

PromptRE: Weakly-Supervised Document-Level Relation Extraction via Prompting-Based Data Programming

Chufan Gao¹, Xulin Fan¹, Jimeng Sun¹, Xuan Wang²

¹University of Illinois Urbana-Champaign ²Virginia Tech
chufan2@illinois.edu, xuanw@vt.edu

Abstract

Relation extraction aims to classify the relationships between two entities into pre-defined categories. While previous research has mainly focused on sentence-level relation extraction, recent studies have expanded the scope to document-level relation extraction. Traditional relation extraction methods heavily rely on human-annotated training data, which is time-consuming and labor-intensive. To mitigate the need for manual annotation, recent weakly-supervised approaches have been developed for sentence-level relation extraction while limited work has been done on document-level relation extraction. Weakly-supervised document-level relation extraction faces significant challenges due to an imbalanced number "no relation" instances and the failure of directly probing pretrained large language models for document relation extraction. To address these challenges, we propose PromptRE, a novel weakly-supervised document-level relation extraction method that combines prompting-based techniques with data programming. Furthermore, PromptRE incorporates the label distribution and entity types as prior knowledge to improve the performance. By leveraging the strengths of both prompting and data programming, PromptRE achieves improved performance in relation classification and effectively handles the "no relation" problem. Experimental results on ReDocRED, a benchmark dataset for document-level relation extraction, demonstrate the superiority of PromptRE over baseline approaches.

1 Introduction

Relation extraction is a crucial task in natural language processing that aims to classify the relationships between two entities (e.g., Pacific Fair and Queensland) into pre-defined categories (e.g., located in). It has various downstream applications such as question answering (Veena et al.,

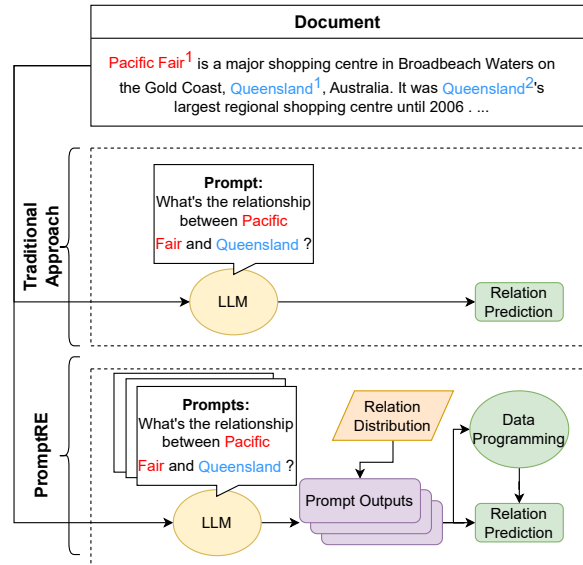


Figure 1: Differences between the naive approach and PromptRE for weakly-supervised document-level relation extraction. We investigate various prompts and different ways to combine the prompting outputs using data programming. Furthermore, PromptRE incorporates the entity type and relation distributions as prior knowledge to improve the classification performance.

2017) and knowledge graph construction (Distiawan et al., 2019).

While previous research has mainly focused on relation extraction within a single sentence, recent studies have expanded the scope to document-level relation extraction (Yao et al., 2019). Traditional relation extraction methods (Tan et al., 2022a; Ma et al., 2023) heavily rely on human annotation for training data, which is time-consuming and labor-intensive. To mitigate the need for manual annotation, recent weakly-supervised approaches (Sainz et al., 2021; Yang and Agrawal, 2023) have been developed for relation extraction with minimal or no manual annotation. For example, Qu et al. (2018) extracted textual patterns from seed examples and used those patterns as weak supervisions for re-

lation extraction. [Sainz et al. \(2021\)](#) represented each relation class using a label verbalizer and then solving the relation extraction task by a textual entailment model. [Wang et al. \(2022a\)](#) analyzed an "extremely unlabeled" scenario where each relation type had only one instance, reducing the training set to about five thousand labeled relation triplets. However, these methods were primarily designed and evaluated for sentence-level relation extraction, which limits their generalizability to document-level relation extraction datasets like ReDocRED ([Tan et al., 2022b](#)), where the presence of a substantial number of "No Relation" or NA classes poses additional challenges.

To address this limitation, we study the problem of weakly-supervised document-level relation extraction. Recent large language models (LLMs) have achieved great success in a wide range of natural language processing tasks ([Brown et al., 2020](#); [Touvron et al., 2023](#)). We investigate the ability of the pretrained large language models on the document-level relation extraction task. We focus on three pretrained large language models: UnifiedQA ([Khashabi et al., 2020, 2022](#)), LLaMA, LLaMA2 ([Taori et al., 2023](#); [Touvron et al., 2023](#)), and ChatGPT ([Ouyang et al., 2022](#)).

UnifiedQA is a T5 model ([Raffel et al., 2020](#)) pretrained on four different question-answering settings: extractive, abstractive, multiple-choice, and yes/no questions. UnifiedQA performs comparably to specialized state-of-the-art models on most relation extraction datasets. **ChatGPT**, developed by OpenAI, is a powerful generative large language model known for its impressive generalization capabilities. However, the closed-source nature of the ChatGPT model limits its accessibility for downstream applications. We utilize the text output of ChatGPT without accessing its internal embedding space or doing any model fine-tuning. **LLaMA** is a collection of foundation language models ranging from 7B to 65B parameters trained on only publicly available datasets. After fine-tuning on an instruction-following dataset ([Taori et al., 2023](#)), LLaMA and LLaMA2 are able to produce reasonable responses to the input instructions. We use LLaMA-7B and LLaMA2-7B, which has a good balance between model performance and efficiency.

We propose PromptRE, a novel weakly-supervised document-level relation extraction method that combines prompting-based techniques with data programming (Figure 1). Given a known type-relation distribution, we first investigate var-

ious ways of prompting the pretrained large language models for relation classification. We then investigate different ways to select the most confident outputs using data programming, a technique that combines multiple sources of weak supervision. By leveraging the strengths of both prompting and data programming, we achieve improved performance in relation classification and effectively handle the "no relation" problem. Furthermore, we leverage ChatGPT as a summarizer to extract relevant information about the entities of interest. This plays a crucial role especially when dealing with lengthy documents which contain unnecessary extra information. To the best of our knowledge, we are the first to propose weakly-supervised document-level relation extraction. Our contributions are summarized as follows:

1. We propose the first weakly-supervised method, PromptRE, for the document-level relation extraction task.
2. PromptRE is a novel method that combines various types of prompting outputs with data programming. PromptRE further incorporates the label distribution and entity types as prior knowledge to improve performance.
3. Extensive experiments on the ReDocRED dataset demonstrate the capability of PromptRE over baseline methods. Ablations provide a comprehensive study on weakly-supervised inference ability. Multiple case studies show the incompleteness of existing document relation extraction datasets.

2 Related Work

2.1 Document-level Relation Extraction

Document-level relation extraction is a crucial task in natural language processing, as more than 40.7% of relations require multiple sentences to extract ([Yao et al., 2019](#)). Consider an example of document-level relation extraction in Figure 2. The task is to identify the relationship between a pair of entities ("Pacific Fair" and "Queensland") in the input document. Each of the entities has two mentions in the text (denoted by superscripts). To infer their relationship, it is evident that the mention-mention pair involving the first mention of each entity provides the most valuable information for extracting the relationship between them.

Pacific Fair¹ is a major shopping centre in Broadbeach Water on the Gold Coast, Queensland¹, Australia. It was Queensland²'s largest regional shopping centre until 2006. Pacific Fair² was developed by Hooker Retail Developments and opened in 1977 on what was swampland with 96 specialty stores and two anchor tenants.

Figure 2: Sample document relation extraction task from DocRED (Yao et al., 2019). The red text indicates the head entity, and the blue text indicates the tail entity. Here, the head is related to the tail by "P131: located in the administrative territorial entity".

Compared to sentence-level relation extraction, document-level relation extraction requires reasoning over multiple sentences which requires neural models to model long-range information. Additionally, entities may contain multiple mentions, which could include irrelevant information. However, this also allows for more information to model the relationship between entity-entity pairs.

Pretrained language models, such as BERT-based models (Xu et al., 2021), have demonstrated significant success in document-level relation extraction. For example, BERT-based methods have employed techniques like hierarchical inference networks (Tang et al., 2020), improved co-reference reasoning (Ye et al., 2020), and adaptive thresholding. Additionally, graphical neural networks (GNNs) (Zeng et al., 2020) have also been utilized for modeling document-level relation extraction. GNNs are used for feature learning on a coreference graph (Sahu et al., 2019), edge-oriented learning techniques (Christopoulou et al., 2019), utilizing attention mechanisms (Guo et al., 2019), and applying iterative refinement strategies for aggregating multi-hop information (Nan et al., 2020). Moreover, several works have proposed new loss functions to tackle the class-imbalance problem in document-level relation extraction (Zhou et al., 2021; Tan et al., 2022a).

However, previous research on document-level relation extraction has relied heavily on human annotation for generating training data, which can be a time-consuming and labor-intensive process. Limited work has been conducted on document-level relation extraction methods that do not require human annotation.

2.2 Weakly-Supervised Relation Extraction

Weakly supervised methods have been extensively explored for relation extraction (Jiang, 2009; Huang and Wang, 2017; Qu et al., 2018; Wang et al., 2018; Li et al., 2018). For example, Huang and Wang (2017) utilized residual connections and convolutional neural networks (CNNs) to select relevant candidates to enhance supervised relation classification. Qu et al. (2018) extracted textual patterns from seed examples to provide additional supervision. Phi et al. (2018) introduced a ranking-based approach for seed selection, improving bootstrapping and distantly supervised relation extraction. Sainz et al. (2021) proposed representing each relation class using a label verbalizer and addressing the relation extraction task with a textual entailment model. Wang et al. (2022a) analyzed an "extremely unlabeled" scenario where each relation type had only one instance and reduced the training set to a smaller number of labeled relation triplets (but still contained more than 5000 training triplets).

However, these methods were either primarily designed and evaluated for sentence-level relation extraction or still require many labels, which limits their generalizability to our weakly-supervised document-level relation extraction task.

3 Methodology

We propose PromptRE, a weakly supervised document-level relation extraction method that combines large language model prompting with data programming. An illustration of the overall framework of PromptRE is shown in Figure 3.

3.1 Problem Definition

In our task formulation, we consider a document D consisting of M sentences (s_1, s_2, \dots, s_M) and N entities (e_1, e_2, \dots, e_N). Given this document D , a specified entity pair (e_{head}, e_{tail}), and a set of positive entity-entity relations (r_1, r_2, \dots, r_k), the objective is to predict a set of relations ($\hat{r}_1, \hat{r}_2, \dots, \hat{r}_p$) between the pair of entities based on the information extracted from the document. Note that each entity can have multiple occurrences within the document D .

3.2 Entity-Oriented Document Preprocessing

One challenge in document-level relation extraction is the long context. Models need to be able to find and focus on the information specific to the

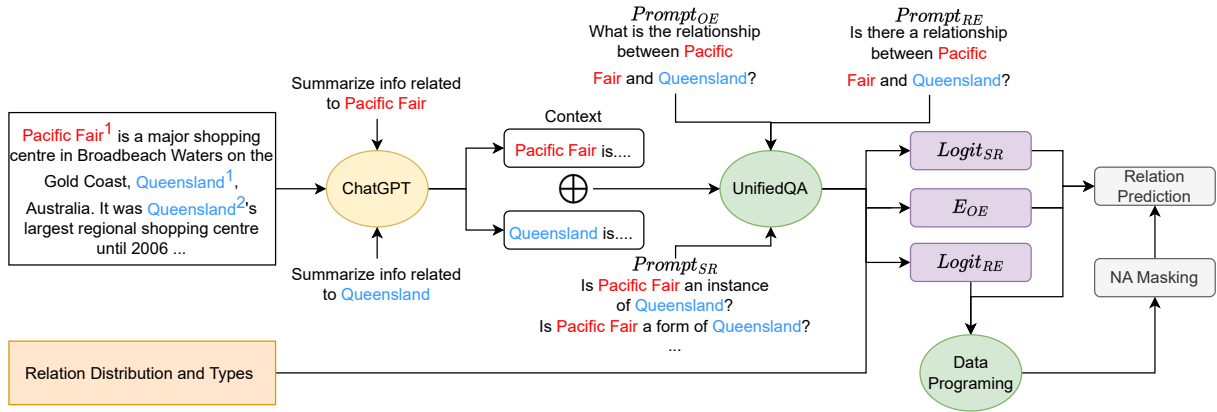


Figure 3: Overall framework of PromptRE. Given an example document and an expected relation distribution, we first summarize the relevant portions of the text regarding both entities and concatenate (\oplus) them together for entity-relevant context. Then, we use a variety of prompts to obtain: 1. The model prediction of each valid relation, 2. the open-ended model relation prediction, and 3. The model predicts the existence of a relationship. The three outputs are then used for relation prediction and data programming for addressing the "no relation" issue (referred to as NA Masking).

given pair of entities. In our PromptRE pipeline, we leverage the power of ChatGPT to solve this problem. For example, given an entity "Pacific Fair" and the document shown in Figure 2, we ask ChatGPT "Based on the given paragraph, summarize the information about "Pacific Fair" \n Pacific Fair is a major shopping center in Broadbeach Waters...".

ChatGPT will generate a natural language summary of the information about the entity in the paragraph. We concatenate the summary of head and tail entities to form the text description of the two entities (denoted as <context>). In later stages, this summary is utilized as the context in place of the original document for relation prediction.

3.3 Relation Prediction Via Prompting

For predicting the relation class, we explore two approaches: Relation Specific Prompting and Open Ended Prompting. For both approaches, we query all non-identity entity-entity pairs in the documents. That is, if one document has n_e entities, we query $n_e(n_e - 1)$ times.

Relation-Specific Prompting We prompt the large language models over all possible relation classes, for all possible entity pairs. For each relation class, we hand-craft a yes-no question. For example, for the relation class, "instance of", our hand-crafted version of the question is "Is "Pacific Fair" an instance of "Queensland" <context> ?". To quantify the certainty of the large language model, we obtain a prediction score $Logit_{SR}$ by subtracting the logit

of the 'no' output from the logit of the 'yes' output. This logit score $Logit_{SR}$ is calculated over each relation class and normalized to obtain a predicted relation class distribution.

Open-Ended Prompting With the open-ended approach, we only prompt the large language models once for each entity pair with the question "What's the relationship between "Pacific Fair" and "Queensland" <context> ?". From there, we obtain the entity pair embedding as follows:

$$E_{OE} = \frac{\sum_{i=1}^{|embed|} LLM(P_{OE})_{embed}}{|embed|}$$

of the large language model output. Note that $E_{OE} \in \mathbb{R}^{N_{hidden}}$, where N_{hidden} is of hidden dimension size from the large language model.¹

In addition, we encode each of the relation classes with a relation embedding $E_{rels} \in \mathbb{R}^{N_{rels} \times N_{hidden}}$ where N_{rels} is the total number of relation classes. We use cosine similarity between the entity pair embedding E_{OE} and the relation embeddings E_{rels} to compute a score over each relation, which is then normalized to obtain a predicted class relation distribution.

Utilizing the Type Distribution In our problem setup, we assume that we know the distribution of the relation classes given the types. We argue that this assumption, while strong, is reasonable. For

¹ N_{hidden} is 1024 for UnifiedQA and 4096 for LLaMA and LLaMA2.

example, it makes sense that no "Person - Person" entity pair could have the relation of "country of citizenship", as it does not make logical sense. It would be much more reasonable if the entity types were "location - person". We assume that this implicit knowledge is provided by the expert on the domain on which this framework is applied, and therefore we add the relevant relation distribution to the predicted probabilities of the raw scores, given the entity type pairs. In our experiments, we estimate this relation distribution from the ReDocRED dataset. More details can be found in Appendix A.

Multi-Label Prediction Since each document has multiple possible labels, we take only the top p percentile of confident predictions over all of the valid relation classes. Note that in this step, we do not consider the No-Relation class.

3.4 Addressing the No-Relation Issue

Although we can extract potential relations via relation prompting from the previous section, we face the issue of false positives in relation prediction due to the large number of No-Relation classes. To address this, we design PromptRE to choose only the most confident relation predictions.

Relation Existence Prompting To obtain the model prediction for the existence of a relation in the input text, we prompt the model with the following prompt: "Is there a relationship between "Pacific Fair" and "Queensland" <context> ?" To quantify the certainty of the large language model, we obtain a prediction score $Logit_{RE}$ by subtracting the logit of the 'no' output from the logit of the 'yes' output. This score is used to preserve only the most confident model predictions.

Data Programming (DP) We combine multiple sources of weak supervision to select highly confident predictions from the previous step of relation prompting. Data programming is a framework to create denoised pseudo-labels from multiple sources of weak supervision from labeling functions (Ratner et al., 2016, 2019).

A labeling function (LF) is a noisy heuristic that takes in data and assigns labels to unlabelled data or abstains from making a prediction. For example, $f(\text{text}) = \text{return SPAM if "http" in text else ABSTAIN}$ is a labeling function for spam detection.

At a high level, we frame the problem as dependency graph G_{source} where each labeling function λ_i is dependently conditioned on the true label Y . In our case, we assume conditional independence of all $\lambda_i|Y$. For this case, the dependency graphs will have observable cliques $\mathcal{O} = \{\lambda_i, i \in n_{lf}\} \subset C$, where n_{lf} is the number of labeling functions.

From here, we can analyze the covariance matrix of an observable subset of the cliques in G_{source} , leading to a matrix completion approach for recovering estimated accuracies μ (used in the final label model to predict $P(\mathbf{Y}|\boldsymbol{\lambda})$).

We assume that $\mu = \mathbb{E}(\psi(C))$ where $\psi(C)$ is vector of indicator random variables for all combinations of all but one of the labels emitted by each variable in clique C .

The norm of the covariance of observed LFs cliques O and separator set S cliques $Cov(\psi(O) \cup \psi(S))$ can be used to recover μ .

$$Cov(\psi(O) \cup \psi(S)) = \Sigma = \begin{bmatrix} \Sigma_O & \Sigma_{OS} \\ \Sigma_{OS}^T & \Sigma_S \end{bmatrix} \quad (1)$$

Its inverse is:

$$K = \Sigma^{-1} = \begin{bmatrix} K_O & K_{OS} \\ K_{OS}^T & K_S \end{bmatrix} \quad (2)$$

Applying block matrix inversion, we get:

$$K_O = \Sigma_O^{-1} + c \Sigma_O^{-1} \Sigma_{OS} \Sigma_{OS}^T \Sigma_O^{-1}$$

$$c = (\Sigma_S - \Sigma_{OS}^T \Sigma_O^{-1} \Sigma_{OS})$$

Let $z = \sqrt{c} \Sigma_O^{-1} \Sigma_{OS}$, then

$$K_O = \Sigma_O^{-1} + z z^T$$

Solving for z can directly recover μ via Algorithm 1 in Ratner et al. (2019).

Reducing NA Predictions via Weak Supervision

To address the "No relation" issue, we attempt to combine multiple sources of weak supervision through data programming to obtain a stronger prediction. We consider three sources of weak supervision below.

1. The first source is the logit of relation-existence prompting $Logit_{RE}$. A higher logit indicates a better likelihood of a relationship between the pair of entities. Additionally, by rephrasing the prompt in different ways, we obtain different views of the model opinion on the existence of a relationship. Other paraphrases could include

Table 1: Statistics of the Re-DocRED dataset as well as the entire paraphrased ChatGPT Summary for every unique entity pair. Although the total number of unique documents is large, they are constructed by concatenating relevant information regarding both entities (and only require n_e calls per document).

Stats	Re-DocRED			ChatGPT Summary
	Train	Dev	Test	Dev
# Docs	3,053	500	500	1,193,092
Avg. # Entities	19.4	19.4	19.6	19.4
Avg. # Triples	28.1	34.6	34.9	34.6
Avg. # Sentences	7.9	8.2	7.9	5.3

"Is there a direct relationship between e_{head} and e_{tail} ?", "Does e_{head} have any connection to e_{tail} ?", and more.

- The second source is the average logit of relation-specific prompting $Logit_{SR}$. The motivation is that if the entity pair has a low average logit for every relation-specific prompt, then it is not relevant to any of the relation classes and there is likely no relationship between the pair of entities.
- The third source is the average cosine similarity between the entity pair embedding E_{OE} and the relation embedding E_{Rels} . Similar to the previous motivation for relation-specific prompting, if an entity pair embedding is very dissimilar from every relation embeddings, then there is likely no relationship between the pair of entities.

To summarize, we combine the three sources of weak supervision as input to the data programming model. Then, we take the $argmax$ from the probabilistic predictions of the data programming model and it as a mask to ensure that only the most probable predictions remain. Following the approach of Ratner et al. (2019), we also fit a logistic regression model on $X = E_{OE}$ and label model predictions $\hat{Y} \sim P(Y|\lambda)$ in order to smooth the decision boundaries.

4 Experiments

4.1 Dataset

To evaluate our methodology, we use ReDocRED (Tan et al., 2022b), an open-access, document-level relation extraction dataset that improves upon the popular DocRED dataset (Yao et al., 2019) by

resolving incompleteness, addressing logical inconsistencies, and correcting coreferential errors. Table 1 shows the amount of training data available for all data splits as well as the ChatGPT-paraphrased entity-relevant text summary. Note that we primarily use the Dev set of ReDocRED for our experiments for computational practicality.

4.2 Experimental Settings

For the large language models for relation extraction, we compared UnifiedQA (Khashabi et al., 2020, 2022) (both *3b* and *large* versions) and Alpaca-lora²—a reproduction of the Stanford Alpaca LLaMA model (Taori et al., 2023; Touvron et al., 2023; Wang et al., 2022b) using LoRa (Hu et al., 2021).

Some experiments can only be run with a subset of these models. For example, $logits_{SR}$ is highly expensive to compute as it requires $(n_e^2 - n_e) \times N_{rels}$, so we only run the UnifiedQA-Large for this score computation. For all other score computations, we may use all models: UnifiedQA-3b, UnifiedQA-3b, and LLama-7b as the base models. Additionally, we also perform weak supervision experiments without $logits_{SR}$ due to its high cost (See App. C for more details). In our experiments, we use precision, recall, and F1 scores as the evaluation metrics for the performance comparison. More details about these evaluation metrics can be found in Appendix B.

4.3 Results

Table 2 shows the main results of our experiments. We observed that the logit performance of prompting every entity-entity pair with the relevant relation prompt does not perform as well as using the cosine similarity of the open-ended QA embeddings and the prompt embeddings. We suspect that this may be due to several reasons, including the lack of regularization of the score output. Additionally, it is possible that using cosine similarity allows the model to capture a more semantically meaningful snapshot of its response, rather than just a single scalar value.

As expected, using the ground truth NA labels leads to a large improvement over relaxing the assumption. It demonstrates the difficulty in determining the existence of relations in documents under weak supervision and points out an exciting direction for future research.

²<https://github.com/tloen/alpaca-lora>

Table 2: We compare all results as ran on UnifiedQA-large, UnifiedQA-3b, and LLaMA-7b denoted by $\{large, 3b, llama, llama2\}$ for different models. Simple RE denotes using the thresholded output of $Logit_{RE}$ without data programming. MV denotes using the baseline majority vote label model. DP denotes using data programming for weak supervision. Knowing the True NA Mask indicates using the ground truth relation existence labels. Bold denotes best performance.

Methods	F1	Ign F1	Precision	Recall
Weakly Supervised Methods				
Logits _{large} + Simple RE	5.5975	4.8830	3.4246	15.3147
Embed _{large} Sim. + Simple RE	9.2030	7.9314	5.6304	25.1794
Embed _{large} Sim. + MV	9.5576	8.5099	7.9800	11.9128
Embed _{large} Sim. + DP (PromptRE)	10.2232	8.7384	6.3969	25.4397
Embed _{3b} Sim. + Simple RE	9.1290	7.8723	5.5852	24.9769
Embed _{3b} Sim. + MV	9.4973	8.4576	7.9296	11.8375
Embed _{3b} Sim. + DP (PromptRE)	10.1465	8.6738	6.3489	25.2488
Embed _{llama} Sim. + Simple RE	9.2136	7.9386	5.6369	25.2083
Embed _{llama} Sim. + MV	6.6330	6.0858	7.7816	5.7799
Embed _{llama} Sim. + DP (PromptRE)	9.9368	8.5486	6.4909	21.1814
Embed _{llama2} Sim. + Simple RE	9.3214	8.0442	5.7029	25.5034
Embed _{llama2} Sim. + MV	8.1840	7.4837	8.4589	7.9264
Embed _{llama2} Sim. + DP (PromptRE)	10.5586	9.0371	6.5623	27.0019
Knowing the True NA Mask				
Embed _{large} Sim. + True NA Mask	46.6324	42.1369	38.3962	59.3670
Embed _{3b} Sim. + True NA Mask	46.4416	41.9580	38.2390	59.1240
Embed _{llama} Sim. + True NA Mask	46.6915	42.1909	38.4448	59.4423
Embed _{llama2} Sim. + True NA Mask	46.7824	42.2882	38.5197	59.5580
Supervised Methods				
DREEAM (Ma et al., 2023)	80.73	79.66	-	-
KD-DocRE (Tan et al., 2022a)	78.28	77.60	-	-

Table 3: Experimental results with or without using ChatGPT for entity-oriented document preprocessing. *large* and *3b* denote the UnifiedQA model we use to compute cosine similarities.

w/o ChatGPT	F1	Ign F1	Precision	Recall
<i>large</i> + Simple RE	9.1924	7.9171	5.6240	25.1504
<i>large</i> + DP	9.8235	8.6655	6.6853	18.5142
<i>3b</i> + Simple RE	9.1861	7.9125	5.6201	25.1331
<i>3b</i> + DP	9.9087	8.7183	6.5299	20.5334
w/ ChatGPT				
<i>large</i> + DP	10.2232	8.7384	6.3969	25.4397

Table 4: Experimental results with Relation Type Distribution using Logits_{large} as the baseline model.

No Type Dist.	F1	Ign F1	Precision	Recall
Simple RE	0.3997	0.3609	0.2445	1.0935
DP	0.3604	0.3332	0.2278	0.8621
Only Type Dist.				
Simple RE	3.4499	2.8987	1.8514	25.2430
DP	4.6606	3.8770	2.5747	24.5429

Effect of Language Model Size The performance comparison between different model sizes is shown in Table 2. One observation is that the UnifiedQA-large model performs better than the UnifiedQA-3b model across all the metrics. Khashabi et al. (2022) observed similar results between the large and 3b models, so this is not entirely unexpected. Additionally, it is interesting to see that even the LLaMA-7b model, the largest

model we consider, provides no significant benefit over using the UnifiedQA-large model. This implies that prompting by itself does not work well for document-level relation extraction, as we see that the relations predicted are generally not of the same format as the true relation classes (see Section 4.4).

Effect of Entity-Oriented Document Preprocessing In this ablation study, We investigate the effects of using ChatGPT for entity-oriented document preprocessing (inspired by (Yu et al., 2022)). From Table 3, we see that using the ChatGPT-generated entity-oriented document summaries yields better performance compared to using the original document as input for relation extraction. However, this suggests that further research could potentially avoid expensive calls to the private model without sacrificing too much performance.

Effect of Relation Type Distribution We further investigate the effect of relation type distribution on relation class prediction. As shown in Table 4, the type distributions are strong signals as prior knowledge to enhance the performance of weakly-supervised relation extraction. We observe that using type distribution only without any other weak supervision achieves almost half of the performance of PromptRE in Table 2. Our results from Table 2 further show that the combination of

Table 5: An example of a relation prediction of a portion of labeled positive relation tuples. Each relation tuple has the format: (Head entity, Tail entity, Ground Truth Relation). The predicted text is shown below after prompting the model using an open ended prompt (e.g. "Is there a relationship between e_{head} and e_{tail} ?").

Labeled Relation Tuple: (Ross Patterson Alger; August 20 , 1920; date of birth)
 Ross Patterson Alger was born on August 20, 1920 in Prelate, Saskatchewan. He became a politician and served as mayor of Calgary from 1977 to 1980.

Labeled Relation Tuple: (Ross Patterson Alger; January 16 , 1992; date of death)
 The relationship between "Ross Patterson Alger" and "January 16 , 1992" is that Ross Patterson Alger died on January 16, 1992 due to cancer which was initially detected in 1985...

Labeled Relation Tuple: (Canadian; Alberta; contains administrative territorial entity)
 Yes, there is a relationship between "Canadian" and "Alberta". Ross Patterson Alger was born in Saskatchewan and moved to Alberta...

Unlabeled Relation Tuple: (Ross Patterson Alger; 1942; no relation)
 Ross Patterson Alger was born in Saskatchewan, received a bachelor of commerce degree from the University of Alberta in 1942, and died in 1992 after battling cancer. He served in the Royal Canadian Air Force during World War II.

Unlabeled Relation Tuple: (Ross Patterson Alger; Calgary City Council; no relation)
 Ross Patterson Alger was an alderman on Calgary City Council from 1971 to 1974 before being elected as the mayor in 1977.

prompting + type distribution performs the best.

4.4 Case Studies

We analyze some example outputs from the predictions of the LLaMA-7b model, as shown in Table 5. We see that practically, the LLaMA model output is biased towards much longer and more detailed text than is required for the relation prediction problem. This could explain why the embeddings between the answers and the relation text would be difficult to correlate, leading to worse performance. Furthermore, the third example shows an instance of an indirect relation. It is true that Ross Patterson Alger was born in Canada and moved to another part of Canada. This is a common failure case with the responses—the relation prediction is too specific to the original text. The final two examples indicate a weakness in the dataset. As with any relation extraction dataset, ReDocRED is not complete, and the large language model was able to pick up on two relations not in the ground truth labeled set—"received a bachelor's degree from" alderman" respectively.

5 Conclusion

In this paper, we investigate several methods to integrate prompting and data programming for relation classification and evaluate our model on ReDocRED. Results show that our best results yield around 10.2 F1 on the development set, a promising result for almost no supervision. Since this is a novel application, further research is required to investigate strategies for improvement. Some ideas include the following.

The NA Issue: The large number of "no relations" continues to be an issue for less than supervised methods for document relation extraction on existing datasets. Further work should focus on more efficient and accurate ways to mine distant labels to address this issue. One major roadblock that coincides with NA is the lack of complete labels in the dataset as shown in the case studies. Future work could improve on existing document relation extraction datasets accordingly.

Extending to the Few-Shot Case: It is usually possible to query human experts for a few examples of the required classification task. Researching ways to take maximum advantage of a small set of labels would also be highly practical, and would not require much extra effort on the annotators. This could also tune a model to better address the "No Relation" issue.

Final Thoughts: We find that weakly supervised document-level relation extraction is a uniquely difficult problem due to the incomplete labels in popular datasets, and we propose PromptRE to attempt to solve it by combining prompting and data programming. We show the effect of tuning different experimental setups, including model size, entity-oriented summarization, and the effect of our relation-type distribution assumption. Case studies support the finding that existing document-level relation extraction datasets may be severely lacking in label completeness. Although the results are a considerable margin from contemporary supervised methods, we hope that this work can serve as a stepping stone in this novel area of less-than-supervised document relation extraction.

Limitations

Although we investigated multiple different LLMs and parameters and the type relation distribution for relation prediction as well as addressing the false positives, the performance we attained is still limited compared to supervised methods on the same task. Additionally, relation prediction is dependent on the prompt choice, as we see from the open-ended prompts performing better than asking specific relationships. Data Programming is also dependent on high-quality sources of weak supervision, as we see from the improvement in performance when not considering the logits in Table 4. Effectively mapping the output of language models to the concrete label space without training remains a hard problem for future work to tackle.

Ethical Statement

Based on the methodology we have currently employed, we do not foresee any significant ethical concerns. All the documents and models utilized in our study were obtained from open-source domains, ensuring a transparent and accessible source of information. Additionally, PromptRE requires no LLM training, eliminating the risk of model drift. Additionally, the task of relation extraction is a widely recognized and well-studied problem across various natural language processing applications.

However, it is crucial to acknowledge a minor factor, namely the presence of potential hidden biases within the pretrained language models used in our analysis. These biases may stem from the data on which the models were trained, which could have inadvertently introduced implicit human biases. While our usage of these pretrained language models enables us to identify relationships between arbitrary entities, it is conceivable that biases may emerge if one were to explore sensitive relation classes and entities.

References

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Fenia Christopoulou, Makoto Miwa, and Sophia Ananiadou. 2019. Connecting the dots: Document-level

neural relation extraction with edge-oriented graphs. *arXiv preprint arXiv:1909.00228*.

- Bayu Distiawan, Gerhard Weikum, Jianzhong Qi, and Rui Zhang. 2019. Neural relation extraction for knowledge base enrichment. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 229–240.
- Zhijiang Guo, Yan Zhang, and Wei Lu. 2019. Attention guided graph convolutional networks for relation extraction. *arXiv preprint arXiv:1906.07510*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Yi Yao Huang and William Yang Wang. 2017. Deep residual learning for weakly-supervised relation extraction. *arXiv preprint arXiv:1707.08866*.
- Jing Jiang. 2009. Multi-task transfer learning for weakly-supervised relation extraction. *ACL*.
- D. Khashabi, S. Min, T. Khot, A. Sabharwal, O. Tafjord, P. Clark, and H. Hajishirzi. 2020. Unifiedqa: Crossing format boundaries with a single qa system.
- Daniel Khashabi, Yeganeh Kordi, and Hannaneh Hajishirzi. 2022. Unifiedqa-v2: Stronger generalization via broader cross-format training. *arXiv preprint arXiv:2202.12359*.
- Qi Li, Xuan Wang, Yu Zhang, Fei Ling, Cathy H Wu, and Jiawei Han. 2018. Pattern discovery for wide-window open information extraction in biomedical literature. In *2018 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 420–427. IEEE.
- Youmi Ma, An Wang, and Naoaki Okazaki. 2023. Dreeam: Guiding attention with evidence for improving document-level relation extraction. *arXiv preprint arXiv:2302.08675*.
- Guoshun Nan, Zhijiang Guo, Ivan Sekulić, and Wei Lu. 2020. Reasoning with latent structure refinement for document-level relation extraction. *arXiv preprint arXiv:2005.06312*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Van-Thuy Phi, Joan Santoso, Masashi Shimbo, and Yuji Matsumoto. 2018. Ranking-based automatic seed selection and noise reduction for weakly supervised relation extraction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 89–95.

- 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*, pages 1257–1266.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Alexander Ratner, Braden Hancock, Jared Dunmon, Frederic Sala, Shreyash Pandey, and Christopher Ré. 2019. Training complex models with multi-task weak supervision. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4763–4771.
- Alexander J Ratner, Christopher M De Sa, Sen Wu, Daniel Selsam, and Christopher Ré. 2016. Data programming: Creating large training sets, quickly. *Advances in neural information processing systems*, 29.
- Sunil Kumar Sahu, Fenia Christopoulou, Makoto Miwa, and Sophia Ananiadou. 2019. Inter-sentence relation extraction with document-level graph convolutional neural network. *arXiv preprint arXiv:1906.04684*.
- Oscar Sainz, Oier Lopez de Lacalle, Gorka Labaka, Ander Barrena, and Eneko Agirre. 2021. Label verbalization and entailment for effective zero- and few-shot relation extraction. *arXiv preprint arXiv:2109.03659*.
- Qingyu Tan, Ruidan He, Lidong Bing, and Hwee Tou Ng. 2022a. Document-level relation extraction with adaptive focal loss and knowledge distillation. *arXiv preprint arXiv:2203.10900*.
- Qingyu Tan, Lu Xu, Lidong Bing, Hwee Tou Ng, and Sharifah Mahani Aljunied. 2022b. Revisiting docred-addressing the false negative problem in relation extraction. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8472–8487.
- Hengzhu Tang, Yanan Cao, Zhenyu Zhang, Jiangxia Cao, Fang Fang, Shi Wang, and Pengfei Yin. 2020. Hin: Hierarchical inference network for document-level relation extraction. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 197–209. Springer.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- G Veena, S Athulya, Salma Shaji, and Deepa Gupta. 2017. A graph-based relation extraction method for question answering system. In *2017 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, pages 944–949. IEEE.
- Xuan Wang, Yu Zhang, Qi Li, Yinyin Chen, and Jiawei Han. 2018. Open information extraction with meta-pattern discovery in biomedical literature. In *Proceedings of the 2018 ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, pages 291–300.
- Ye Wang, Xinxin Liu, Wenxin Hu, and Tao Zhang. 2022a. A unified positive-unlabeled learning framework for document-level relation extraction with different levels of labeling. *arXiv preprint arXiv:2210.08709*.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Han-naneh Hajishirzi. 2022b. Self-instruct: Aligning language model with self generated instructions. *arXiv preprint arXiv:2212.10560*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Benfeng Xu, Quan Wang, Yajuan Lyu, Yong Zhu, and Zhendong Mao. 2021. Entity structure within and throughout: Modeling mention dependencies for document-level relation extraction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14149–14157.
- Hsiu-Wei Yang and Abhinav Agrawal. 2023. Extracting complex named entities in legal documents via weakly supervised object detection. *arXiv preprint arXiv:2305.05836*.
- Yuan Yao, Deming Ye, Peng Li, Xu Han, Yankai Lin, Zhenghao Liu, Zhiyuan Liu, Lixin Huang, Jie Zhou, and Maosong Sun. 2019. Docred: A large-scale document-level relation extraction dataset. *arXiv preprint arXiv:1906.06127*.
- Deming Ye, Yankai Lin, Jiaju Du, Zhenghao Liu, Peng Li, Maosong Sun, and Zhiyuan Liu. 2020. Coreferential reasoning learning for language representation. *arXiv preprint arXiv:2004.06870*.
- Wenhao Yu, Dan Iter, Shuohang Wang, Yichong Xu, Mingxuan Ju, Soumya Sanyal, Chenguang Zhu, Michael Zeng, and Meng Jiang. 2022. Generate rather than retrieve: Large language models are strong context generators. *arXiv preprint arXiv:2209.10063*.
- Shuang Zeng, Runxin Xu, Baobao Chang, and Lei Li. 2020. Double graph based reasoning for document-level relation extraction. *arXiv preprint arXiv:2009.13752*.

Wenxuan Zhou, Kevin Huang, Tengyu Ma, and Jing Huang. 2021. Document-level relation extraction with adaptive thresholding and localized context pooling. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 14612–14620.

A Parameter Settings

All models were run on an NVIDIA A6000 with 48 gigabytes of VRAM. Still, around 10 days were required to fully run the experiments. For particularly expensive computations, like *Logits_{SR}*, only the fastest model—UnifiedQA-large—could be feasibly run.

All models were downloaded from Huggingface (Wolf et al., 2019). We used the default setup of the pretrained models and did not do further finetuning. All the step mentioned in the methodology section works on the output of the pretrained models.

Supervised results DREEEAM (Ma et al., 2023) and KD-DocRE (Tan et al., 2022a) were taken from the original source papers.

B Evaluation Metrics

To keep in tradition with existing document relation extraction work, we report both F1 and Ign_F1 as computed by the official metrics from ReDocRED. F1 refers to micro-averaged F1 score that combines precision P and recall R

$$F1 = \frac{2PR}{P + R}$$

$$P = \frac{\text{length of correct (h,t,rel) preds}}{\text{length of all (h,t,rel) preds}}$$

$$R = \frac{\text{length of correct (h,t,rel) preds}}{\text{length of correct (h,t,rel)}}$$

Where (h, t, rel) denotes a tuple of the predicted head, tail, and relation. Ign_F1 is computed similarly to above but ignores the samples in the DocRED’s distantly supervised training set. (Note that we do not use any distantly labeled data).

C Effect of Relation-Specific Prompts

In this ablation study, we investigate the usefulness of *Logits_{SR}* on both the relation class prediction part as well as the addressing of the NA issue. Because each entity-entity pair has to prompt with all relations, it is quite expensive to perform. Thus, we only perform experiments with the fastest model we consider—UnifiedQA large. From Table 2, we see that by itself, using the logits do not perform as well as embedding similarity.

For the case of reducing NAs, we actually do not include it as a source of weak supervision in the data programming framework due to its inefficiency. However, if we *did* include it, we would see that performance drops as well, as shown in Table 6.

Table 6: Experimental results via prompting the model for each specific relation using the baseline model *Logits_{large}*.

	F1	Ign F1	Precision	Recall
Simple RE	5.5975	4.8830	3.4246	15.3147
MV	4.5163	4.3027	6.6494	3.4193
DP	6.2096	5.4219	3.9254	14.8519

Ross Patterson Alger (August 20 , 1920 - January 16 , 1992) was a politician in the Canadian province of Alberta , who served as mayor of Calgary from 1977 to 1980 . Born in Prelate , Saskatchewan , he moved to Alberta with his family in 1930s . He received a bachelor of commerce degree from the University of Alberta in 1942 . He served with the Royal Canadian Air Force during World War II . After the war , he received an MBA from the University of Toronto . He settled in Calgary and started a career in accounting . In 1958 , he was a public school board trustee , and later became the chairman . From 1971 to 1974 , he was an alderman on Calgary City Council . In 1974 , he ran for mayor losing to Rod Sykes . He was elected mayor in 1977 and served one term until 1980 . During Alger ’s term , notable accomplishments include the construction of the Ctrain ’s first leg , the bid for the XV Olympic Winter Games , and planning for the Olympic coliseum . His brother was Harry Alger . Alger died of cancer in 1992 , which had first been diagnosed in 1985 .

Figure 4: Original document for Case Study 1.

D Relation Distribution Calculation

We test our assumptions of the relation type distribution. Specifically, we how the performance changes with more or less expert annotated documents. The results are shown in Table 7. Recall that in total, we have 500 documents, so 1% of all documents represent only 5 annotated documents. This reinforces our assumption that creating this relation/type distribution is not exorbitantly expensive. Furthermore, this computation is only an *estimate* of the actual input that domain experts would provide, so it is possible that real world performance would be better or worse depending on the distribution of true types and relations.

Table 7: The Performance of UnifiedQA-large on the varying percentages of the data we use for to compute the expert-provided relation/type distribution.

	1%	F1	Ign F1	Precision	Recall
Embed _{large} Sim. + Simple RE		6.8663	5.8224	4.2008	18.7862
Embed _{large} Sim. + DP		7.8207	6.5522	4.9439	18.7052
10%					
Embed _{large} Sim. + Simple RE		8.6363	7.4106	5.2837	23.6288
Embed _{large} Sim. + DP		9.6495	8.1978	6.0999	23.0791
25%					
Embed _{large} Sim. + Simple RE		8.9535	7.7211	5.4778	24.4966
Embed _{large} Sim. + DP		9.8986	8.4370	6.2575	23.6751
50%					
Embed _{large} Sim. + Simple RE		9.1121	7.8481	5.5748	24.9306
Embed _{large} Sim. + DP		9.9107	8.4176	6.2651	23.7040
Original (100%)					
Embed _{large} Sim. + Simple RE		9.2030	7.9314	5.6304	25.1794
Embed _{large} Sim. + DP		10.2232	8.7384	6.3969	25.4397

Mess of Blues is an album by Jeff Healey . It was released in 2008 less than two weeks after his death and just three weeks shy of his 42nd birthday . Four of the album 's tracks were recorded live in front of audiences , two of the live tracks at the Islington Academy in London , and the other two live tracks at Jeff Healey 's Roadhouse in Toronto . The other six tracks were recorded at Studio 92 in Canada by Norm Barker and Richard Uglow . The whole album features the band which normally accompanied Jeff at his club , Jeff Healey 's Roadhouse . The song " Mess of Blues " , which appears on the album was written by Doc Pomus and Mort Shuman and was originally recorded by Elvis Presley .

Figure 5: Original document for Case Study 2.

E Additional Case Studies

We analyze a few more examples of the model output in Table 8 with the original document in Figure 5. While the first 3 labeled examples make intuitive sense and are correct, the lack of label completeness is evident in the latter 3 examples. For example, in the last example, Richard Uglow is the recorder of part of the album, but the pair "Richard Uglow" and "Mess of Blues" is marked as "No relation" in the dataset. This type of weakly labeled situation is a particularly difficult problem for weakly-supervised methods like ours because, without training, it's very hard to inform the model with subtle details regarding the relation label space.

Table 8: An example of a relation prediction of a portion of labeled positive relation tuples. Each relation tuple has the format: (Head entity, Tail entity, Ground Truth Relation). The predicted text is shown below after prompting the model using an open-ended prompt (e.g. "Is there a relationship between <Head> and <Tail>?").

Labeled Relation Tuple: (Mess of Blues; 2008; publication date)

The album Mess of Blues was released in 2008, less than two weeks after Jeff Healey's death and just three weeks shy of his 42nd birthday. It features four live tracks recorded in London and Toronto, and six studio tracks recorded in Canada.

Labeled Relation Tuple: (Mess of Blues; Doc Pomus; lyrics by)

The song "Mess of Blues" was originally recorded by Elvis Presley and appears on Jeff Healey's posthumous album. Doc Pomus is mentioned as a songwriter who co-wrote the song with Mort Shuman.

Labeled Relation Tuple: (Elvis Presley; Mess of Blues; notable work)

song "Mess of Blues" on the album was originally recorded by Elvis Presley.

Unlabeled Relation Tuple: (Mess of Blues; two weeks; no relation)

The album Mess of Blues by Jeff Healey was released two weeks after his death. He died just three weeks shy of his 42nd birthday.

Unlabeled Relation Tuple: (Mess of Blues; Islington Academy; no relation)

The album "Mess of Blues" was recorded at the Islington Academy in London.

Unlabeled Relation Tuple: (Mess of Blues; Richard Uglow; no relation)

Yes, there is a relationship between "Mess of Blues" and "Richard Uglow". Richard Uglow recorded six tracks of Jeff Healey's album "Mess of Blues" at Studio 92 in Canada along with Norm Barker.
