment model inference.

3.1.1. Label Verbalization

The label verbalization process creates templates of classes (i.e., relation types and argument roles) and then uses them to generate hypotheses. The

templates can be easily created because relation labels and argument roles naturally implicate such verbalization templates. For example, the relation `per:schools_attended` can be verbalized as `{subj} studied in {obj}`, where `{subj}` and `{obj}` are placeholders for subject and objective entities. For example, `giver` can be expressed as `{arg} gave something to someone`, where `{arg}` is the placeholder for an argument span.

3.1.2. Textual Entailment Model Inference

For each input sentence, LaVeEntail constructed hypotheses that are generated by verbalization templates of all relation types (or argument roles), and fed them to a TE model, and obtained entailment scores of all hypotheses. LaVeEntail inferred that the predicted relation (or role) type of the input sentence is the relation (or role) type whose hypothesis yields the highest entailment score. Figure 1 shows the inference procedure of relation extraction.

Entity type information is helpful to infer relation types (Tran et al., 2020). A relation naturally indicates entity types of subject and object. For instance, the relation `per:city_of_death` implicates that the entity type of subject and object should be `PERSON` and `CITY` respectively. In the inference stage, when the entity type information is given, we could rule out some relation types that are impossible to be ground truth. LaVeEntail created entity type constraint(s) for each relation according to the meaning of the relation. If the entity types in the input sentence do not match the entity type constraints of a relation, then the entailment score(s) of all hypotheses related to this relation is set to zero. In the case where there is no relation between two entities, a threshold-based approach is used to detect `no_relation`. If the entailment scores of all hypotheses are less than a threshold, the prediction is `no_relation`.

3.2. Clean Data Detection

The clean data detection module aims to select relatively clean data for subsequent finetuning from silver standard data annotated by LaVeEntail. To alleviate the impact of imbalanced noisy data, we introduce an iteratively weighted negative learning (IWNL) algorithm. Additionally, we employ a class-aware data selector (CADS) to choose clean samples from a boarder range of classes.

3.2.1. Iteratively Weighted Negative Learning

Negative Learning (Kim et al., 2019) loss is robust to noise. Different from positive learning loss (e.g., cross entropy loss) which tells the model

what is correct, the negative learning loss provides the model with the complementary label(s), telling what is not correct, e.g., the input image is not a dog. The complementary label is randomly selected from the label space excluding the input label (possibly noisy). For noisy data, the probability of selecting the ground truth as the complementary label is low. Hence, using negative learning loss can decrease the risk of overfitting noisy labels. The formula of NL loss is shown as follows,

$$\mathcal{L}_{neg} = - \sum_{d \in D} \sum_{i=1}^{|\mathcal{Y}|} \hat{\mathbf{y}}_i^d \log(1 - \mathbf{p}_i^d), \quad (1)$$

where d is a sample in the dataset D , $|\mathcal{Y}|$ is the number of relation types, $\hat{\mathbf{y}}^d$ is a one-hot vector with the complementary label being one, $\hat{\mathbf{y}}_i^d$ is the i -th element of $\hat{\mathbf{y}}^d$, \mathbf{p}^d is the output probability distribution of a sample d , and \mathbf{p}_i^d is the i -th element of \mathbf{p}^d .

The original NL loss in equation (1) treats each class equally, which may not be appropriate when dealing with real-world datasets that exhibit severe class imbalance. In these datasets, majority classes often have a significantly higher number of data samples compared to minority classes. Consequently, the model encounters much fewer samples in the minority classes, leading to underfitting (i.e., high loss values) during the training process. It poses a challenge to distinguish between clean and noisy samples in minority classes as they both have high loss values. We propose a iteratively weighted NL loss to alleviate this issue, giving more weight on minority classes.

$$\mathcal{L}_{neg}^j = - \sum_{d \in D} \sum_{i=1}^{|\mathcal{Y}|} w_i^j \cdot \hat{\mathbf{y}}_i^d \log(1 - \mathbf{p}_i^d) \quad (2)$$

$$w_i^j = w_i^{j-1} \cdot e^{1 - \frac{c_i^{j-1}}{c_A^{j-1}}} \quad (3)$$

$$w_i^0 = \frac{\sum_{k=1}^{|\mathcal{Y}|} c_k^0}{c_i^0} \quad (4)$$

where \mathcal{L}_{neg}^j represents the negative loss for j -th epoch, w_i^j denotes the weight for class i in j -th epoch, which is dynamically updated by prediction in previous epoch according to equation (3). c_i^{j-1} is the quantity of class i in $j-1$ th epoch and c_A^{j-1} is the average quantity across all classes in $j-1$ th epoch. The quantity of class i is the number of samples that are predicted as class i . Initial weight w_i^0 is computed according to labels of silver data, as described in equation (4). According to Eq. (4) and Eq. (4), minority classes have more weight. If the dataset is initially balanced, our IWNL algorithm can degenerate to the original negative learning algorithm.

Iterative weight can fix the problem that initial weight may be wrong. Our silver dataset is noisy

and is likely to give wrong weight in initial. Wrong weight can degrade the performance of Clean-LaVe. Iterative weight can make some correction to wrong proportions in initial.

We use BERT (Devlin et al., 2019) as the relation classifier when using IWNL loss. Although it is possible to train a TE model as a relation classifier using IWNL loss, its performance falls short of that of BERT-based classifier. After being trained with IWNL loss, the classifier attempts to assign high confidence scores to clean data while give low confidence scores to noisy data. These confidence scores can be leveraged for subsequent data selection.

3.2.2. Class-Aware Data Selector

A straightforward approach to selecting clean data is sorting all samples according to their confidence scores and then selecting a fixed proportion η of whole data as the clean data set. Given that $\mathcal{S}(D_s)$ is the total confidence scores of all samples in D_s , and η is a hyperparameter representing selection proportion, the clean data set D_{clean} are selected as follows,

$$D_{clean} = \arg \max_{D_s: |D_s| = \eta \cdot |D_{silver}|} \mathcal{S}(D_s). \quad (5)$$

However, it does not consider class diversity. Samples in some classes can yield very high confidence scores while some classes have very low confidence scores. Large quantities of samples in those classes are selected in D_{clean} while some classes even do not have any clean data in D_{clean} , which harms performance badly.

We propose a class-aware data selection algorithm that considers confidence scores as well as class diversity. First, we select a proportion η of data with high confidence scores. This step can ensure that samples with low noise levels are selected. Next, we select m more samples to encourage diversity. η is used to select high confidence data and m is used to make supplement according to label distribution. For each class, we select some samples with high confidence scores in this class. The number of selected samples for a class is proportional to the number of samples that are predicted as the class. The class-aware data selection algorithm is presented in Algorithm 1.

The low confidence scores observed in certain classes can be attributed to two factors: either they are minority classes and suffer from underfitting, or the samples in these classes are noisy. The class-Aware data selector serves as a compensatory mechanism to mitigate the impact of underfitting in minority classes. However, class-aware data selectors carry the risk of inadvertently noisy data. This risk becomes even more pronounced when the dataset is balanced, as evidenced by the experimental results in Table 2.

Algorithm 1 Class-Aware Data Selector

Input: silver standard data set D_{silver} , proportion η , diversity number m , the set of classes \mathcal{C} , the total confidence scores function $\mathcal{S}(\cdot)$.

- 1: $D_{clean} = \emptyset$.
- 2: Obtain D_{clean} using Eq. 5 by setting the proportion to η .
- 3: $D_{rest} = D_{silver} - D_{clean}$, divide D_{rest} into $|\mathcal{C}|$ subsets according to class predictions. The subset for class c is denoted as D^c .
- 4: **for** c in \mathcal{C} **do**
- 5: $D_{clean}^c = \arg \max_{D_s: |D_s| = \frac{|D^c|}{|D_{rest}|} \cdot m} \mathcal{S}(D_s)$
- 6: $D_{clean} = D_{clean} \cup D_{clean}^c$
- 7: **end for**

Output: clean data set D_{clean} .

Algorithm 2 Finetuning and Inference

Input: silver standard data set D_{silver} , test set D_{test} , textual entailment model \mathcal{M} .

- 1: Obtain D_{clean} using Algorithm 1.
- 2: Generate premise hypothesis pairs dataset D'_{clean} based on D_{clean} .
- 3: Finetune \mathcal{M} using D'_{clean} , and obtain finetuned model \mathcal{M}' .
- 4: Use \mathcal{M}' to infer relation (role) types on D_{test} .

Output: relation (role) types of samples on D_{test} .

3.3. Finetuning and Inference

After running clean data detection algorithms, we obtain D_{clean} which consists of the input sentence and its relation (role) type pairs. Since the input formats of the TE and RE (or EAC) tasks are different, we need to convert D_{clean} to premise-hypothesis pairs so that we can use D_{clean} to finetune the TE model.

For each relation (or role) in the RE (or EAC) task form, we should create entailment, contradiction, and neutral hypothesis for the TE task. The entailment hypothesis is generated with the templates that describe the ground truth relation (or role), the neutral hypothesis is generated by randomly select a template that does not describe the ground truth relation (role) and the contradiction hypothesis is generated using the template “{subj} and {obj} are not related,” or “{arg}” is not an argument of {trg} where {trg} is the trigger word.

We use premise hypothesis pairs constructed from D_{clean} to finetune the off-the-shelf TE model. Finally, we use the finetuned TE model to infer relation (role) types on the test set. The complete algorithm is presented in Algorithm 2.

Dataset	Relation Types	Entity Types	Distribution	Instances		
				Train	Dev	Test
TACRED	42	17	Skewed	68124	22631	15509
Wiki80	80	29	Uniform	40320	10080	5600
Smiler-It	22	-	Skewed	73228	746	1510
Smiler-Po	21	-	Skewed	16651	180	344
Smiler-Kr	28	-	Skewed	18538	173	382
ACE05-E+	22	7	Skewed	4815	603	573

Table 1: The statistics of datasets.

4. Experiment

4.1. Experimental Settings

We follow the same zero-shot setting with LaVeEntail, all data used for training are unlabeled, and only 1% development set are available for adjusting hyperparameters. For the zero-shot cross-lingual setting, we do not use any annotated training data from both the source language and target language. Traditionally, zero-shot cross-lingual methods (Huang et al., 2022a; Ma et al., 2023, 2022; Wu et al., 2020) used a large number of annotated data from the source language.

For the zero-shot RE task, we evaluate our method on the TACRED (Zhong et al., 2018) and Wiki80 (Xiao, 2022) dataset. For the zero-shot cross-lingual RE task, we evaluate our method on the Smile (Seganti et al., 2021) dataset which contains 14 languages. We evaluate Clean-LaVe in three languages, i.e., Italian, Polish, and Korean. For the zero-shot EAC task, we evaluate our method on ACE05-E+ (Lin et al., 2020). The statistics of datasets are shown in Table 1.

The TE model we used for RE and EAC tasks is microsoft/deberta-v2-xlarge-mnli (He et al., 2021). We use mDeBERTa-v3-base-xnli-multilingual-nli-2mil7 (Laurer et al., 2022) for the cross-lingual RE task as it can process multiple languages.

For each dataset, we manually created verbalization templates for each relation or argument role, as well as entity type constraints. The constraints we used on TACRED dataset are different from those of LaVeEntail. We delete the constraints which leak the information of the ground truth. For example, there is only one relation type that has the constraint where the subject entity type is PERSON and the object entity type is TITLE. The sentence that satisfies this constraint has a very large probability of being inferred as `per:title` relation because other relations are ruled out. (Tran et al., 2020) showed that entity types are a strong inductive bias. However, in LaVeEntail, the inductive bias is not learned by the algorithm itself but by manually designed type constraints. It leads to artificially inflated performance, so we deleted those type constraints. Besides, we also conduct experiments using original type constraints in §4.5.

¹LaVeEntail direct infers on the test set and does not involve any training process, resulting in zero variance.

4.2. Compared Methods

To demonstrate the effectiveness of our method, we compare our model with the following baselines:

We consider training a supervised relation classification model on silver standard data using **CE** (Cross Entropy loss) and different noise-robust losses including **GCE** (Generalized Cross Entropy loss) (Zhang and Sabuncu, 2018), **SCE** (Symmetric Cross Entropy loss) (Wang et al., 2019), and **Co-Regularization** (Zhou and Chen, 2021) as baselines.

We also include the following representative noise-robust learning algorithms which identify noisy data or clean data and convert the problem to a semi-supervised learning problem as baseline methods: **O2U** (Overfitting to Underfitting) (Huang et al., 2019), and **DivideMix** (Li et al., 2020).

We also consider the SOTA zero-shot RE method **QA4RE** (Zhang et al., 2023) and the SOTA zero-shot EAC method **Global_Constraints** (Lin et al., 2023) as baselines. QA4RE is based on ChatGPT (AI, 2023), which is easily adapted to the zero-shot EAC task. Global_Constraints is delicately designed for the zero-shot EAC task.

LaVeEntail (Sainz et al., 2021) utilized an off-the-shelf textual entailment model to directly infer the test data. **Labeled Data Finetune** randomly select a proportion of labeled training data to finetune the off-the-shelf TE model. **Clean-LaVe** is our proposed method. Additionally, we conduct comparisons by removing the weighted negative learning module, the class-aware data selector, and both of them respectively to assess their impact on the results. **Silver-LaVe** can be considered as Clean-LaVe without a clean data detection module, which uses all silver standard data to finetune an off-the-shelf TE model.

4.3. Result Analysis

As shown in the first and second block of Table 2, Clean-LaVe outperforms noise-robust loss based methods and semi-supervised based noisy labels learning methods across all datasets. Directly applying noisy labels learning methods on silver standard data is straightforward but not effective. Hence, there is a need to investigate how to use silver data.

As shown in the third block, Clean-LaVe outperforms the SOTA methods by 3% ~15% on all datasets except on the Smiler-It. On Smiler-It, QA4RE outperforms Clean-LaVe by 1%. Despite facing a stronger competitor based on ChatGPT, Clean-LaVe delivers commendable overall performance.

As shown in the fourth block, Clean-LaVe can gain significant improvement compared to LaVeEntail by 10% ~16%. Additionally, our method is

	RE		Cross-lingual RE			EAC
	TACRED	Wiki80	Smiler-It	Smiler-Po	Smiler-Kr	ACE05-E+
CE	45.35±0.58	40.76±0.29	40.79±0.12	41.56±0.19	49.75±0.50	71.79±0.96
GCE (Zhang and Sabuncu, 2018)	45.93±0.67	41.28±0.61	47.27±0.21	<u>45.99±0.60</u>	53.35±0.49	71.61±0.79
SCE (Wang et al., 2019)	45.82±0.92	41.12±0.24	40.97±0.70	40.41±0.31	47.79±0.09	71.88±0.26
Co-Regularization (Zhou and Chen, 2021)	48.86±0.34	28.48±0.42	42.17±0.61	41.86±0.19	50.16±0.70	<u>72.93±0.17</u>
O2U (Huang et al., 2019)	47.52±0.81	42.62±0.03	41.12±0.23	44.47±0.66	49.67±0.49	69.83±0.06
DivideMix (Li et al., 2020)	49.78±0.80	<u>45.52±0.26</u>	41.94±0.78	43.79±0.69	52.48±0.80	69.13±0.64
Global_Constraints (Lin et al., 2023) - QA4RE (Zhang et al., 2023)	-	-	-	-	-	66.1*
	<u>58.55±0.05</u>	43.93±0.09	56.42±0.84	38.09±0.19	<u>56.08±0.73</u>	64.74±0.84
LaVeEntail ¹ (Sainz et al., 2021)	52.18	41.16	39.96	37.84	44.30	71.60
Labeled Data Finetune (1%)	56.61±1.29	47.39±0.33	51.85±0.96	46.44±0.66	47.33±0.91	76.21±1.50
Labeled Data Finetune (5%)	63.72±1.03	53.89±0.46	52.56±0.57	49.56±0.54	55.30±0.14	78.87±0.17
Silver-LaVe	54.67±0.58	44.57±0.31	48.91±0.55	50.60±0.38	54.64±0.81	80.18±0.08
Clean-LaVe	63.36 ±1.03	51.53±0.53	55.09±0.05	52.99 ±0.88	59.41 ±0.84	81.22 ±0.38
- Iteratively Weighted Negative Learning	58.66±0.93	48.44±0.44	54.20±0.97	48.09±0.59	57.18±0.95	78.07±0.82
- Class-Aware Data Selector	59.55±0.98	52.52 ±0.21	54.97±0.33	50.14±0.64	57.34±0.26	78.14±0.61
- Above Both	56.41±1.82	52.34±0.38	54.28±0.65	45.99±0.41	54.97±0.74	77.37±0.67

Table 2: Results of zero-shot classification tasks. We report the average of micro F1 scores in 3 runs. The best F1 scores are marked in **bold**. SOTA baselines are highlighted with underline. Results marked with * are retrieved from the original paper.

comparable to or even outperforms the supervised LaVeEntail with 5% labeled data.

As shown in the last block, we surprisingly find Silver-LaVe outperforms LaVeEntail by 2% ~13%. It indicates that, to some content, our proposed framework (i.e., finetuning pre-trained model with silver standard data) can be beneficial, regardless of the quality of silver standard data. Clean-LaVe outperforms Silver-LaVe, indicating the effectiveness of the clean data detection module.

We also provide results after removing Iteratively Weighted Negative Learning, Class-Aware Data Selector, and both of them respectively. After removing the IWNL component, we observe decreases in performance across all datasets, which validates the effectiveness of this component. After removing the CADS, we observe decreases in performance across all datasets except Wiki80, which validates the effectiveness of this component. Removing the CADS leads to a slight improvement (1%) on Wiki80. This improvement can be attributed to the fact that Wiki80 is a balanced dataset. The reason has been stated in §3.2.2.

4.4. Case Analysis

We conduct an in-depth analysis on TACRED regarding the effectiveness of Iteratively Weighted Negative Learning (IWNL) and Class-Aware Data Selector (CADS). The clean data detection accuracy (referred to as detection accuracy for brevity) is the percentage of clean data whose predictions are equal to ground truths. We sort the class according to the number of samples in the class in descending order and consider the former (latter) half as the majority (minority) classes. The pro-


portion η indicates how much proportion of data is selected as clean data.

Iteratively Weighted Negating Learning (IWNL) can alleviate the effect of underfitting and improve the clean data detection accuracy of minority classes. As depicted in Figure 2(a) (left), IWNL yields consistently higher detection accuracy scores than w/o IWNL on minority classes. Without IWNL, the detection accuracy for minority classes is almost negligible when the selection proportion is small. As depicted in Figure 2(a) (right), the performance of IWNL on majority classes is comparable to w/o IWNL.


Class-Aware Data Selector (CADS) can encourage clean samples from a broader range of classes. As depicted in Figure 2(b) (left), there are more orange bars than blue bars, indicating CADS selects samples from more classes, especially from minority classes. As depicted in Figure 2(b) (right), the detection accuracy scores of classes that are only selected by CADS are satisfying overall. For classes that are selected by both CADS and w/o CADS, the accuracy scores of some classes increase but some decrease after applying CADS. The possible reason for decreased accuracy is that it involves noisy data, as we have discussed in §3.2.2.

4.5. Full Constraints Comparison

As previously mentioned, we remove some constraints defined in LaVeEntail to prevent information leakage. In this section, we compare our method with LaVeEntail using full constraints. As table 3 is shown, our method still outperforms LaVeEntail given full constraints. Under full type



(a) Analysis of IWNL on minority and majority classes.



(b) Analysis of CADs.

Figure 2: Analysis of IWNL and CADs on TACRED.

constraints, LaVeEntail and Clean-LaVe obtain improvement compared to partial constraints results in table 2. But the improvement is inflated.

	Pr.	Rec.	F1
LaVeEntail ⁺	63.20*	59.80*	61.40*
Clean-LaVe ⁺	72.60	59.98	65.59 \uparrow 4.1

Table 3: The results of using original constraints defined in LaVeEntail. Upper + means full constraints and the results of LaVeEntail marked with * are retrieved from the original paper.

4.6. Cross-lingual Silver Standard Data

We further explore the potential of cross-lingual silver standard data. We combine the silver data from the source language (i.g., English) with the silver data from the target language and fine-tune the TE model. Note that we do not use any labeled data from the source language as well as the target language. Results show that by using cross-lingual silver data from English, Clean-LaVe can further improve 0.2% - 1.7%.

4.7. Hyper-parameter Analysis

As Section 3.2.2 mentioned, Class-Aware Data Selector introduces two hyper-parameters to control the selection. η controls the number of clean data

	Italian	Polish	Korean
LaVeEntail	39.96	37.84	44.30
Clean-LaVe	55.09	52.99	59.41
Clean-LaVe + En Silver	56.81 \uparrow 1.72	53.23 \uparrow 0.24	60.74 \uparrow 1.33

Table 4: The average of F1 scores in 3 runs of Clean-LaVe on the zero-shot cross-lingual RE task.

and m controls the number of samples from diverse classes. We analyse these hyper-parameters on 1% development set on each dataset.

Clean Data Selection Proportion η . We select $\eta \cdot |D_{silver}|$ data to finetune the TE model. The search range for η is $[0.01, 0.03, \dots, 0.1]$, while keeping another hyper-parameter, m , fixed at 0 to eliminate its influence. As shown in Figure 3 (left column), with the increase of parameter η , the performance of the Clean-LaVe method increases first and then decreases. When η is too small, although D_{clean} has a low noise level, it only contains a few samples and classes, thus the model performance is barely satisfactory. When η is too large, it easily involves too many noisy samples, thus deteriorating performance.

Diversity Number m . We evaluate the effects of hyper-parameter m which controls the number of samples from diverse classes. The search range for m is $[100, 200, \dots, 1300]$, while we maintain the value of η fixed at the best value found during previous tuning. As shown in Figure 3 (right column), with the increase of m , the performance generally increases since when m is too large, it easily in-

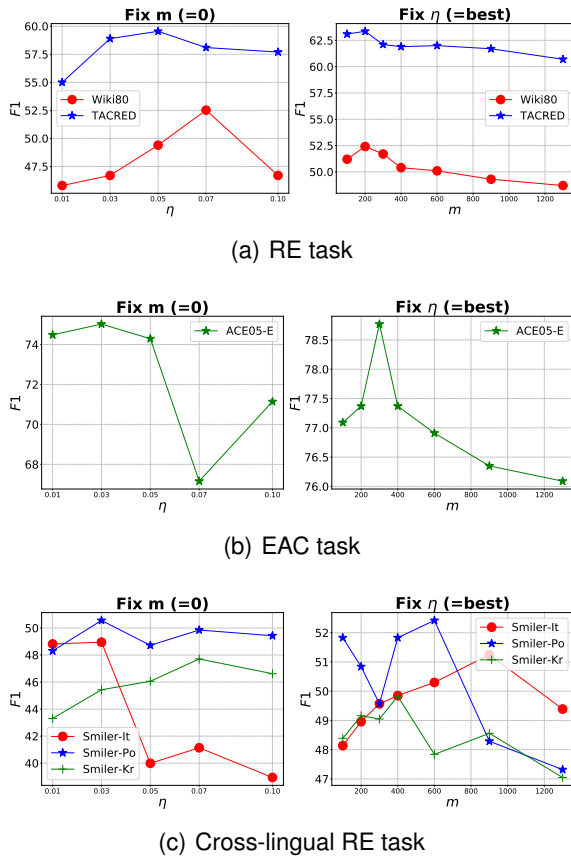


Figure 3: Results of different η and m .

volves too many noisy samples, thus deteriorating performance.

We perform manual quality checks on the filtered data using IWNL and CADs methods. As shown in Table 5, opting for larger values of η and m results in the inclusion of more classes and more noise. With the optimal hyper-parameter combination, we observe significant improvements in both data accuracy and class count. This enhancement underscores the effectiveness of our IWNL and CADs methods in improving the quality of filtered data.

Fix $m (=0)$					
η	0.01	0.05	0.1	0.5	1
Data Accuracy	92.10	86.02	80.16	68.51	64.03
Class Count	4	18	26	32	40
Fix η at best (=0.05)					
m	100	200	500	700	1000
Data Accuracy	79.48	76.40	75.15	74.70	74.20
Class Count	26	34	39	40	40

Table 5: The reliability of the data filtered by IWNL and CADs on TACRED. The results under the best hyper-parameter are marked in bold.

5. Conclusion

We propose a framework named Clean-LaVe to first detect a small amount of clean data from silver standard data and then use them to fine-tune the pre-trained model. We propose a Iteratively Weighted Negative Learning algorithm and Class-Aware Data Selector in clean data detection process to alleviate the imbalanced issue and to broaden the range of classes during selection. The experimental results demonstrate the effectiveness of our proposed method.

6. Future Work

- For further improvement, our work can be easily combined with LLM, using LLM as pretrained models.
- Clean-Lave has great potential to be extend to other tasks and settings which is worthwhile to be explored.

Acknowledgements

This research was supported by the Guangzhou Basic and Applied Basic Research Foundation (Grant No. 2023A04J1687).

7. Bibliographical References

- Open AI. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Görkem Algan and Ilkay Ulusoy. 2021. Image classification with deep learning in the presence of noisy labels: A survey. *Knowledge Based System*, 215:106771.
- Eric Arazo, Diego Ortego, Paul Albert, Noel O’Connor, and Kevin McGuinness. 2019. Un-supervised label noise modeling and loss correction. In *Proceedings of ICML*, pages 312–321.
- Devansh Arpit, Stanisław Jastrzębski, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, et al. 2017. A closer look at memorization in deep networks. In *Proceedings of ICML*, pages 233–242.
- Nontawat Charoenphakdee, Jongyeong Lee, and Masashi Sugiyama. 2019. On symmetric losses for learning from corrupted labels. In *Proceedings of ICML*, pages 961–970.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Benoît Frénay and Michel Verleysen. 2013. Classification in the presence of label noise: a survey. *IEEE transactions on neural networks and learning systems*, 25(5):845–869.
- Aritra Ghosh, Himanshu Kumar, and PS Sastry. 2017. Robust loss functions under label noise for deep neural networks. In *Proceedings of the AAAI*, pages 1919–1925.
- Ankur Goswami, Akshata Bhat, Hadar Ohana, and Theodoros Rekatsinas. 2020. Unsupervised relation extraction from language models using constrained cloze completion. In *Findings of EMNLP*, pages 1263–1276.
- Bo Han, Quanming Yao, Tongliang Liu, Gang Niu, Ivor W. Tsang, James T. Kwok, and Masashi Sugiyama. 2020. A survey of label-noise representation learning: Past, present and future. *arXiv preprint*.
- Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi Sugiyama. 2018. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Proceedings of NeurIPS*, pages 8536–8546.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. Deberta: Decoding-enhanced bert with disentangled attention. In *Proceedings of ICLR*.
- Jinchi Huang, Lie Qu, Rongfei Jia, and Binqiang Zhao. 2019. O2u-net: A simple noisy label detection approach for deep neural networks. In *Proceedings of ICCV*, pages 3326–3334.
- Kuan-Hao Huang, I Hsu, Premkumar Natarajan, Kai-Wei Chang, Nanyun Peng, et al. 2022a. Multilingual generative language models for zero-shot cross-lingual event argument extraction. In *Proceedings ACL*, pages 4633–4646.
- Lifu Huang, Heng Ji, Kyunghyun Cho, and Clare R Voss. 2018. Zero-shot transfer learning for event extraction. In *Proceedings of ACL*, pages 2160–2170.
- Yingsong Huang, Bing Bai, Shengwei Zhao, Kun Bai, and Fei Wang. 2022b. Uncertainty-aware learning against label noise on imbalanced datasets. In *Proceedings of AAAI*, volume 36, pages 6960–6969.
- Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. 2018. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *Proceedings of ICML*, pages 2304–2313.
- Lifeng Jin, Linfeng Song, Kun Xu, and Dong Yu. 2021. Instance-adaptive training with noise-robust losses against noisy labels. In *Proceedings of EMNLP*, pages 5647–5663.
- Youngdong Kim, Junho Yim, Juseung Yun, and Junmo Kim. 2019. Nlnl: Negative learning for noisy labels. In *Proceedings of ICCV*, pages 101–110.
- Moritz Laurer, Wouter van Atteveldt, Andreu Salleras Casas, and Kasper Welbers. 2022. [Less Annotating, More Classifying – Addressing the Data Scarcity Issue of Supervised Machine Learning with Deep Transfer Learning and BERT - NLI](#). *Preprint*. Publisher: Open Science Framework.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. 2017. Zero-shot relation extraction via reading comprehension. In *Proceedings of CoNLL*, pages 333–342.
- Junnan Li, Richard Socher, and Steven CH Hoi. 2020. Dividemix: Learning with noisy labels as semi-supervised learning. In *Proceedings of ICLR*.

- Zizheng Lin, Hongming Zhang, and Yangqiu Song. 2023. Global constraints with prompting for zero-shot event argument classification. In *Findings of EACL*, pages 2482–2493.
- Jian Liu, Yubo Chen, Kang Liu, Wei Bi, and Xiaojiang Liu. 2020. Event extraction as machine reading comprehension. In *Proceedings of EMNLP*, pages 1641–1651.
- Chenwei Lou, Jun Gao, Changlong Yu, Wei Wang, Huan Zhao, Weiwei Tu, and Ruifeng Xu. 2022. Translation-based implicit annotation projection for zero-shot cross-lingual event argument extraction. In *Proceedings of SIGIR*, pages 2076–2081.
- Di Lu, Shihao Ran, Joel Tetreault, and Alejandro Jaimes. 2023a. Event extraction as question generation and answering. In *Proceedings of ACL*, pages 1666–1688.
- Keming Lu, I Hsu, Wenxuan Zhou, Mingyu Derek Ma, Muhao Chen, et al. 2022. Summarization as indirect supervision for relation extraction. *arXiv preprint arXiv:2205.09837*.
- Yang Lu, Yiliang Zhang, Bo Han, Yiu-ming Cheung, and Hanzi Wang. 2023b. Label-noise learning with intrinsically long-tailed data. In *Proceedings of ICCV*, pages 1369–1378.
- Qing Lyu, Hongming Zhang, Elinor Sulem, and Dan Roth. 2021. Zero-shot event extraction via transfer learning: Challenges and insights. In *Proceedings of ACL*, pages 322–332.
- Shengfei Lyu and Huanhuan Chen. 2021. Relation classification with entity type restriction. In *Proceedings of ACL*, pages 390–395.
- Yueming Lyu and Ivor W Tsang. 2019. Curriculum loss: Robust learning and generalization against label corruption. In *Proceedings of ICLR*.
- Jun-Yu Ma, Beiduo Chen, Jia-Chen Gu, Zhen-Hua Ling, Wu Guo, Quan Liu, Zhigang Chen, and Cong Liu. 2022. Wider & closer: Mixture of short-channel distillers for zero-shot cross-lingual named entity recognition. *arXiv preprint arXiv:2212.03506*.
- Jun-Yu Ma, Jia-Chen Gu, Zhen-Hua Ling, Quan Liu, Cong Liu, and Guoping Hu. 2023. Shine: Syntax-augmented hierarchical interactive encoder for zero-shot cross-lingual information extraction. *arXiv preprint arXiv:2305.12389*.
- Xingjun Ma, Hanxun Huang, Yisen Wang, Simone Romano, Sarah M. Erfani, and James Bailey. 2020. Normalized loss functions for deep learning with noisy labels. In *Proceedings of ICML*, volume 119, pages 6543–6553.
- Eran Malach and Shai Shalev-Shwartz. 2017. Decoupling "when to update" from "how to update". In *Proceedings of NeurIPS*, pages 960–970.
- Sneha Mehta, Huzefa Rangwala, and Naren Ramakrishnan. 2022. Improving zero-shot event extraction via sentence simplification. In *Workshop of EMNLP*, pages 32–43.
- Aditya Krishna Menon, Ankit Singh Rawat, Sashank J. Reddi, and Sanjiv Kumar. 2020. Can gradient clipping mitigate label noise? In *Proceedings of ICLR*.
- Abiola Obamuyide and Andreas Vlachos. 2018. Zero-shot relation classification as textual entailment. In *Proceedings of the FEVER*, pages 72–78.
- Sinno Jialin Pan and Qiang Yang. 2009. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, pages 1345–1359.
- Mahdi Rahimi and Mihai Surdeanu. 2023. Improving zero-shot relation classification via automatically-acquired entailment templates. In *Proceedings of Workshop on Repl4NLP*, pages 187–195.
- Scott Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. 2015. Training deep neural networks on noisy labels with bootstrapping. In *Proceedings of ICLR Workshop*.
- Oscar Sainz, Oier Lopez de Lacalle, Gorra Labaka, Ander Barrena, and Eneko Agirre. 2021. Label verbalization and entailment for effective zero- and few-shot relation extraction. In *Proceedings of EMNLP*, pages 1199–1212.
- Oscar Sainz, Itziar Gonzalez-Dios, Oier Lopez de Lacalle, Bonan Min, and Eneko Agirre. 2022a. Textual entailment for event argument extraction: Zero- and few-shot with multi-source learning. In *Findings of NAACL-HLT*, pages 2439–2455.
- Oscar Sainz, Haoling Qiu, Oier Lopez de Lacalle, Eneko Agirre, and Bonan Min. 2022b. Zs4ie: A toolkit for zero-shot information extraction with simple verbalizations. In *Proceedings of NAACL-HLT*, pages 27–38.
- Yongliang Shen, Xiaobin Wang, Zeqi Tan, Guangwei Xu, Pengjun Xie, Fei Huang, Weiming Lu, and Yueting Zhuang. 2022. Parallel instance query network for named entity recognition. In *Proceedings of ACL*, pages 947–961.

- Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. 2019. Meta-weight-net: Learning an explicit mapping for sample weighting. In *Proceedings of NeurIPS*, pages 1917–1928.
- Sunil Thulasidasan, Tanmoy Bhattacharya, Jeff Bilmes, Gopinath Chennupati, and Jamal Mohd-Yusof. 2019. Combating label noise in deep learning using abstention. *arXiv preprint arXiv:1905.10964*.
- Thy Thy Tran, Phong Le, and Sophia Ananiadou. 2020. Revisiting unsupervised relation extraction. In *Proceedings of ACL*, pages 7498–7505.
- Thy Thy Tran, Phong Le, and Sophia Ananiadou. 2021. One-shot to weakly-supervised relation classification using language models. In *Proceedings of AKBC*.
- Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. 2019. Symmetric cross entropy for robust learning with noisy labels. In *Proceedings of ICCV*, pages 322–330.
- Qianhui Wu, Zijia Lin, Börje F Karlsson, Jian-Guang Lou, and Biqing Huang. 2020. Single-/multi-source cross-lingual ner via teacher-student learning on unlabeled data in target language. In *Proceedings of ACL*, pages 6505–6514.
- Sen Yang, Dawei Feng, Linbo Qiao, Zhigang Kan, and Dongsheng Li. 2019. Exploring pre-trained language models for event extraction and generation. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 5284–5294.
- Quanming Yao, Hansi Yang, Bo Han, Gang Niu, and James Kwok. 2019. Searching to exploit memorization effect in learning from corrupted labels. *arXiv preprint arXiv:1911.02377*.
- Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. 2019. How does disagreement help generalization against label corruption? In *Proceedings of ICML*, pages 7164–7173.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. 2017. Understanding deep learning requires rethinking generalization. In *Proceedings of ICLR*.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. 2021a. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115.
- Hongming Zhang, Haoyu Wang, and Dan Roth. 2021b. Zero-shot label-aware event trigger and argument classification. In *Findings of ACL-IJCNLP*, pages 1331–1340.
- Kai Zhang, Bernal Jimenez Gutierrez, and Yu Su. 2023. Aligning instruction tasks unlocks large language models as zero-shot relation extractors. In *ACL*, pages 794–812.
- Zhilu Zhang and Mert Sabuncu. 2018. Generalized cross entropy loss for training deep neural networks with noisy labels. In *Proceedings of NeurIPS*, pages 8792–8802.
- Jun Zhao, WenYu Zhan, Xin Zhao, Qi Zhang, Tao Gui, Zhongyu Wei, Junzhe Wang, Minlong Peng, and Mingming Sun. 2023. Re-matching: A fine-grained semantic matching method for zero-shot relation extraction. In *Proceedings of ACL*, pages 6680–6691.
- Zexuan Zhong and Danqi Chen. 2021. A frustratingly easy approach for entity and relation extraction. In *Proceedings of NAACL*, pages 50–61.
- Wenxuan Zhou and Muhao Chen. 2021. Learning from noisy labels for entity-centric information extraction. In *Proceedings of EMNLP*, pages 5381–5392.
- Enwei Zhu and Jinpeng Li. 2022. Boundary smoothing for named entity recognition. In *Proceedings of ACL*, pages 7096–7108.

8. Language Resource References

- Lin, Ying and Ji, Heng and Huang, Fei and Wu, Lingfei. 2020. *A joint neural model for information extraction with global features*. PID <https://aclanthology.org/2020.acl-main.713/>.
- Seganti, Alessandro and Firląg, Klaudia and Skowronska, Helena and Sattawa, Michał and Andruszkiewicz, Piotr. 2021. *Multilingual Entity and Relation Extraction Dataset and Model*. PID <https://aclanthology.org/2021.eacl-main.166>.
- Hongmin Xiao. 2022. *Wiki80*. PID <https://figshare.com/articles/dataset/Wiki80/19323371>.
- Victor Zhong and Yuhao Zhang and Danqi Chen and Gabor Angeli and Christopher Manning. 2018. *TAC Relation Extraction Dataset*. ISLRN 927-859-759-915-2. PID <https://catalog.ldc.upenn.edu/LDC2018T24>.