

NERO: A Neural Rule Grounding Framework for Label-Efficient Relation Extraction

Wenxuan Zhou¹, Hongtao Lin¹, Bill Yuchen Lin¹, Ziqi Wang²
Junyi Du¹, Leonardo Neves³, Xiang Ren¹

¹University of Southern California ²Tsinghua University ³Snapchat Inc.

{zhouwenx, lin498, yuchen.lin, junyidu, xiangren}@usc.edu ²ziqu-wan16@mails.tsinghua.edu.cn ³lneves@snap.com

ABSTRACT

Deep neural models for relation extraction tend to be less reliable when perfectly labeled data is limited, despite their success in label-sufficient scenarios. Instead of seeking more instance-level labels from human annotators, here we propose to annotate frequent surface patterns to form *labeling rules*. These rules can be automatically mined from large text corpora and generalized via a soft rule matching mechanism. Prior works use labeling rules in an exact matching fashion, which inherently limits the coverage of sentence matching and results in the low-recall issue. In this paper, we present a neural approach to ground rules for RE, named NERO, which jointly learns a relation extraction module and a soft matching module. One can employ any neural relation extraction models as the instantiation for the RE module. The soft matching module learns to match rules with *semantically similar* sentences such that raw corpora can be automatically labeled and leveraged by the RE module (in a much better coverage) as augmented supervision, in addition to the exactly matched sentences. Extensive experiments and analysis on two public and widely-used datasets demonstrate the effectiveness of the proposed NERO framework, comparing with both rule-based and semi-supervised methods. Through user studies, we find that the time efficiency for a human to annotate rules and sentences are similar (0.30 vs. 0.35 min per label). In particular, NERO’s performance using 270 rules is comparable to the models trained using 3,000 labeled sentences, yielding a 9.5x speedup. Moreover, NERO can predict for unseen relations at test time and provide interpretable predictions. We release our code¹ to the community for future research.

ACM Reference Format:

Wenxuan Zhou, Hongtao Lin, Bill Yuchen Lin, Ziqi Wang, Junyi Du, Leonardo Neves, and Xiang Ren. 2020. NERO: A Neural Rule Grounding Framework for Label-Efficient Relation Extraction. In *Proceedings of The Web Conference 2020 (WWW ’20)*, April 20–24, 2020, Taipei, Taiwan. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3366423.3380282>

1 INTRODUCTION

Relation extraction (RE) plays a key role in information extraction tasks and knowledge base construction, which aims to identify

¹<https://github.com/INK-USC/NERO>

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW ’20, April 20–24, 2020, Taipei, Taiwan

© 2020 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-7023-3/20/04.

<https://doi.org/10.1145/3366423.3380282>

Rules

p_1 (SUBJ-PER, 's child, OBJ-PER) → PER:CHILDREN
 p_2 (SUBJ-PER, is known as, OBJ-PER) → PER:ALTERNATIVE NAMES
 p_3 (SUBJ-ORG, was founded by, OBJ-PER) → ORG:FOUNDED_BY

Corpus

		Matching Score
s_1	Microsoft was founded by Bill Gates in 1975.	1.0
s_2	Microsoft was created by Bill Gates in 1975.	0.9
s_3	In 1975, Bill Gates launched Microsoft.	0.8

Hard Matching

Figure 1: Current rule-based methods mostly rely on exact/hard matching to raw corpus and suffer from limited coverage. For example, the rule body of p_3 only matches sentence s_1 but is also similar to s_2 and s_3 , which express the same relation as p_3 . A “soft rule matching” mechanism is desirable to make better use of the corpus for label generation.

the relation between two entities in a given sentence. For example, given the sentence “Bill Gates is the founder of Microsoft” and an entity pair (“Bill Gates”, “Microsoft”), a relation classifier is supposed to predict the relation of ORG:FOUNDED_BY. Recent advance in neural language processing has shown that neural models [36–38] gained great success on this task, yielding state-of-the-art performance when a large amount of well-annotated sentences are available. However, these supervised learning methods degrade dramatically when the sentence-level labels are insufficient. This problem is partially solved by the use of distant supervision based on knowledge bases (KBs) [20, 30]. They usually utilize an existing KB for automatically annotating an unlabeled corpus with an oversimplified assumption — two entities co-occurring in a sentence should be labeled as their relations in the KB regardless of their contexts. While distant supervision automates the labeling process, it also introduces noisy labels due to context-agnostic labeling, which can hardly be eliminated.

In addition to KBs, labeling rules are also important means of representing domain knowledge [25], which can be automatically mined from large corpora [14, 21]. Labeling rules can be seen as a set of pattern-based heuristics for matching a sentence to a relation. These rules are much more accurate than using KBs as distant supervision. The traditional way of using such rules is to perform exact string matching, and a sentence is either able or unable to be matched by a rule. This kind of hard-matching method inherently limits the generalization of the rules for sentences with similar semantics but dissimilar words, which consequently causes the low-recall problem and data-insufficiency for training neural models.

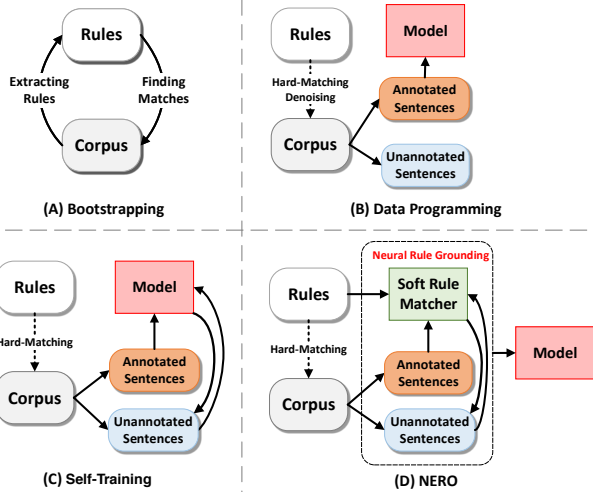


Figure 2: Comparison between previous work and the proposed NERO framework. (A) Bootstrapping. (B) Data Programming. (C) Self-Training. (D) NERO. The neural rule grounding process enables the soft-matching between rules and unmatched sentences using the soft rule matcher.

For example, in Fig. 1, the rule P_3 can only find S_1 as one hard-matched instance for training the model while other sentences (e.g., S_2 and S_3) should also be matched for they have a similar meaning.

Many attempts have been made to RE using labeling rules and unlabeled corpora, as summarized in Fig. 2. *Rule-based bootstrapping* methods [2, 4, 15] extract relation instances from the raw corpus by a pre-defined rule mining function (e.g., TF-IDF, CBOW) and expanding the rule set in an iterative way to increase the coverage. However, they still make predictions by performing hard-matching on rules, which is context-agnostic and suffers from the low-recall problem. Also, their matching function is not learnable and thus has trouble in extracting semantically similar sentences. Another typical option is *data programming* [9, 25], which aims to annotate the corpus using rules with model fitting. It trains a neural RE model using the hard-matched sentences by rules and then reduces the noise of rules with their proposed algorithms. However, it does not consider the massive data that fail to be annotated by hard-matching. *Self-training* [26], as a semi-supervised framework, does attempt to utilize unmatched sentences by using confident predictions of a learned model. Then, they train the model again and again by iteratively generating more confident predictions. However, it does not explicitly model the soft matching ability of rules over unmatched sentences, making the generated labels noisy and unreliable.

In this paper, we want to explicitly exploit labeling rules over unmatched sentences as supervision for training better RE models, for which we propose a NEural Rule grOunding (NERO) framework. The NERO framework has two major modules: a sentence-level relation classifier and a soft rule matcher. The former aims to learn the neural representations of sentences and classify which relation it talks about, which serves as the outcome of NERO. We first apply our collected rules on a raw corpus by *hard matching* and use hard-matched sentences as the main training data for NERO. The unmatched sentences are assigned with pseudo labels by the soft rule

matcher, which is a *learnable* module that produces matching scores for unmatched sentences with collected rules. The key intuition behind our soft rule matcher is that the distances between rules and sentences can be modeled by simple cosine computations in a new space transformed from their neural representations (e.g. word embeddings), which can be learned by a contrastive loss. Jointly training the two modules reinforce the quality of pseudo-labels and then improve the performance of the relation classification module. Specifically, in a batch-mode training process, we use the up-to-date learned soft matcher to assign all unmatched sentences with weighted labels, which further serve as an auxiliary learning objective for the relation classifier module.

Extensive experiments on two public datasets, TACRED [38] and Semeval [11], demonstrate that the proposed NERO consistently outperforms baseline methods using the same resources by a large margin. To investigate the label efficiency of NERO, we further conduct a user study to compare models trained with rules and labeled sentences, respectively, both created under the same time constraint. Results show that to achieve similar performance, NERO requires about ten times fewer human efforts in annotation time than traditional methods of gathering labels.

The **contributions** of our work are summarized as follows: (1) Our proposed method, NERO, is among the first methods that can learn to generalize labeling rules for training an effective relation extraction model, without relying on additional sentence-level labeled data; (2) We propose a learnable rule matcher that can semantically ground rules for sentences with similar meaning that are not able to be hard-matched. This method also provides interpretable predictions and supports few-shot learning ability on unseen relations; (3) We conduct extensive experiments to demonstrate the effectiveness of our framework in a low-resource setting and show careful investigations on the efficiency of using human efforts.

2 PROBLEM FORMULATION

We first introduce basic concepts and their notations and then present the problem definition as well as the scope of the work.

Relation Extraction. Given a pair of entity strings ($e_{\text{subj}}, e_{\text{obj}}$), identified as the subject and object entities respectively, the task of relation extraction (RE) is to predict (classify) the relation $r \in \mathcal{R} \cup \{\text{NONE}\}$ between them, where \mathcal{R} is a pre-defined set of relations of interest. Specifically, here we focus on *sentence-level* RE [11, 18, 38], which aims to predict the relation of entity mention pairs in a sentence s —i.e., identifying the label $r \in \mathcal{R}$ for an *instance* in the form of $(e_{\text{subj}}, e_{\text{obj}}; s)^2$ and yielding the prediction $(e_{\text{subj}}, r, e_{\text{obj}}; s)$.

Labeling Rules. As an alternative to instance-level human-annotated data (e.g., instance-label pairs), labeling rules formalize human’s domain knowledge in a structured way, and can be either directly used for making predictions over test instances [2], or applied to generate labeled instances for model learning [25]. Formally, an inductive labeling rule consists of a *rule body* and a *rule head*. In the context of relation extraction, we consider the cases where the rule body is a textual pattern $p = [\text{SUBJ-TYPE}; c; \text{OBJ-TYPE}]$, in which c denotes a word sequence between the two entities in a sentence (i.e., context), and SUBJ-TYPE and OBJ-TYPE specify the entity types of the subject and object entities required by the rule body. An

²We use “instance” and “sentence” exchangeably in the rest of the paper.

instance $(e_{\text{subj}}, e_{\text{obj}}; s)$ is *hard-matched* by a rule p if and only if p 's context c exactly match the context between e_{subj} and e_{obj} in s and the entity types of e_{subj} and e_{obj} match p 's associated entity types. For example, sentence s_1 in Fig. 1 is hard-matched by rule p_3 .

Such labeling rules can either be hand-written by domain experts, or automatically generated from text corpora with a knowledge base that shares the relations of interests. In this work, we adopt a hybrid approach (as detailed in Sec. 3.2): 1) we first extract surface patterns (e.g. frequent patterns, relational phrases, etc.) as candidate rule bodies from raw corpora (after running named entity recognition tool) using pattern mining tools, and 2) ask human annotators to assign relation labels (i.e., rule heads) to the rule bodies. Then, we apply these human-annotated rules over instances from a raw corpus to generate labeled instances (as training data). Compared to directly annotating instances, our study will show that annotating rules is much more label-efficient, as one rule may match many instances in the corpus with reliable confidence, generating multiple labeled instances. Moreover, annotating rules may require similar (or even less) time than annotating instances since they are shorter in length, as shown in our later analysis.

EXAMPLE 1 (LABELING RULE GENERATION). *Given instances such as {(Bill Gates, Microsoft; "Bill Gates founded Microsoft."), (Bill Gates, Steve Jobs; "Guests include Bill Gates and Steve Jobs.")}, the candidate patterns include [{"SUBJ-PERSON founded OBJ-PERSON"}, {"SUBJ-PERSON, OBJ-PERSON"}], and the rule set after human annotation includes [{"SUBJ-PERSON founded OBJ-PERSON"} \rightarrow FOUNDED_BY]. The second pattern is filtered out by annotators due to its low quality.*

Problem Definition. In this paper, we focus on sentence-level relation extraction using labeling rules created by a semi-automatic method. Specifically, given a raw corpus consisting of a set of instances $S = \{(e_{\text{subj}}^i, e_{\text{obj}}^i; s^i)\}_{i=1}^N$, our approach first extracts surface patterns $\{[\text{SUBJ-TYPE}; c; \text{OBJ-TYPE}]\}$ as candidate rule bodies, and then annotates them to generate labeling rules $\mathcal{P} = \{p^k\}_{k=1}^{|\mathcal{P}|}$ with p^k denoting $[\text{SUBJ-TYPE}^k; c^k; \text{OBJ-TYPE}^k] \rightarrow r_k$ and $r_k \in \mathcal{R} \cup \{\text{NONE}\}$. Next, our framework aims to leverage both the rules \mathcal{P} and the corpus S to learn an effective relation classifier $f: S \rightarrow \mathcal{R} \cup \{\text{NONE}\}$ which can predict relation for new instances in unseen sentences.

In addition to predicting relations already observed in the training phase, we aim to **predict unseen relations** solely based on new rules specified for the unseen relations at the test time—i.e., a few-shot learning scenario. Such a generalization ability is desirable for domain adaptation. In this setting, the model is trained on a corpus S using labeling rules defined on relation set \mathcal{R}_c . In the testing phase, the model is expected to make predictions on unseen relation types \mathcal{R}_u given some corresponding labeling rules.

Our Focus. Instead of relying on pre-existing knowledge bases to conduct distant supervision, we focus on the scenarios where labeling rules are relatively cheaper and faster to obtain by automatically inducing from large text corpora (as compared to manually curating a KB from scratch). We study neural rule grounding to integrate information from rules and corpora into learning an effective relation extraction model that can generalize better than solely using rules in a hard-matching manner. We mainly use surface pattern rules (e.g., similar to the ones used in [27]) and leave the study of more complex rules (e.g., regular expression) as future work.

3 NEURAL RULE GROUNDING (NERO)

This section introduces the important concepts in NERO, and provides details of the framework design and learning objectives. We first present our key motivation, then give an overview of NERO framework (Sec. 3.1) and introduce the details of each component (Secs. 3.2-3.4). Lastly, we show how we jointly learn model parameters in NERO for sequence encoding and rule matching (Sec. 3.5).

Motivation of Soft Rule Matching. Traditional methods of using labeling rules for relation extraction are mainly based on exact string/pattern matching (i.e., *hard-matching*), where a sentence can be annotated by a rule if and only if it has exactly the same surface form as the rule body. While being easy to implement and yielding relatively high precision for the prediction, these hard-matching methods only achieve sub-optimal performance because of the severe low recall issue (see Fig. 1 for an example). Such a low coverage of using rules can only produce a minimal number of "labeled" sentences for learning a RE model, which may be far from enough for data-hungry neural models. To address this problem, we would like to study "soft-matching" for exploiting the rules for relation extraction. A soft-rule matcher is proposed to semantically match a rule even when the surfaces are different (e.g. "founded" v.s. "created"). Towards improving relation extraction with soft-matching, we would like to learn a SRM such that we can assign *pseudo-labels* to sentences. Further, we can use these pseudo-labeled sentences to learn a sentence-level relation classifier.

3.1 Framework Overview

The NERO framework consists of a relation classifier (RC) and a soft rule matcher (SRM). The RC is for learning neural representations for sequences of words, from either an instance or a rule body (Sec. 3.3). The SRM gives a matching score to indicate the similarity between a rule body and an instance (Sec. 3.4). Previous works [4] mostly adopt a *fixed* metric function (e.g., cosine similarity measure between the sequence embeddings). In contrast, our SRM is a *learnable* neural network and thus can be more effective in modeling the semantic similarity between sentences and rule bodies. As shown in Fig. 3, we first apply our collected rules to a raw corpus S to perform *hard matching*. S will then be divided into two subsets: "*hard-matched sentences*" S_{matched} and "*unmatched sentences*" $S_{\text{unmatched}}$. Lastly, we use SRM and rules to iteratively generate pseudo labels over $S_{\text{unmatched}}$, while jointly learning RC and SRM with the pseudo labels to promote a common representation space (Sec. 3.5).

3.2 Labeling Rule Generation

As introduced in Sec. 2, our candidate rule bodies (i.e., surface patterns) are automatically extracted from the raw corpus. In this work, we adopt a simple yet effective pattern mining method. Given a raw corpus, we first replace entities with entity type masks SUBJ/OBJ-NER, where "NER" denotes entity type (same as the procedure in [38]). Then we pick the word sequences between and including the two entities as candidate rules. So each candidate rule is just a short phrase / sentence containing two masked entities. To reduce the annotation efforts and ensure the popularity of rules, we convert all words to their root form using Porter stemming algorithm [33] and only keep the rules whose stemmed form appears at least N times in the whole corpus. Finally, we ask human annotators to

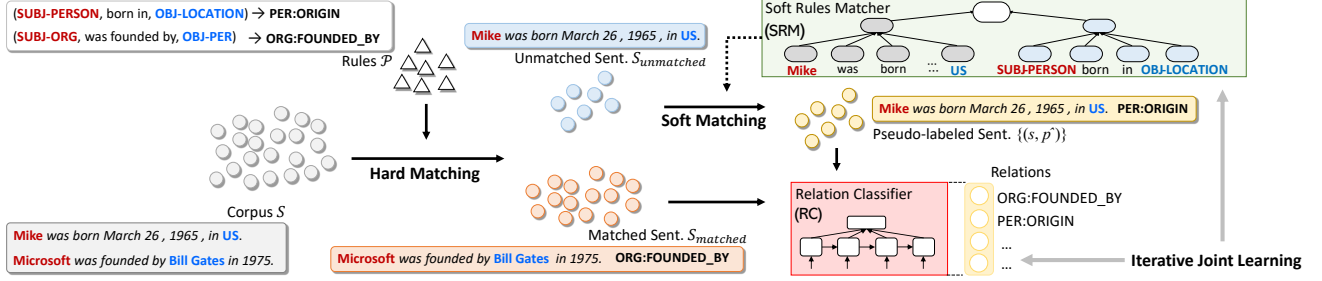


Figure 3: Overview of the NERO framework. Each unmatched sentence is first annotated by the soft rule matcher (SRM) to generate pseudo labels, and then fed into the relation classifier (RC) to update the model parameters. The whole framework is trained iteratively and jointly, based on multiple loss functions as introduced in Sec. 3.5.

select rules that indicate a relation and assign labels to them. This process is similar to label a sentence with a pre-defined relation or NONE. In spite of the existence of more complex rule extraction methods such as shortest dependency path [21, 24] and meta patterns [14], we adopt frequent word sequences in this work because of their simplicity and high readability. Our framework can also be adapted to other rule generation methods as long as they can generate human-readable candidate rules. Regardless of their high generalization ability, labeling rules may bias the label distribution, e.g. PER:TITLE counts for 19% in the official training data while only counts for 5.5% in the matched sentences. This domain gap poses another challenge for learning with rules.

3.3 Relation Classifier (RC)

The relation classifier (named as RC) is the end product of our NERO framework, which can represent a relation instance (e_1, e_2, s) into vector embeddings. This module is general and can use various designs (e.g., attention CNN [32], attention RNN [38, 39], and GCN [7, 37]). In this work, we stick with the LSTM + ATT model [39]. Given a sequence of n words, we first look up their word embeddings as $\{\mathbf{x}_t\}_{t=1}^n$. Then we use a bi-directional LSTM network [12] to obtain the contextualized embeddings of each word $\{\mathbf{h}_t\}_{t=1}^n$ where $\mathbf{h}_t \in \mathbb{R}^{d_h}$. Finally, we employ an attention layer [3] to get a sentence representation \mathbf{c} :

$$\{\mathbf{h}_t\}_{t=1}^n = \text{BiLSTM}(\{\mathbf{x}_t\}_{t=1}^n) \quad (1)$$

$$\alpha_t = \frac{\exp(\mathbf{v}^T \tanh(\mathbf{A}\mathbf{h}_t))}{\sum_{t'=1}^n \exp(\mathbf{v}^T \tanh(\mathbf{A}\mathbf{h}_{t'}))} \quad (2)$$

$$\mathbf{c} = \sum_{t=1}^n \alpha_t \mathbf{h}_t \quad (3)$$

where $\mathbf{A} \in \mathbb{R}^{d_a \times d_h}$, $\mathbf{v} \in \mathbb{R}^{d_a}$ are learnable model parameters for the attention mechanism. We can get the relation type probability distribution by feeding $\mathbf{W}_{\text{rc}}\mathbf{c}$ into a SoftMax layer:

$$\text{RC}(s, e_{\text{subj}}, e_{\text{obj}}) = \text{SoftMax}(\mathbf{W}_{\text{rc}}\mathbf{c}),$$

where $\mathbf{W}_{\text{rc}} \in \mathbb{R}^{d_h \times |\mathcal{R}|}$ is a matrix to learn. In other words, we are using the output vector of the RC for modeling the conditional distribution of relation $\mathbb{P}_{\theta_{\text{RC}}}(r = i | s) = \text{RC}(s, e_{\text{subj}}, e_{\text{obj}})[i]$, meaning that i -th element of the RC output vector as the probability of the i -th relation for the input sentence s . The θ_{rc} denotes the set of model parameters of RC that we need to learn, including \mathbf{A} , \mathbf{v} , \mathbf{W}_{rc} , and the weights of the BiLSTM.

3.4 Soft Rule Matcher (SRM)

The core component of the NERO is the soft rule matcher (named to be SRM), which is defined as a neural function for modeling the matching score between a sentence s and a rule pattern p . We formulate the Soft Rule Matcher as a function $\text{SRM} : (s, p) \rightarrow [-1, 1]$. As shown in Fig. 4, we first map the sentence s and rule p into the same embedding space by the RC, and then apply a distance metric on their embeddings to get the matching score ranging from -1 to 1 . Note that a matching score of 1 indicates the sentence can be hard-matched by the rule. Because the rules are typically short phrases, we use word-level attention – a weighted continuous bag-of-words model to calculate matching scores. We explored more designs in later experiments and this part can be further studied in future work. We represent a sentence and a rule pattern as the sequence of the word embeddings, $\{\mathbf{x}_t^s\}_{t=1}^n$ and $\{\mathbf{x}_t^p\}_{t=1}^m$, respectively. By applying the attention on word embeddings instead of contextualized embedding generated by LSTM, we aim to reduce over-fitting given the scarcity of rules per relation. Note that we are using a different set of parameters for this word embedding-level attention (\mathbf{B} and \mathbf{v}). Specifically, we have

$$\mathbf{z}_s = \sum_{t=1}^n \frac{\exp(\mathbf{u}^T \tanh(\mathbf{B}\mathbf{x}_t^s))}{\sum_{t'=1}^n \exp(\mathbf{u}^T \tanh(\mathbf{B}\mathbf{x}_{t'}^s))} \mathbf{x}_t^s, \quad (4)$$

$$\mathbf{z}_p = \sum_{t=1}^m \frac{\exp(\mathbf{u}^T \tanh(\mathbf{B}\mathbf{x}_t^p))}{\sum_{t'=1}^m \exp(\mathbf{u}^T \tanh(\mathbf{B}\mathbf{x}_{t'}^p))} \mathbf{x}_t^p, \quad (5)$$

$$\text{SRM}(s, p) = \frac{(\mathbf{D}\mathbf{z}_s)^T (\mathbf{D}\mathbf{z}_p)}{\|\mathbf{D}\mathbf{z}_s\| \|\mathbf{D}\mathbf{z}_p\|}, \quad (6)$$

where \mathbf{z}_s , \mathbf{z}_p are the representations for the sentence s and the rule p respectively, and \mathbf{D} is a trainable diagonal matrix which denotes the importance of each dimension. Similar idea has been shown effective in other tasks for capturing semantically close phrases, and it generalizes well when only limited training data is available [35]. In practice, a sentence may be very long and contain much irrelevant information. So for the sentence, we only keep the words between (and including) the subject and object entities, which is a common practice in other rule-based RE systems [4, 24].

3.5 Joint Module Learning

We now formulate the overall learning process of NERO in a low-resource learning scenario, where we only have collected rules but no sentence-level labels at all. Given a raw corpus S (i.e. a

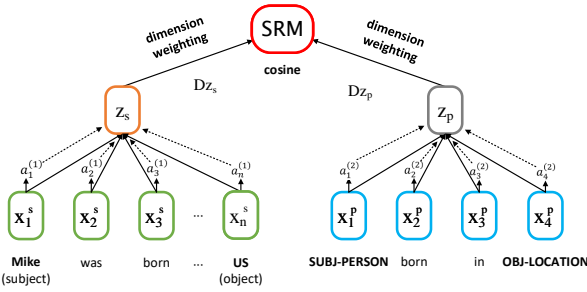


Figure 4: Detailed architecture of the soft rule matcher (SRM). The cosine similarity between two embeddings indicates the degree of matching between rules and sentences.

large set of sentences without human-annotated labels) and a set of inductive labeling rules \mathcal{P} , we first exactly match every rule on \mathcal{S} and get a set of **hard-matched sentences** named $\mathcal{S}_{\text{matched}}$. These hard-matched labels are pretty accurate³, which can serve the purpose as training examples. The other **unmatched sentences**, denoted as $\mathcal{S}_{\text{unmatched}}$ are also informative for our model, while they are ignored by most prior works. Although not able to be exactly matched by any rules, the unmatched sentences can semantically align the meaning of some rules and thus help improve the RE performance. By applying the SRM upon them to create pseudo-labels, we can further utilize the hidden useful information as the supervision for learning the final RE model. Apart from that, the SRM is also optimized for improving the quality of matching in the meantime. Finally, supervision signals from $\mathcal{S}_{\text{matched}}$, $\mathcal{S}_{\text{unmatched}}$, and \mathcal{P} together train the RE model in a joint learning schema.

Learning with Hard-Matched Sentences ($\mathcal{S}_{\text{matched}}$). As the primary goal of the learning, we want to minimize the error of the RC in classifying relations in the sentences in $\mathcal{S}_{\text{matched}}$ respect to their matched relation labels. Given a hard-matched sentence $s \in \mathcal{S}_{\text{matched}}$ and its associated relation $r_s \in \mathcal{R}$, we aim to minimize the cross-entropy loss L_{matched} as follows.

$$L_{\text{matched}}(\theta_{RC}) = \mathbb{E}_{s \sim \mathcal{S}_{\text{matched}}} [-\log \mathbb{P}_{\theta_{RC}}(r = r_s | s)]. \quad (7)$$

Learning with Rules (\mathcal{P}). To incorporate our collected rules for exploiting the unmatched sentences $\mathcal{S}_{\text{unmatched}}$ as well, we propose two additional auxiliary tasks for imposing rules in the RC and SRM. First, we can treat the rule body $p \in \mathcal{P}$ as a ‘‘sentence’’ and its rule head r_p as its associated label, thus forming a labeled instance (p, r_p) . With such rule-reformatted instances as training data, we aim to minimize the error for RC in classifying them as follows. This objective helps RC to explicitly reinforce the memory of RC about the collected rules.

$$L_{\text{rules}}(\theta_{RC}) = \mathbb{E}_{p \sim \mathcal{P}} [-\log \mathbb{P}_{\theta_{RC}}(r = r_p | p)] \quad (8)$$

More importantly, we present how we utilize the rules for learning the soft rule matcher (SRM) such that we can exploit the unmatched sentences as distant supervision. Our key motivation here is that a good matcher should be able to clustering rules with the same type, and thus we expect the SRM to increase the distances between rules with different types and reduce the variance of the distances between rules with the same types. Simply put, given two

³We conduct an experimental analysis on the quality of these matches in Sec. 5.




Figure 5: The contrastive loss in learning the SRM, which increases the matching score of rules with the same relation type and decrease the matching score otherwise.

rules, their similarity score should be high if they belong to the same relation type and low otherwise.

We use the contrastive loss [22] for this objective. Given a rule $p \in \mathcal{P}$ and its relation $r_p \in \mathcal{R}$, we divide all the other rules with the same relation type as $\mathcal{P}_+(p) = \{p' | r_{p'} = r_p\}$, and the ones with different types as $\mathcal{P}_-(p) = \{p' | r_{p'} \neq r_p\}$. The contrastive loss aims to pull the rule p closer to its most dissimilar rule in $\mathcal{P}_+(p)$ and in the meantime push it away from its most similar rule in $\mathcal{P}_-(p)$. Mathematically, the loss is defined as follows:

$$L_{\text{clus}} = \mathbb{E}_{p \sim \mathcal{P}} \left[\max_{p_i \in \mathcal{P}_+(p)} \text{dist}_+(p, p_i) - \min_{p_j \in \mathcal{P}_-(p)} \text{dist}_-(p, p_j) \right], \quad (9)$$

where the measurement of the most dissimilar same-type rule and most similar different-type rule can be measured with the distances defined as below:

$$\begin{aligned} \text{dist}_+(p, p_i) &= \max(\tau - \text{SRM}(p, p_i), 0)^2, \\ \text{dist}_-(p, p_j) &= 1 - \max(\text{SRM}(p, p_j), 0)^2. \end{aligned}$$

τ is a hyper-parameter for avoiding collapse of the rules’ representations. Minimizing L_{clus} pushes the matching scores of rules that are of the same relation type up to τ and pull the matching score down to 0 otherwise. Without labels about the alignment between rules \mathcal{P} and unmatched sentences $\mathcal{S}_{\text{unmatched}}$, this objective can be seen as a secondary supervision for SRM to train the parameters θ_{SRM} using \mathcal{P} itself only by this contrastive loss.

Learning with Unmatched Sentences ($\mathcal{S}_{\text{unmatched}}$). We use the SRM to label unmatched sentences $\mathcal{S}_{\text{unmatched}}$ (see Fig. 3). As the labels generated by the SRM may be noisy, we propose to use pseudo-labeling [16] with instance weighting [13] to alleviate the noise. Each unmatched sentence is weighed by the matching score from the SRM. Specifically, for each unmatched sentence $s \in \mathcal{S}_{\text{unmatched}}$, we first apply the SRM to compute its matching scores with each rule $p \in \mathcal{P}$ and then assign the pseudo-label to the sentence s as the relation of the highest-score rule $r_{\hat{p}}$, where

$$\hat{p} = \arg \max_{p \in \mathcal{P}} \text{SRM}(s, p).$$

The corresponding weight for a particular pseudo-labeled instance, (s_i, \hat{p}_i) , is the SoftMax over all the sentences (in a mini-batch when training; see Sec. 4.1 for more details):

$$w_s = \frac{\exp(\sigma \text{SRM}(s, \hat{p}_i))}{\sum_{s' \in \mathcal{S}_{\text{unmatched}}} \exp(\sigma \text{SRM}(s', \hat{p}_j))}, \quad (10)$$

where σ represents the temperature of the SoftMax function. The loss function for each batch of unmatched sentences would be:

$$L_{\text{unmatched}}(\theta_{RC}) = \mathbb{E}_{s \sim \mathcal{S}_{\text{unmatched}}} [-w_s \log \mathbb{P}_{\theta_{RC}}(r = r_{\hat{p}} | s)]. \quad (11)$$

Algorithm 1: Optimization of NERO model**Input:** A raw corpus \mathcal{S} , pre-defined relations $\mathcal{R} \cup \{\text{NONE}\}$.**Output:** A relation classifier $f: \mathcal{S} \rightarrow \mathcal{R} \cup \{\text{NONE}\}$.Extract candidate rules from \mathcal{S} with pattern mining tools.Ask human annotators to select and label the candidate rules to get \mathcal{P} .Partition \mathcal{S} into $\mathcal{S}_{\text{matched}}$ and $\mathcal{S}_{\text{unmatched}}$ by hard-matching with \mathcal{P} .**while** L in Eq. 3.5 not converge **do** Sample batch $\mathcal{B}_m = \{(s_i, r_i)\}_{i=1}^n$ from $\mathcal{S}_{\text{matched}}$. Update L_{matched} by Eq. 7. Sample batch $\mathcal{B}_u = \{s_j\}_{j=1}^m$ from $\mathcal{S}_{\text{unmatched}}$. **foreach** $s \in \mathcal{B}_u$ **do** Find highest-scored rule \hat{p} and pseudo label $r_{\hat{p}}$ by SRM. Update $L_{\text{unmatched}}$ by Eq. 12. Update L_{rules} by Eq. 8. **foreach** $p \in \mathcal{P}$ **do** Calculate SRM(p, p') for each $p' \in \mathcal{P} - \{p\}$. Update L_{clus} . $L = L_{\text{matched}} + \alpha \cdot L_{\text{rules}} + \beta \cdot L_{\text{clus}} + \gamma \cdot L_{\text{unmatched}}$. Update model parameters w.r.t. L .

Joint Optimization Objective. The whole framework is jointly trained under the overall loss function:

$$L(\theta_{\text{RC}}, \theta_{\text{SRM}}) = L_{\text{matched}}(\theta_{\text{RC}}) + \alpha \cdot L_{\text{rules}}(\theta_{\text{RC}}) + \beta \cdot L_{\text{clus}}(\theta_{\text{SRM}}) + \gamma \cdot L_{\text{unmatched}}(\theta_{\text{RC}}).$$

To sum up, the proposed NERO framework has two trainable components: a relation classifier (RC) and a soft rule matcher (SRM) with their specific parameters (θ_{RC} and θ_{SRM}). Our primary task is to minimize the error of the RC on the hard-matched sentences (L_{matched}). To exploit the unmatched sentences with collected rules, we first let RC to explicitly learn the rules with the objective of rule classification (L_{rules}) and also learn the SRM with the help of contrastive loss (L_{clus}) for clustering rules. We jointly learn the two modules by connecting them through pseudo-labeling on the unmatched sentences with the SRM and then expect the RC has better performance on such pseudo-labeled unmatched sentences as well ($L_{\text{unmatched}}$).

4 MODEL LEARNING AND INFERENCE

In this section, we introduce the specific training details of our proposed framework, and how the model conduct inference.

4.1 Parameter Learning of NERO

NERO starts with a raw corpus \mathcal{S} and pre-defined relations $\mathcal{R} \cup \{\text{NONE}\}$. It first extracts candidate rules from \mathcal{S} with pattern mining tools, then asks human annotators to select and label the rules to get the rule set \mathcal{P} . Before training, NERO does a quick pass of \mathcal{S} by hard-matching to split it into hard-matched sentences $\mathcal{S}_{\text{matched}}$ and unmatched sentences $\mathcal{S}_{\text{unmatched}}$. Both of them are utilized to train the relation classifier RC in a joint-training manner. Algorithm 1 summarizes the overall training procedure.

The joint training objective is efficiently optimized by batch-mode training. Specifically, in each training step we randomly sample a batch of hard-matched sentences \mathcal{B}_m from $\mathcal{S}_{\text{matched}}$, and then calculate the average cross-entropy loss in the batch by L_{matched} . Similarly, we sample \mathcal{B}_u from $\mathcal{S}_{\text{unmatched}}$, and generate the pseudo-labels with the entire rule set \mathcal{P} , and calculate normalized weights

Dataset	# Train / Dev / Test	# Relations	# Rules	# matched Sent.
TACRED [38]	75,049 / 25,763 / 18,659	42	270	1,630
SemEval [11]	7,199 / 800 / 1,864	19	164	1,454

Table 1: Statistics for TACRED and SemEval datasets.

for $L_{\text{unmatched}}$ by:

$$w_s = \frac{\exp(\sigma \text{SRM}(s, \hat{p}_i))}{\sum_{s' \in \mathcal{B}_u} \exp(\sigma \text{SRM}(s', \hat{p}_j))}, \quad (12)$$

$$L_{\text{unmatched}}(\theta_{\text{RC}}) = \frac{1}{|\mathcal{B}_u|} \sum_{s \in \mathcal{B}_u} \left[-w_s \log \mathbb{P}_{\theta_{\text{RC}}}(r = r_{\hat{p}} | s) \right].$$

The normalized instance weights ensure the scale of $L_{\text{unmatched}}$ to be stable across different steps. Finally, we calculate L_{rules} and L_{clus} with the entire \mathcal{P} . For L_{rules} , the loss is averaged for all $p \in \mathcal{P}$. For L_{clus} , we take all $p \in \mathcal{P}$ to calculate the contrastive loss w.r.t. $\mathcal{P} - \{p\}$, and average the losses. We do not sample \mathcal{P} due to the relatively small number of rules and simple structure of the soft rule matcher (SRM). We stop training when the joint training loss L converges and take RC as the output. Note that word vectors are also trainable parameters that are initialized by pre-trained embeddings.

4.2 Model Inference

Model inference in NERO aims at predicting the relation of a new sentence. Intuitively, we can give out a prediction by passing the sentence through the RC and return the relation with the highest probability as the prediction. As an alternative, we can also predict the relation using the SRM. Given a sentence, we can first find the most similar rule and return its rule head as the prediction. This alternative way of inference can be applied for predicting *unseen relations* given new rules in the testing time. In experiments, the first method shows much better performance, since RC can capture rich contextual information while the SRM cannot.

For predicting “NONE” examples, our model filters out the predictions that our model is least certain⁴ about. Specifically, we measure the uncertainty using the entropy of softmax distribution or the similarity score produced by SRM.

5 EXPERIMENTS

In this section, we first introduce the datasets and compared baseline methods. Then, we illustrate the detailed setup and present extensive experiments with discussion and analysis. Finally, we conduct user studies to investigate the efficiency of our proposed approach.

5.1 Data Preparation

We choose two public and widely-used sentence-level relation extraction datasets in our experiments (see Tab. 1) as follows. For both datasets, we construct the rule set as illustrated in Section 3.2⁵.

- **TACRED** [38] (TAC relation extraction dataset) contains more than 100,000 sentences categorized into 42 relation types. Among the sentences, 79.5% of the examples are labeled as NONE. We construct 270 rules which 1,630 hard-matched sentences in the official training data;

⁴A threshold δ tuned on dev set is used.

⁵We filter candidate patterns by requiring frequency ≥ 3 .

- **SemEval 2010 Task 8** [11] contains about 10,000 sentences with 19 relation types, where 17.4% of the sentences are NONE. We construct 164 rules which hard-match 1,454 instances in the official training data.

5.2 Compared Methods

Recall Sec. 3.5, we first apply hard-matching on the unlabeled sentences and partition them into hard-matched sentences (S_{matched}) and unmatched sentences ($S_{\text{unmatched}}$). Our framework is applicable in both training with S_{matched} and training with $S_{\text{matched}} \cup S_{\text{unmatched}}$. Thus, we compare our models with *rule-based methods*, *supervised methods* and *semi-supervised methods*. For both semi-supervised and supervised methods, we use hard-matched sentences as the training data for comparing under the same setting. The experiments evaluate the ability to use rules for improving RE.

Rule-based Baseline Methods. We apply the following models on $S_{\text{matched}} \cup S_{\text{unmatched}}$: (1) CBOW-GloVe adopts continuous bag-of-words [19] on GloVe embeddings [23] to represent a sentence or rule body, which labels a sentence using its most similar rule (in cosine distance). (2) BREDS [4] is a rule-based bootstrapping method originally designed for corpus-level RE. Given some entity pairs as seeds, it alternates between extracting new rules and new entity pairs. BREDS performs soft matching between rules and sentences using average word embeddings. In our experiments, we apply the set of rules learned by BREDS to perform hard-matching over the test sentences at the prediction time. (3) The Neural Rule Engine (NRE [17]) is an unsupervised method of soft-matching. It generalizes given rules by first performing unigram matching for each token in the rules using CNN, then accumulates the matching scores of all tokens along the parsed tree structure. When used for prediction, NRE performs soft-matching on the given rules.

Supervised Baseline Methods. We apply the following methods on S_{matched} : (1) PCNN [36] represents each token using both word embeddings and positional embeddings. Convolution and piecewise max pooling layer are performed to produce the sentence embedding. (2) LSTM+ATT adopts bi-directional LSTM and attention mechanism [3] to produce a sentence embedding, which is fed into a fully-connected layer and a softmax classifier to predict the relation. (3) PA-LSTM [38] extends the LSTM model by incorporating positional information into attention mechanism and achieved state-of-the-art performance on TACRED. (4) Data Programming [9, 25] denoises the conflicting rules by learning their accuracy and correlation structures. (5) LSTM+ATT ($S_{\text{matched}} + \mathcal{P}$) extends the LSTM+ATT model by also using rules as training data. It serves as the base model for all semi-supervised baselines.

Semi-Supervised Baseline Methods. To make fair comparisons with NERO, we also regard labeling rules as training data (same to L_{rules}) and use LSTM+ATT ($S_{\text{matched}} + \mathcal{P}$) as the base model. We apply the following methods on $S_{\text{matched}} \cup S_{\text{unmatched}} \cup \mathcal{P}$: (1) Pseudo-Labeling [16] trains the network using labeled and unlabeled data simultaneously. Pseudo-Labels are created for unlabeled data by picking up the class with the maximum predicted probability and are used as if they are labeled data during training. (2) Self-Training [26] iteratively trains the model using the labeled dataset and expands the labeled set using the most confident predictions among the unlabeled set. This procedure stops when

Method / Dataset	TACRED			SemEval		
	Precision	Recall	F ₁	Precision	Recall	F ₁
Rules	85.0	11.4	20.1	81.2	17.2	28.5
BREDS [4]	53.8	20.3	29.5	62.0	24.5	35.1
CBOW-GloVe	27.9	45.7	34.6	44.0	52.8	48.0
NRE [17]	65.2	17.2	27.2	78.6	18.5	30.0
PCNN [36]	44.5 ± 0.4	24.1 ± 2.8	31.1 ± 2.6	59.1 ± 1.4	43.0 ± 0.7	49.8 ± 0.5
LSTM+ATT	38.1 ± 2.7	39.6 ± 2.7	38.8 ± 2.4	64.5 ± 2.8	53.3 ± 2.8	58.2 ± 0.8
PA-LSTM [38]	39.8 ± 2.5	40.2 ± 2.0	39.0 ± 0.6	64.0 ± 3.6	54.2 ± 2.5	58.5 ± 0.6
Data Programming [25]	39.2 ± 1.3	40.1 ± 2.0	39.7 ± 0.9	61.8 ± 2.1	54.8 ± 1.1	58.1 ± 0.7
LSTM+ATT ($S_{\text{matched}} + \mathcal{P}$)	39.2 ± 1.7	45.5 ± 1.7	42.1 ± 0.9	63.4 ± 2.1	55.0 ± 0.3	58.8 ± 0.9
Pseudo-Labeling [16]	34.5 ± 4.1	37.4 ± 5.1	35.3 ± 0.8	59.4 ± 3.3	55.8 ± 2.1	57.4 ± 1.3
Self-Training [26]	37.8 ± 3.5	41.1 ± 3.1	39.2 ± 2.1	62.3 ± 2.0	53.0 ± 2.7	57.1 ± 1.0
Mean-Teacher [31]	46.0 ± 2.7	41.6 ± 2.2	43.6 ± 1.3	62.3 ± 1.5	54.5 ± 1.2	57.9 ± 0.5
DualRE [18]	40.2 ± 1.5	42.8 ± 2.0	41.7 ± 0.5	63.7 ± 2.8	54.6 ± 2.1	58.6 ± 0.8
NERO w/o $S_{\text{unmatched}}$	41.9 ± 1.8	44.3 ± 3.8	42.9 ± 1.4	61.4 ± 2.4	56.2 ± 1.9	58.6 ± 0.6
NERO-SRM	45.6 ± 2.2	45.2 ± 1.2	45.3 ± 1.0	54.8 ± 1.6	55.2 ± 2.0	54.9 ± 0.6
NERO	54.0 ± 1.8	48.9 ± 2.2	51.3 ± 0.6	66.0 ± 1.5	55.8 ± 0.9	60.5 ± 0.7

Table 2: Performance comparison (in %) of relation extraction on the TACRED and SemEval datasets. We report the mean and standard deviation of the evaluation metrics by conducting 5 runs of training and testing using different random seeds. We use LSTM+ATT ($S_{\text{matched}} + \mathcal{P}$) as the base model for all semi-supervised baselines and our models.

unlabeled data is exhausted. (3) Mean-Teacher [31] assumes that data points with small differences should have similar outputs. We perturb each unlabeled sentence using word dropout and regularize their outputs to be similar. (4) DualRE [18] jointly trains a relation prediction and a retrieval module which mutually enhance each other by selecting high-quality instances from unlabeled data.

Variants of NERO. (1) Nero w/o $S_{\text{unmatched}}$ removes the loss on unmatched sentences and only use S_{matched} and \mathcal{P} in model training. It only performs hard-matching on rules. (2) Nero-SRM has the same training objective as NERO, but uses the soft rule matcher (SRM) to make predictions. Given a sentence, NERO-SRM finds the most similar rule and returns the rule head as the prediction. It can be taken as a context-agnostic version of NERO (Sec. 4.2).

5.3 Experiment Settings

Implementation. We implement most baselines from the scratch using Tensorflow 1.10 [1] except for those that have released their codes (like PA-LSTM and DualRE). We adapt the baselines to our setting. For supervised and semi-supervised baselines, we use the hard-matched sentences (S_{matched}) as the “labeled data” and the unmatched sentences ($S_{\text{unmatched}}$) as the “unlabeled data”.

Training details. We use pre-trained Glove embeddings [23] to initialize the word embeddings and fine-tune them during training. For NERO, We set the batch size to 50 for hard-matched sentences, and 100 for unmatched sentences. For other baselines, we set the batch size to 50. To exclude the possibility that NERO takes advantage of a larger batch size, we also tried a batch size of 150 for other baselines but it showed no difference in performance. To reduce over-fitting, we adopt the entity masking technique [38] and replace the subject and object with SUBJ/OBJ-NER and regard them as normal tokens. We use a two-layer bi-directional LSTM as the encoder. The hidden dimension is 100 for the LSTM, and 200 for the attention layers. We set β, γ, τ to 0.05, 0.5, and 1.0 respectively. We set α to 1.0 for TACRED, and 0.1 for SemEval. The temperature θ in instance weighting is set to 10. We apply dropout [28] after the




Figure 6: Performance w.r.t. different number of rules and human-annotated labels on TACRED. We show different models’ F1 scores and number of rules or labels used for training the corresponding model.

LSTM with a rate of 0.5. All models are optimized using AdaGrad [6] with initial learning rate 0.5 and decay rate 0.95. For models that require the prediction of NONE (including NERO and all supervised / semi-supervised baselines), we select the threshold (see 4.2) from range [0.0, 1.0] based on the dev set.

5.4 Performance Comparison

Rule-based Models. We first compare the models that solely use rules, as shown in Table 2. Hard-matching with rules achieves a high precision (85% on TACRED) but suffers from the low-recall problem due to its failure in matching instances with similar meanings (only 11.4% on TACRED). Bootstrapping methods (such as BREDS) and unsupervised soft-matching methods (such as NRE) manage to cover more instances (6% gain in recall on TACRED dataset), but they fail to capture contextual information.

Models Trained on S_{matched} . Neural networks are capable of fitting the hard-matched data, but they suffer from over-fitting due to the small data size. Adding the pattern encoder loss (L_{rules}) and the soft matching loss (L_{clus}) boosts the performance by a huge margin, as indicated by LSTM+ATT ($S_{\text{matched}} + \mathcal{P}$) and NERO w/o $S_{\text{unmatched}}$. However, their performance is still far from satisfactory, since they cannot utilize large amounts of unmatched sentences. Data programming does not bring any improvement because our rule mining and labeling method rarely introduces conflicted labeling rules. Hence, data programming with our rules is the same as LSTM+ATT.

Models Trained on $S_{\text{matched}} \cup S_{\text{unmatched}} \cup \mathcal{P}$. Table 2 shows the results for methods that incorporate unmatched data and rules. For methods that actively create labels for unmatched data, namely pseudo-labeling, self-training and DualRE, their performance is even worse than the supervised counterparts (LSTM+ATT ($S_{\text{matched}} + \mathcal{P}$)). It is because in a low-resource setting, the model performance is so low that the created labels are very noisy. Mean-Teacher manages to improve model performance on TACRED, but only by a small margin. Our proposed model, NERO, is relatively indifferent to the noise due to the soft rule matcher, which learns directly from similar textual sentences. This results in a significant gain in precision (14% on TACRED) and recall (4% on TACRED) over

the semi-supervised baselines. Compared to TACRED, the improvement on SemEval is much smaller (only 1.7% in F1). We hypothesize it is because the sentences in SemEval are quite short and contain very simple rule patterns. Thus, the soft-matched sentences can hardly provide additional information. We also try to predict with the soft rule matcher, but the performance is lower in F1 (45.3% on TACRED and 54.9% on SemEval) since it cannot capture contextual information.

5.5 Performance Analysis

1. Performance on Various Amounts of Rules and Human Annotated Labels. To show that our rules are more powerful than human-annotated labels, we show the performance of models trained with different number of rules (30%, 50%, 70%, and 100%) and the number of labels required to reach comparable performance using supervised baseline (LSTM+ATT). We random sample subsets of rules w.r.t. each relation to avoid extreme cases where some relations are not sampled at all. We conduct 5 runs of sampling and training and report the average performance. As shown in Figure 6, NERO consistently outperforms other methods by a huge margin, and its performance increases with the number of rules (from 42.3% to 51.3%). Mean-Teacher also outperforms the base model especially when the number of rules is small, but its performance does not increase much (from 41.3% to 43.6%). It shows that Mean-Teacher generalizes well in low-resource setting but cannot fully utilize the knowledge in rules. In terms of label efficiency, one rule in NERO is equivalent to 10 human-annotated labels in model performance. Even in the base model LSTM+ATT ($S_{\text{matched}} + \mathcal{P}$), one rule is still equivalent to about 4 labels. It demonstrates the superior generalization ability of rules over labels.

2. Performance on Various Amounts of the Raw Corpus. To test the robustness of our model, we plot the performance curve in Fig. 7 when different amounts of raw corpus are available on TACRED. Again, NERO outperforms all methods, and its F_1 score positively correlates to the amount of available data (from 43.4% to 51.3%). We also observe similar phenomena in other methods (all based on LSTM+ATT ($S_{\text{matched}} + \mathcal{P}$) model), which show that matching on rules provides additional knowledge to model training. Just memorizing these rules with a neural model leads to very bad performance. Also, we observe that with only 10% of data, our final model (NERO) already outperforms the best supervised baseline NERO w/o $S_{\text{unmatched}}$. This indicates that the soft matching module can utilize the raw corpus and rules in a more data-efficient way.

3. Performance on Unseen Relations. To show that NERO has the capacity of predicting unseen relations, we evaluate NERO on relations unseen in training. In the experiment, we randomly sample 5 relations as unseen relations and repeat the experiment 10 times. In each time, we use the same sets of unseen relations and the same random seed for all methods. We remove the sampled relations (both sentences and rules) in training while keeping only them in testing. We sample sentences of other relations as NONE relation in testing with the same ratio as the raw corpus (79.5% for TACRED and 17.4% for SemEval). Moreover, since the special tokens SUBJ/OBJ-NER are out-of-vocabulary in both Glove and BERT, we only use the words between two entities in encoding and only use rules with the same entity types to the sentence in inference. To




Figure 7: Performance of different semi-supervised models trained using various amounts of unlabeled sentences randomly sampled from the raw corpus on TACRED.

Method	TACRED			SemEval		
	P	R	F_1	P	R	F_1
Rule (exact match)	100	6.1	10.8	83.2	17.7	28.2
CBOW-GloVe	52.4	86.3	64.7	40.3	45.5	34.7
BERT-base (frozen)	66.2	76.8	69.5	37.8	33.2	35.3
NERO	61.4	80.5	68.9	43.0	54.1	45.5

Table 3: Performance on predicting unseen relations. NERO applies the learned soft rule matcher on unseen relation rules to make predictions.

make predictions by NERO, we find the most similar rule using soft rule matcher. Baseline models include exact rule matching, cosine similarity of average Glove embedding, and cosine similarity of pre-trained BERT embedding [5]. As shown in Table 3, our soft rule matcher achieves similar or better performance compared to all baselines. It shows that instead of just memorizing the rules and the corresponding relation types, the soft rule matcher also learns knowledge about how to better perform rule matching and can even generalize to unseen relations.

5.6 Model Ablation Study

1. Effect of Different Loss Functions. To show the efficacy of our proposed losses (L_{clus} , $L_{unmatched}$ and L_{rules}), we conduct ablation study on loss functions by removing one loss function at a time. As shown in Table 4, all three losses contribute to the final performance, while $L_{unmatched}$ helps the most, which proves the effectiveness of soft-matching. When removing the contrastive loss L_{clus} , the F1 score drops from 51.3 to 46.4, which shows the effectiveness of a trainable SRM. The L_{rules} loss also brings improvement, similar to including more training instances.

Objective	Precision	Recall	F_1
L (ours)	54.0	48.9	51.3
$-L_{rules}$	50.0	47.7	49.0
$-L_{clus}$	50.9	43.0	46.4
$-L_{unmatched}$	41.9	44.3	42.9

Table 4: Ablation Study of Different Training Objectives on TACRED dataset. We remove each loss term one at a time.

2. Effect of Different Soft-matching Models. Besides word-level attention, we also try other soft matching mechanisms including CBOW-Glove, LSTM+ATT, and BERT-base (both frozen and fine-tuned). All methods first summarize the sentences / rules into vector representations, then calculate the matching scores by cosine similarity. We report their performance in Table 5. We observe that our method achieves the best performance despite its simplicity. To our surprise, BERT-base performs much worse than our method. In the experiment we find that BERT tends to given high matching scores for all sentence-rule pairs so our framework cannot distinguish the false pseudo-labels from reliable ones.

Objective	Precision	Recall	F_1
CBOW-Glove	49.4	43.5	46.2
LSTM+ATT	56.2	46.0	50.6
BERT-base (frozen)	45.6	47.6	46.5
BERT-base (fine-tuned)	50.3	45.8	47.9
Word-level attention (ours)	54.0	48.9	51.3

Table 5: Ablation study of different soft-matching models for NERO on the TACRED dataset.

3. Sensitivity Analysis of Model Hyper-parameters. The most important hyper-parameters in NERO are α , β , γ , and τ . We test the sensitivity of these parameters on TACRED. We adjust one hyper-parameter at a time and remain the other three unchanged. The results are shown in Figure 8. We observe that for τ and β , there exists a wide interval where the F1 score remains stable. For α and γ , the F1 score hardly changes with the choice of their values. This is because that $L_{matched}$, L_{clus} , and L_{rules} quickly drop to 0 and $L_{unmatched}$ dominates the joint loss in the training phase, and the impact of γ is lessened by adaptive subgradient method (AdaGrad).

5.7 Case Study

Study on Label Efficiency. To test the label efficiency of rules and sentences in real scenarios, we ask 5 college students in computer science to annotate frequent candidate rules (with frequency ≥ 3) and the unlabeled sentences, during a 40-minute period. For annotating sentences, they are required to provide labels from the




Figure 8: Sensitivity analysis of τ , α , β , and γ on TACRED. We report the mean and standard deviation F1 score by conducting 5 runs of experiments using different random seeds.



Figure 9: Output visualization of SRM. Left: attention weights of words and the soft matching scores between a rule and three sentences. Right: cosine similarity matrix between word embeddings learned with the contrastive loss.

pre-defined relation set $R \cup \text{NONE}$. For annotating candidate rules, they are required to filter out uninformative rules (e.g. “SUBJ-PERSON and OBJ-PERSON”) and assign labels to the remaining ones. For each candidate rule, they are further given an example of hard-matched sentence to help understand the meaning and context of the rule. Then we use each user’s annotations to train the RE model. We use **LSTM+ATT** for labeled sentences and **NERO** for rules.

We report the average number of selected rules / labeled sentences and performance of trained RE model every 5 minutes in Figure 10. We observe that getting rules is nearly as fast as getting labels but is much more efficient. With the same annotation time, NERO can already learn a good RE classifier (with 41.3% F1), while LSTM+ATT cannot learn anything from the labels. In just 40 minutes, NERO can achieve comparable performance to 1000 labels (refer to Figure 6), which requires about 320 minutes to annotate.




Figure 10: Study on label efficiency. Average number of rules / sentences labeled by annotators (dashed line) are shown on the x-axis over the left-hand side; and the performance of models trained with these corresponding labeled rules / sentences (solid line) are shown on the x-axis over the right-hand side. We use NERO and LSTM+ATT as the base model for the labeling rules and sentences, respectively.

Interpretability of Soft Rule Grounding. Fig. 9 gives an example of the similarity matrix, showing that our soft matching module learns meaningful representations of rules for soft matching. The left part of the figure presents 1) which words are more important for building attention-pooled rule/sentence representations, and 2) the soft matching scores between the rule and three sentences. We also show two concrete examples of attention matrix in the right part. For each word-pair between a rule and a sentence, we visualize the cosine similarity matrix between word representations

learned with the contrastive loss. While the word “chairman” in our rule matches perfectly with the same surface word in the sentence, it also has a high similarity score with “Chief” and “Executive”.

6 RELATED WORK

Relation Extraction. Traditional corpus-level RE systems take a rule-based approach, in which rules are either handcrafted [10] or automatically learned from large corpora [2, 4]. Although feature-based soft matching methods (e.g. TF-IDF, CBOW) are used in rule mining, the inference (i.e. labeling new sentences) is still done by hard-matching, which leads to high precision but low recall due to the low coverage of rules. While we use a trainable soft rule matcher, which is more powerful than prior feature-based methods. And we apply soft-matching to both training and inference phase in a consistent manner. Recent RE models [36, 38] successfully apply deep neural networks with the help of large scale datasets such as SemEval [11] and TACRED [38]. However, their performance degrades heavily when labeled data is insufficient. Some work [9, 29] proposes to train a RE classifier from natural language explanations, which are first transformed to labeling rules by a semantic parser then applied to raw corpora by hard-matching. But none of them utilizes the data that cannot be matched by rules, which is of massive amount. This motivates us to propose a framework that can fully utilize the information from both rules and raw corpora.

Semi-Supervised Learning. Our work is relevant to semi-supervised learning, if we consider rule-annotated data as labeled among the corpus, e.g., self-training [26], mean-teacher [31], and semi-supervised VAE [34]. In relation extraction, one line of work proposes to alleviate the low recall problem in rule-based approach using bootstrapping framework [2] and revised pattern matching [4, 8]. However, their settings are tailored for corpus-level RE. Another line of work applies self-training framework in supervised learning models. However, These models turn out to be ineffective in rule-labeled data due to potentially large difference in label distribution, and the generated pseudo-labels may be quite noisy.

7 CONCLUSION

In this paper, we proposed a novel framework, named NERO, for label-efficient relation extraction. We automatically extracted frequent patterns from large raw corpora and asked users to annotate them to form labeling rules. To increase the coverage of these rules, we further proposed the soft-matching mechanism, where unlabeled sentences are annotated by their most semantically similar labeling rules and weighted in training the RE model. Experiments on two public datasets proved the effectiveness of our framework.

REFERENCES

- [1] Martin Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <https://www.tensorflow.org/> Software available from tensorflow.org.
- [2] Eugene Agichtein and Luis Gravano. 2000. Snowball: Extracting relations from large plain-text collections. In *Proceedings of the fifth ACM conference on Digital libraries*. ACM, 85–94.
- [3] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473* (2014).
- [4] David S Batista, Bruno Martins, and Mário J Silva. 2015. Semi-supervised bootstrapping of relationship extractors with distributional semantics. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. 499–504.
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [6] John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* 12, Jul (2011), 2121–2159.
- [7] Zhijiang Guo, Yan Zhang, and Wei Lu. 2019. Attention Guided Graph Convolutional Networks for Relation Extraction. *arXiv preprint arXiv:1906.07510* (2019).
- [8] Pankaj Gupta, Benjamin Roth, and Hinrich Schütze. 2018. Joint bootstrapping machines for high confidence relation extraction. *arXiv preprint arXiv:1805.00254* (2018).
- [9] Braden Hancock, Paroma Varma, Stephanie Wang, Martin Bringmann, Percy Liang, and Christopher RAI. 2018. Training Classifiers with Natural Language Explanations. *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (2018). <https://doi.org/10.18653/v1/p18-1175>
- [10] Marti A Hearst. 1992. Automatic acquisition of hyponyms from large text corpora. In *Proceedings of the 14th conference on Computational linguistics-Volume 2*. Association for Computational Linguistics, 539–545.
- [11] Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Sebastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. 2009. Semeval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals. In *Proceedings of the Workshop on Semantic Evaluations: Recent Achievements and Future Directions*. Association for Computational Linguistics, 94–99.
- [12] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [13] Jing Jiang and ChengXiang Zhai. 2007. Instance weighting for domain adaptation in NLP. In *Proceedings of the 45th annual meeting of the association of computational linguistics*. 264–271.
- [14] Meng Jiang, Jingbo Shang, Taylor Cassidy, Xiang Ren, Lance M Kaplan, Timothy P Hanratty, and Jiawei Han. 2017. Metapad: Meta pattern discovery from massive text corpora. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 877–886.
- [15] Rosie Jones, Andrew McCallum, Kamal Nigam, and Ellen Riloff. 1999. Bootstrapping for text learning tasks. In *IJCAI-99 Workshop on Text Mining: Foundations, Techniques and Applications*, Vol. 1.
- [16] Dong-Hyun Lee. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks.
- [17] Shen Li, Hengru Xu, and Zhengdong Lu. 2018. Generalize Symbolic Knowledge With Neural Rule Engine. *arXiv preprint arXiv:1808.10326* (2018).
- [18] Hongtao Lin, Jun Yan, Meng Qu, and Xiang Ren. 2019. Learning Dual Retrieval Module for Semi-supervised Relation Extraction. In *The Web Conference*.
- [19] 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*. 3111–3119.
- [20] 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*. Association for Computational Linguistics, 1003–1011.
- [21] Ndapandula Nakashole, Gerhard Weikum, and Fabian Suchanek. 2012. PATTY: a taxonomy of relational patterns with semantic types. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*. Association for Computational Linguistics, 1135–1145.
- [22] Paul Neculouiu, Maarten Versteegh, and Mihai Rotaru. 2016. Learning Text Similarity with Siamese Recurrent Networks. In *Rep4NLP@ACL*.
- [23] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*. 1532–1543.
- [24] Meng Qu, Xiang Ren, Yu Zhang, and Jiawei Han. 2018. Weakly-supervised Relation Extraction by Pattern-enhanced Embedding Learning. In *Proceedings of the 2018 World Wide Web Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 1257–1266.
- [25] 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*. 3567–3575.
- [26] Chuck Rosenberg, Martial Hebert, and Henry Schneiderman. 2005. Semi-supervised self-training of object detection models. (2005).
- [27] Benjamin Roth, Tassilo Barth, Michael Wiegand, Mittul Singh, and Dietrich Klakow. 2014. Effective slot filling based on shallow distant supervision methods. *arXiv preprint arXiv:1401.1158* (2014).
- [28] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research* 15, 1 (2014), 1929–1958.
- [29] Shashank Srivastava, Igor Labutov, and Tom M. Mitchell. 2017. Joint Concept Learning and Semantic Parsing from Natural Language Explanations. In *EMNLP*.
- [30] Mihai Surdeanu, Julie Tibshirani, Ramesh Nallapati, and Christopher D Manning. 2012. Multi-instance multi-label learning for relation extraction. In *Proceedings of the 2012 joint conference on empirical methods in natural language processing and computational natural language learning*. Association for Computational Linguistics, 455–465.
- [31] Antti Tarvainen and Harri Valpola. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in neural information processing systems*. 1195–1204.
- [32] Linlin Wang, Zhu Cao, Gerard de Melo, and Zhiyuan Liu. 2016. Relation Classification via Multi-Level Attention CNNs. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, 1298–1307. <https://doi.org/10.18653/v1/P16-1123>
- [33] Peter Willett. 2006. The Porter stemming algorithm: then and now. *Program* 40, 3 (2006), 219–223.
- [34] Weidi Xu, Haoze Sun, Chao Deng, and Ying Tan. 2017. Variational autoencoder for semi-supervised text classification. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- [35] Mo Yu and Mark Dredze. 2015. Learning composition models for phrase embeddings. *Transactions of the Association for Computational Linguistics* 3 (2015), 227–242.
- [36] Daojian Zeng, Kang Liu, Yubo Chen, and Jun Zhao. 2015. Distant supervision for relation extraction via piecewise convolutional neural networks. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. 1753–1762.
- [37] Yuhao Zhang, Peng Qi, and Christopher D Manning. 2018. Graph convolution over pruned dependency trees improves relation extraction. *arXiv preprint arXiv:1809.10185* (2018).
- [38] Yuhao Zhang, Victor Zhong, Danqi Chen, Gabor Angeli, and Christopher D Manning. 2017. Position-aware attention and supervised data improve slot filling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. 35–45.
- [39] Peng Zhou, Wei Shi, Jun Tian, Zhenyu Qi, Bingchen Li, Hongwei Hao, and Bo Xu. 2016. Attention-Based Bidirectional Long Short-Term Memory Networks for Relation Classification. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, Berlin, Germany, 207–212. <https://doi.org/10.18653/v1/P16-2034>