

Revisiting DocRED – Addressing the False Negative Problem in Relation Extraction

Qingyu Tan*^{1,2} Lu Xu*^{1,3} Lidong Bing^{†1} Hwee Tou Ng² Sharifah Mahani Aljunied¹

¹DAMO Academy, Alibaba Group

²Department of Computer Science, National University of Singapore

³Singapore University of Technology and Design

{qingyu.tan, lu.x, l.bing, mahani.aljunied}@alibaba-inc.com

{qtan6, nght}@comp.nus.edu.sg

Abstract

The DocRED dataset is one of the most popular and widely used benchmarks for document-level relation extraction (RE). It adopts a recommend-revise annotation scheme so as to have a large-scale annotated dataset. However, we find that the annotation of DocRED is incomplete, i.e., false negative samples are prevalent. We analyze the causes and effects of the overwhelming false negative problem in the DocRED dataset. To address the shortcoming, we re-annotate 4,053 documents in the DocRED dataset by adding the missed relation triples back to the original DocRED. We name our revised DocRED dataset Re-DocRED. We conduct extensive experiments with state-of-the-art neural models on both datasets, and the experimental results show that the models trained and evaluated on our Re-DocRED achieve performance improvements of around 13 F1 points. Moreover, we conduct a comprehensive analysis to identify the potential areas for further improvement.¹

1 Introduction

The field of relation extraction (RE) is related to knowledge bases (KBs). Most popular relation extraction datasets are constructed from knowledge triples in KBs. For example, the TACRED dataset (Zhang et al., 2017) is constructed by the TAC Knowledge Base Population challenge. The NYT10 (Riedel et al., 2013) dataset matched the Freebase Knowledge base (Bollacker et al., 2008) to the New York Times corpus (Sandhaus, 2008). Wiki20 (Gao et al., 2021) and DocRED (Yao et al., 2019) originated from distant supervision from the Wikidata knowledge base (Vrandečić and Krötzsch, 2014) and Wikipedia articles. By exploiting distant

* Equal contribution. Qingyu Tan and Lu Xu are under the Joint PhD Program between Alibaba and NUS/SUTD.

[†] Corresponding author.

¹Our dataset is publicly available at <https://github.com/tonytan48/Re-DocRED>.

"I Knew You Were Trouble" is a song recorded by American singer - songwriter Taylor Swift for her fourth studio album, *Red* (2012). It was released on **October 9, 2012**, in the **United States** by **Big Machine Records** as the third promotional single from the album. Later, "**I Knew You Were Trouble**" was released as the third single from *Red* on **November 27, 2012**, in the **United States**. It was written by Swift, **Max Martin** and **Shellback**, with the production handled by the latter two ...

It later peaked at number two in **January 2013**, blocked from the top spot by **Bruno Mars'** "**Locked Out of Heaven**". At the inaugural **YouTube Music Awards** in 2013, "**I Knew You Were Trouble**" won the award for **YouTube** phenomenon ...

DocRED: (I Knew You Were Trouble, producer, Max Martin); (Taylor Swift, country of citizenship, the United States) ...

Re-DocRED: (I Knew You Were Trouble, producer, Max Martin); (I Knew You Were Trouble, producer, Shellback) ...

Figure 1: A sample document in our Re-DocRED dataset. From the highlighted evidence, it can be inferred that "I Knew You Were Trouble" was produced by both "Max Martin" and "Shellback", but in the previous incomplete DocRED dataset, only the first producer is labeled.

supervision, relation triple candidates can be retrieved from the knowledge bases for a given piece of text. The retrieved triples are based on entity co-occurrence only, and they may not be related to the context. This is known as the false positive (FP) problem (Bing et al., 2015; Xiao et al., 2020).

However, the false negative (FN) problem in the relation extraction datasets receives less attention as compared to FPs. Without resolving this issue, the annotations of the datasets are incomplete. Recent efforts on addressing the false negative problem are from the model perspective (Chen et al., 2021; Hao et al., 2021), which aim to denoise the false negative data during training. The challenge for these approaches is that both devel-

opment and test sets can be incomplete at the same time. Without a completely annotated dataset, the current evaluation is ill-defined.

Several re-annotation works revised the existing sentence-level relation extraction datasets. [Alt et al. \(2020\)](#) re-annotated a small amount of challenging samples in the development and test sets of the TACRED dataset ([Zhang et al., 2017](#)). [Stoica et al. \(2021\)](#) extended this work and re-annotated the training, development, and test sets of TACRED with a semantically refined label space. Besides, [Gao et al. \(2021\)](#) re-annotated the test sets of two distantly supervised RE datasets. Even though their discussions and examples emphasized the efforts on correcting the false positive samples, all the revised versions of the datasets have a significant increase of positive samples. That is, many samples that are previously labeled as *no_relation* (NA) are re-annotated with relation labels in the revised datasets. We show the detailed statistics in [Table 1](#), which indicates that the false negative problem is prevalent in the sentence-level relation extraction field. Compared to the sentence-level task, document-level RE is more susceptible to the false negative problem. This is primarily because document-level RE involves significantly more entity pairs in a raw text, as shown in [Table 1](#). Note that the objective of the RE task is to determine the relation types for all entity pairs, and the number of entity pairs is quadratic in the number of entities.

In this paper, we address the false negative problem in DocRED. We find that the false negative problem arises due to two reasons. First, although the Wikidata knowledge base provides a good starting point for distantly supervised annotation of DocRED, it is highly sparse and far from complete. There are many relation triples that are not in the Wikidata KB. For example, in [Figure 1](#), the article reflects that the song “I Knew You Were Trouble” was produced by both “Max Martin” and “Shellback”, but the fact that “Shellback” is the other producer is not included in both the knowledge base and the DocRED dataset. Second, the additional relation triples recommended by the RE model and human annotators cannot cover the ground-truth relation triples that are missed from the Wikidata. A detailed discussion is given in [Section 2](#). With the incompletely annotated development and test sets, the previous evaluation does not necessarily provide a fair reference. Therefore, we propose to revise the DocRED dataset to recover the in-

complete annotations through an iterative approach with human in the loop. Specifically, we use multiple state-of-the-art document-level RE models to generate relation candidates and ask human annotators to examine the recommended triples. The details of our iterative human-in-the-loop annotation process are given in [Section 3](#).

Most recently, [Huang et al. \(2022\)](#) also identify the false negative issue in DocRED ([Yao et al., 2019](#)). They combat the problem by annotating 96 documents from scratch with two expert annotators. However, annotating relation triples from scratch is different from revising recommended triples, and it is difficult to scale up to a dataset of a larger size. We provide a comprehensive analysis between our approach and the annotating-from-scratch approach in [Appendix A](#). Compared to [Huang et al. \(2022\)](#), our approach is better in the following aspects. First, our dataset is significantly larger in size (4,053 vs 96 documents). Second, the precision of our annotation is higher. Third, our evaluation dataset contains more triples per document than [Huang et al. \(2022\)](#), indicating that our dataset better tackles the incompleteness problem of DocRED. Fourth, our dataset annotation approach is more scalable and can be extended to an arbitrary number of relation types.

Our contributions can be summarized as follows:

- We identify the overwhelming false negative problem in relation extraction.
- We show that the false negative problem is the cause of performance bottleneck on many relation extraction datasets, and we also provide a high-quality revised version of the document-level relation extraction dataset Re-DocRED.
- Moreover, we conduct comprehensive experimental analysis to identify the potential areas for further improvement.

2 Preliminaries

2.1 Background of the DocRED Dataset

The DocRED dataset ([Yao et al., 2019](#)) is one of the popular and widely-studied benchmark datasets for document-level relation extraction. The dataset contains 5,053 Wikipedia documents, where each document has an average length of 196.7 words, and an average of 19.5 entities. With 97 predefined relation types (including *no_relation*) and an average of 393.6 entity pairs, there exist around

Dataset	# Relations	# Triples	# Samples	Avg. # Words	Avg. # Entities	Avg. # Entity Pairs	NA
Sentence-level							
Wiki20 (Han et al., 2020)	454	15,026	28,897	25.7	2	1	48.0%
Wiki20m (Gao et al., 2021)	81	103,709	137,986	25.6	2	1	24.8%↓
TACRED (Zhang et al., 2017)	42	21,373	106,264	36.4	2	1	79.9%
Re-TACRED (Stoica et al., 2021)	40	33,690	91,467	36.4	2	1	63.2%↓
Document-level							
HacRED (Cheng et al., 2021)	26	56,798	7,731 ¹	122.6	10.7	163.9	95.7%
DocRED (Yao et al., 2019)	96	50,503	4,053 ¹	198.4	19.5	393.6	97.1%
Re-DocRED (Ours)	96	120,664	4,053	198.4	19.4 ²	391.0	94.0%↓

¹ We do not include the blind test sets for these two datasets.

² We also resolve the coreferential errors. Details are given in Appendix B.

Table 1: Statistics of relation extraction datasets. We can see that all re-annotation efforts on sentence-level relation extraction datasets (Wiki20 → Wiki20m, TACRED → Re-TACRED) result in significantly smaller percentages of *no_relation* (NA) samples, indicating that false negative is a common problem in the RE datasets.

38,000 relation triple candidates per document to be annotated. To reduce the annotation workload, a recommend-revise scheme is adopted. Specifically, relation triple candidates are recommended from distantly supervised data and predictions of an RE model. On average for each document, 19.9 triples are suggested from distantly supervised data and 7.8 triples are recommended by relation extraction models. Human annotators are asked to read the documents and review the relation triple candidates, and wrong recommendations will be discarded. As a result, an average of 12.5 triples per document are kept in the DocRED dataset.

2.2 Problem of Incomplete Annotation

Based on our empirical analysis of the top-performing models (Tan et al., 2022; Zhang et al., 2021; Zhou et al., 2021) on the DocRED leaderboard, many predicted triples are correct but are not annotated in the DocRED dataset. Therefore, it is important to review the dataset and identify the true bottleneck of document-level RE. We identify incomplete annotation as the major issue in the DocRED dataset. As mentioned in Section 2.1, it is difficult to annotate a document-level RE dataset from scratch. The DocRED dataset is mainly created by human filtering the recommended relation triples in distantly supervised data and predictions from an RE model. It is worth noting that the construction procedure of DocRED relies on the underlying assumption that the combination of the recommended triples by distant supervision and the RE model contains almost all the ground-truth relation triples in the documents. This assumption is not true since the relations in the Wikidata knowledge base (KB) are sparse and incomplete. That is, many semantically related triples are not reflected in the KB. Therefore, treating such triples

as negative instances (*no_relation*) will introduce false negative examples into the dataset. In the original DocRED production process, only the distantly supervised data is used to train the RE model, which may lead to low-quality prediction. Furthermore, the performance of the previously used RE model is significantly worse than the recent approaches based on pre-trained language models. Therefore, the recommended relation triples from the RE model are likely not enough for covering the ground-truth relation triples for a given text.

By relying heavily on the relation triple candidates generated from the above two methods, the annotation of the previous DocRED dataset is incomplete. While the recommend-revise scheme is the major source of the incompleteness of the original DocRED dataset (Yao et al., 2019), a secondary source of incompleteness comes from logical inconsistency. There are many inverse relation pairs in DocRED. For example, if entity A is annotated as a *sibling* of entity B, then entity B is also a *sibling* of entity A. The lack of inverse relation also contributes to the incompleteness problem.

2.3 The Impact of False Negatives

To identify the difficulty of relation extraction, we conduct a preliminary analysis on two evaluation tasks: Relation Extraction (RE) and Positive Relation Classification (PRC).

Settings of Our Analysis Consider a text T and a set of n entities $\{e_1, \dots, e_n\}$. e_i and e_j denote two different entities, and R is a set of predefined relations, including *no_relation*. The objective of the RE task is to identify the relation type in R for each entity pair (e_i, e_j) . Under the PRC setting, we do not use any entity pairs of *no_relation*, and the model is trained and evaluated with only en-

Dataset	P	R	F1
Sentence-level			
TACRED	93.38	93.38	93.38
Re-TACRED	96.83	96.83	96.83
Incomplete Re-TACRED	96.53	96.53	96.53
Document-level			
DocRED	93.33	90.28	91.78
HacRED	95.18	94.95	95.07

Table 2: Preliminary experimental results of PRC.

tity pairs that have some pre-defined relation types. Such a setting allows us to assess the difficulty of differentiating the predefined relation types. We use the top-performing models for our preliminary experiments, and we report the scores on the development sets in these experiments. For the sentence-level task, we use the typed entity marker (Zhou and Chen, 2021) with RoBERTa-large (Liu et al., 2019) as our baseline. For the document-level task, we use KD-DocRE (Tan et al., 2022) as the baseline. We use RoBERTa-large as the encoder for DocRED and XLM-R-base (Conneau et al., 2020) for HacRED.

Datasets We compare the preliminary investigation results on two sentence-level RE datasets: TACRED (Zhang et al., 2017) and Re-TACRED (Stolica et al., 2021), and two document-level RE datasets: DocRED (Yao et al., 2019) and HacRED (Cheng et al., 2021). Table 1 shows the statistics of the datasets. Re-TACRED revised the TACRED dataset to reduce annotation errors. Besides, we also create an incomplete Re-TACRED dataset, where 30% of the positive labels are replaced with *no_relation* in the training, development, and test sets. This modification artificially creates false negative samples. This is to show the detrimental effect of the incompleteness problem on downstream tasks. DocRED and HacRED are document-level relation extraction datasets in English and Chinese, respectively. For sentence-level relation extraction datasets, the negative ratio is typically lower than that of the document-level datasets, mainly because the complexity of relation extraction is quadratic in the number of entities. Therefore, document-level RE is more prone to the false negative problem.

2.3.1 Preliminary Results of Positive Relation Classification

By assuming that there is a relation between the entity pairs, the positive relation classification (PRC)

Dataset	P	R	F1
Sentence-level			
TACRED	75.70	73.40	74.50
Re-TACRED	90.60	91.30	90.90
Incomplete Re-TACRED	65.61	71.71	68.52
Document-level			
DocRED	64.62	63.53	64.07
HacRED	78.45	77.93	78.19

Table 3: Preliminary experimental results of RE.

task shows the difficulty of differentiating the relation types. Note that document-level PRC is a multi-label classification problem. Hence, precision and recall are not necessarily the same. From Table 2, we can see that all the models perform well. The performance of sentence-level RE and document-level RE are comparable, even though document-level RE has a significantly longer context and requires cross-sentence reasoning. This shows that the difficulty of positive relation classification is not severely affected by sentence boundary or context length. Another finding from Table 2 is that the revised version of TACRED has marginally higher performance than the original version. This is expected as the revised version receives an extra round of human annotation. The performance on Incomplete Re-TACRED is only marginally worse than Re-TACRED, which shows that positive relation classification can achieve a comparatively high performance despite the reduction of training instances. Besides, it is worth noting that the performance on HacRED is higher than the performance on DocRED, even though HacRED claims that it is a semantically harder dataset. Although the baselines are not exactly the same, we can still infer that the difficulty level of classifying the positive relation types on sentence-level and document-level datasets is not significantly different.

2.3.2 Preliminary Results of Relation Extraction

Compared to the setting of positive relation classification, the standard relation extraction task includes all the negative *no_relation* samples during training. We compare the performance of the previous best approaches on the sentence-level and document-level RE datasets in Table 3. We observe that the performance on the standard RE task is lower than that on the PRC task. For a spe-

cific dataset, the performance on PRC is the upper bound of RE performance, since the evaluation of PRC ignores *no_relation*. However, the performance between relation extraction and positive relation classification should not have a large gap, as the positive samples should be semantically different from the negative samples for a well-annotated dataset. For the sentence-level dataset, we can see that the performance on the revised version of TACRED (Re-TACRED) is significantly higher than that of the original TACRED dataset. The performance difference between the positive relation classification and relation extraction is 18.88 (93.38 vs. 74.50) F1 score for TACRED, whereas the gap is only 5.93 (96.83 vs. 90.90) F1 score for Re-TACRED. For the incomplete Re-TACRED, this gap becomes 28.01 (96.53 vs. 68.52). Moreover, from the comparison of Re-TACRED and incomplete Re-TACRED, we can see that the performance of PRC only dropped by 0.3 (96.83 vs. 96.53), whereas the performance of RE dropped by 22.38 (90.90 vs. 68.52). This shows that the incompleteness problem greatly decreases RE performance even if the positive instances are precisely labeled, whereas the performance of PRC is not significantly affected by the incompleteness problem. For the document-level datasets, we observe large gaps between positive relation classification and relation extraction, 27.71 (91.78 vs. 64.07) F1 for DocRED and 16.88 (95.07 vs. 78.19) for HacRED, respectively. Compared to the 5.93 F1 gap on the Re-TACRED dataset, the gaps of PRC and RE on the document-level datasets are significantly larger. These observations indicate that there may exist a substantial number of false negative examples throughout the training, development, and test sets of these datasets. Therefore, it is necessary to re-check the quality of the document-level relation extraction datasets.

3 Revising DocRED with an Iterative Approach

In this section, we describe our iterative human-in-the-loop approach to revise DocRED. The goal is to add the previously unrecognized relation triples to the DocRED dataset. Our iterative approach consists of three steps in each iteration: (1) Training scorer models; (2) Scoring all possible triples; and (3) Human verification. The following sections provide the details of each step.

Split 0	Split 1	Split 2	Split 3
1,000	1,000	1,000	1,053

Table 4: Dataset statistics of each split. Split 0 is the original DocRED development set.

3.1 Our Iterative Approach

Step 1 - Training Scorer Models Even though the annotation of the previous DocRED is incomplete, we can still train neural models on such data. We adopt three top-performing models on the current DocRED leaderboard: KD-DocRE (Tan et al., 2022), DocuNET (Zhang et al., 2021), and AT-LOP (Zhou et al., 2021). To obtain relation triple candidates for all 4,053 documents, we split the original DocRED into 4 different splits (Table 4), with the first split as the original DocRED development set. This is to ensure that the number of training samples is comparable to the original DocRED when any three of the splits are used for training our scorer models, and the remaining split is used for prediction. Following the training paradigm of Tan et al. (2022), we first pre-train each model with the distantly supervised data in DocRED. Then, the pre-trained model is further trained on any three splits of DocRED, and we then use the trained model to make predictions on the remaining split. Therefore, it requires four rounds of training and inference so that we can get the predictions for all the four splits.

Step 2 - Scoring the Triples In this step, we aim to generate a large number of relation triple candidates, so that they could cover the missing annotations in the previous DocRED. With the trained scorer models, we can predict the scores for all the enumerated relation triples. To control the number of relation triple candidates for the next step, we define a threshold score to remove the less confident predictions. Specifically, we set the dynamic threshold to 0.9 of the models for the Adaptive Threshold class (Zhou et al., 2021) for all trained models. The predicted relation triples from all the models are then merged together. Due to the different characteristics of these models, we could generate a large and diverse pool of relation triple candidates for the next step. By default, the original positive triples from DocRED are treated as correct and are not re-annotated.

Dataset	Re-DocRED			DocRED	
	Train	Dev	Test	Train	Dev
# Documents	3,053	500	500	3,053	1,000
Avg. # Entities	19.4	19.4	19.6	19.5	19.6
Avg. # Triples	28.1	34.6	34.9	12.5	12.3
Avg. # Sentences	7.9	8.2	7.9	7.9	8.1

Table 5: Statistics of Re-DocRED and DocRED.

Step 3 - Human Verification After the relation triple candidates are generated from the previous step, each triple candidate will be annotated by humans. The human annotators are asked to read the document and check whether the triples can be inferred from the document. Each triple will be annotated by two annotators, and a third annotator will resolve the conflicting annotations.

3.2 Our Revised Re-DocRED

The above three steps form one round of our iterative approach. We conducted two rounds of annotation in total. For the first round, we annotated 4,053 documents that include all training and evaluation documents. On average, we recommended 11.9 triples for each document and 9.4 triples were accepted, with an acceptance rate of 79.0%. The Fleiss Kappa (Fleiss, 1971) coefficient for round 1 annotation is 0.73, which is considered as substantial agreement. To further improve the recall on the evaluation dataset, we conducted a second round of annotation for the 1,000 evaluation documents. We used the annotated 3,053 training samples from round 1 for round 2 training. In this round, 14.1 triples were recommended and only 6.0 triples were accepted, with an acceptance rate of only 42.6%. The Fleiss Kappa coefficient for round 2 annotation is 0.66. After human annotation, we also add relation triples by manually defining logical rules. In this way, we are able to resolve the problem of logical inconsistency (described in Section 2.2) in the DocRED dataset. These rules mainly consist of inverse relations and co-occurring relations. See Appendix E for more details. Overall, our training documents contain 28.1 triples on average, with 9.4 triples added from human annotation and 6.2 triples from logical rules. Our evaluation documents contain 34.7 triples on average, with 15.7 triples added from human annotation and 6.7 triples from logical rules. We divide the 1,000 evaluation documents into 500 development and 500 test documents. The average number of triples of

the evaluation documents is higher than that of the training documents. This indicates that the evaluation data has more complete annotation compared to the training data. The detailed statistics of the Re-DocRED dataset are shown in Table 5. The average number of triples per document is significantly higher for Re-DocRED compared to the original DocRED. There are 12.3 triples per document in the DocRED dev set and 34.7 triples per document in the dev and test sets of Re-DocRED. This shows that approximately 64.6% of all triples are missing in the original DocRED dataset.

4 Experiments

4.1 Comparison on Relation Extraction

To compare the previous DocRED and our Re-DocRED, we evaluate 4 different approaches on the two datasets. Apart from the three models that are used during our annotation process (Section 3.1), we also compare the performance with an additional approach, JEREX (Eberts and Ulges, 2021). This approach is not included in our data production process and we use it as an independent model to compare DocRED and Re-DocRED. Table 6 shows the experimental results, and the reported metrics are micro-averaged F1 scores and Ign_F1 scores. The latter refers to the F1 score that ignores the triples appearing in the training set. According to the statistics in Table 1, even though our revised Re-DocRED dataset contains many more relation triples, we observe that all the baseline models demonstrate significant performance improvement on both development and test sets. Compared to DocRED, the performance of the baseline models on Re-DocRED increased by more than 12 F1 points. When these models are pre-trained with distantly supervised data², we observe consistent performance improvement on Re-DocRED. We also further analyze the error cases of the SOTA model in Appendix D.

4.2 Comparison on Positive Relation Classification

Following the experimental setting in Section 2.3.1, we compare the positive relation classification performance between the original DocRED and our Re-DocRED with KD-DocRE (Tan et al., 2022) in Table 7. The performance on Re-DocRED is slightly better than that on the original DocRED.

²There are 101K distantly supervised documents in the original DocRED dataset.

Model	DocRED				Re-DocRED				Test Differences	
	Dev		Test		Dev		Test		Δ Ign_F1	Δ F1
	Ign_F1	F1	Ign_F1	F1	Ign_F1	F1	Ign_F1	F1		
JEREX	57.07	58.97	57.14	59.01	71.59	72.68	71.45	72.57	+14.31	+13.56
ATLOP	61.21	63.07	61.44	63.20	76.79	77.46	76.82	77.56	+15.38	+14.36
DocuNET	61.62	63.53	61.80	63.64	77.49	78.14	77.26	77.87	+15.46	+14.23
KD-DocRE	62.08	64.07	62.21	64.07	77.85	78.51	77.60	78.28	+15.39	+14.21
+ Pre-trained with distantly supervised data										
JEREX	60.39	62.24	60.29	62.15	73.34	74.77	73.48	74.79	+13.19	+12.64
ATLOP	63.68	65.61	63.63	65.51	78.32	79.26	78.52	79.46	+14.89	+13.95
DocuNET	63.22	65.25	63.23	65.26	78.20	78.90	78.28	78.99	+15.06	+13.73
KD-DocRE	65.11	67.04	65.21	67.09	79.79	80.56	80.32	81.04	+15.11	+13.95

Table 6: Experimental results using the original DocRED and our Re-DocRED. The reported numbers are average scores over 3 runs. The reported DocRED results use the same splits of development and test sets as Re-DocRED.

Training	Test	P	R	F1
DocRED	DocRED	93.33	90.28	91.78
Re-DocRED	Re-DocRED	95.03	90.55	92.74

Table 7: Positive relation classification performance with the original DocRED and our revised Re-DocRED (using KD-DocRE).

This indicates that our added triples are of comparable quality to the original DocRED data.

4.3 More Analysis

Additional Evaluation Metrics As mentioned in Section 2, document-level relation extraction is a challenging task. Hence, it is necessary to have various performance evaluation metrics so as to conduct a comprehensive evaluation. On top of the standard F1 and Ign_F1 evaluation metrics, we use four additional metrics to assess the models. (1) **Freq. F1** (Tan et al., 2022), which only considers the 10 most common relation types in the training set of Re-DocRED, where these frequent relation types account for around 60% of the relation triples. (2) **LT F1**, which only considers the long-tail (the remaining 86) relation types. (3) **Intra F1** (Nan et al., 2020), which evaluates on relation triples that appear in the same sentence. (4) **Inter F1**, which evaluates on cross-sentence relation triples.

We show the comprehensive evaluation results in Table 8. We observe that there exists a relatively large gap between the **Freq. F1** and **LT F1** metrics, and the difference is around 6 to 10 F1 points. Such behavior shows that the frequent relation types are easier to be recognized compared to the long-tail relation types. Furthermore, we also

find that the performance on triples that appear in the same sentence (**Intra F1**) is better than that on the cross-sentence relation triples (**Inter F1**), by around 2 to 6 F1 points. This is because it is harder to encode long-distance interactions. Therefore, future research can work on matching the performance of the long-tail relation types to the frequent types and also improve the model’s representation capability to capture inter-sentence interactions.

Training with DocRED vs. Re-DocRED We compare the performance of the KD-DocRE model when it is trained on the training sets of DocRED and Re-DocRED, and the results are shown in the upper section of Table 9. Note that for reliability, the evaluations are conducted on the development and test sets of Re-DocRED. Under the default setting (using all the positive and negative samples during training), the model trained with the Re-DocRED training data achieves a comparable precision score as DocRED (89.76 vs. 92.08). However, the recall score of the model trained with Re-DocRED is much better than DocRED (69.40 vs. 32.07). This demonstrates the severe incompleteness of the previous DocRED’s annotation, and models trained with such data fail to capture many triples. After adding back the missed samples, Re-DocRED’s training set is more complete and thus the trained model tends to extract more valid triples. Figure 2 shows an example document with predictions of models trained on the two training sets. We can see that the model trained with Re-DocRED indeed outputs more correct triples. On the other hand, for Re-DocRED, the recall score is still lower than the corresponding precision score. One possible reason is that the evaluation sets of Re-DocRED

Model	Dev		Test					
	Ign_F1	F1	Ign_F1	F1	Freq. F1	LT F1	Intra F1	Inter F1
JEREX	71.59	72.68	71.45	72.57	77.09	66.31	76.10	69.88
ATLOP	76.79	77.46	76.82	77.56	80.78	72.29	80.11	74.92
DocuNET	77.49	78.14	77.26	77.87	81.25	73.32	79.89	76.58
KD-DocRE	77.85	78.51	77.60	78.28	80.85	74.31	79.52	77.18
+ Pre-trained with distantly supervised data								
JEREX	73.34	74.77	73.48	74.79	78.67	68.62	76.90	72.65
ATLOP	78.32	79.26	78.52	79.46	82.23	75.17	80.39	78.44
DocuNET	78.20	78.90	78.28	78.99	82.08	74.19	80.46	77.72
KD-DocRE	79.79	80.56	80.32	81.04	83.17	79.31	82.01	79.87

Table 8: Performance comparison under different metrics on the Re-DocRED dataset.

Training	Negative Sample	P	R	F1
DocRED	All	92.08	32.07	47.57
Re-DocRED	All	89.76	69.40	78.28
DocRED	Partial (10%)	77.96	51.97	62.37
Re-DocRED	Partial (70%)	83.58	75.06	79.09

Table 9: Experimental results when using different numbers of negative samples.

received two rounds of annotation while the training set only received one round. Therefore, the annotation of the evaluation sets of Re-DocRED is more complete compared to the training set.

To further compare the annotation quality, we conduct experiments to randomly sample a subset of the negative samples instead of using all of them. This training strategy is known as negative sampling, and it is shown to be effective in resolving the false negative problem (Li et al., 2021) during training for named entity recognition. The experimental results on negative sampling are given in the lower section of Table 9. When random down-sampling on negative entity pairs is adopted, the experimental results are improved, due to significantly improved recall. We also observe that the improvement is more significant on DocRED. This demonstrates that the incompleteness problem of our revised training data is significantly alleviated. The sampling rate is a hyper-parameter obtained by grid search, and the details of negative sampling are described in Appendix C.

5 Related Work

Relation Extraction There is a series of relation extraction datasets built over the past decades, and they have significantly advanced research on RE. ACE 2005 (Walker et al., 2006) and SemEval

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**...

Subject Entity: Ross Patterson Alger

DocRED: (place of birth, **Prelate**), (date of birth, **August 20 , 1920**), (educated at, **University of Alberta**)...

Re-DocRED: (place of birth, **Prelate**), (date of birth, **August 20 , 1920**), (educated at, **University of Alberta**), (educated at, **University of Toronto**), (military branch, **Royal Canadian Air Force**)...

Figure 2: Sample predictions of the model when it is trained on DocRED and Re-DocRED training sets, and the subject entity of the selected triples is “Ross Patterson Alger”. The highlighted evidence verifies that the model trained on Re-DocRED has captured more correct inter-sentence triples (colored in green) than DocRED.

2010 Task 8 (Hendrickx et al., 2010) created two sentence-level RE datasets by human annotation. However, these two datasets have a relatively small number of relation types and instances. The New York Times corpus (Sandhaus, 2008) is another common relation extraction dataset used in the literature (Riedel et al., 2013; Nayak and Ng, 2019). It has been used for joint entity and relation extraction by (Nayak and Ng, 2020). The large-scale TACRED (Zhang et al., 2017) dataset is created based on the 2009–2014 TAC knowledge base population (KBP) challenges and crowd-sourced human annotations. FewRel (Han et al., 2018) and FewRel 2.0 (Gao et al., 2019) have been proposed to study the transferability and few-shot capability of RE

models. However, early relation extraction datasets mainly focus on sentence-level RE, whereas many relations can only be expressed by multiple sentences. The document-level relation extraction task has been proposed to build RE systems that are able to extract relations from multiple entities and sentences. Yao et al. (2019) have created the DocRED dataset by distant supervision from Wikipedia articles and the Wikidata knowledge base. They sampled 5,053 documents for human annotation. The annotation strategy of DocRED is mainly based on machine recommendation and human filtering. With a similar approach, Cheng et al. (2021) have created a Chinese document-level RE dataset that focuses on hard relation cases.

Machine-Assisted Data Generation Since labeled data is expensive to obtain for complex NLP tasks, there are many research works on generating automatically labeled data. Distant supervision is first used by Mintz et al. (2009) to generate relation extraction data without human efforts. Prior work on automatic data generation mainly relies on rule-based pattern matching (Lehmann et al., 2015) and web crawling (Buck et al., 2014). These types of rule-based methods are susceptible to noise propagation. With the rapid development of pre-trained language models (PLMs; Devlin et al., 2019; Liu et al., 2019; Brown et al., 2020), many recent works leverage PLMs for automatic data generation (Bing et al., 2013; Anaby-Tavor et al., 2020; Zhou et al., 2022; Yang et al., 2020; Kumar et al., 2020; Liu et al., 2021). However, these methods typically depend on a certain set of supervised data. Another line of work utilizes manually designed prompts and instructions to generate data in an unsupervised manner (Wang et al., 2019; Schick and Schütze, 2021; Chia et al., 2022). Although these methods improve the performance on some downstream tasks, the quality of the machine-generated data still does not match human annotation. To mitigate noise from the machine-generated data, West et al. (2021) generate a large amount of commonsense knowledge data and employ human annotators to filter the generated candidates.

6 Conclusions

In conclusion, this paper identifies the causes and effects of the overwhelming false negative problem in relation extraction. We show that the false negative problem is the cause of the performance bottleneck on many RE datasets. We have also provided

a high-quality revised version of document-level RE dataset, namely Re-DocRED. Moreover, we have provided comprehensive experimental analysis for Re-DocRED and identified the potential areas for further improvement. We have also conducted a thorough error analysis on state-of-the-art RE models.

7 Acknowledgements

We would like to thank all our annotators for their efforts, and the anonymous reviewers for their insightful comments and feedback.

8 Limitations

In this section, we discuss the limitations of our work. We have revised a popular document-level relation extraction dataset – the DocRED dataset. We use an iterative human-in-the-loop approach for revising this dataset. However, after two rounds of annotation of the evaluation data, there might still be missing annotations, even though we have added almost twice the number of triples compared to the original DocRED dataset. That is, our re-annotation serves as a reasonable approximation to the ground truth, although it may still not be the gold standard for an ideal document-level RE scenario. We have provided a detailed analysis on annotation quality and comparison with contemporaneous work in Appendix A. The second limitation of our work is that human revision on the training data is only carried out in one round. The completeness level of the training data is lower compared to our evaluation data. Future research can focus on advanced algorithms for learning from imperfectly annotated data, which is actually very common in real applications. Nevertheless, such study still needs high-quality test data for reliable evaluation.

9 Ethical Considerations

This paper focuses on revising the annotation on the false negative examples in the DocRED dataset, a publicly available and widely used benchmark for document-level relation extraction. All the documents and relations are provided in the original DocRED dataset. The annotators involved in this work were paid around 60 CNY per hour, which is more than three times the minimum wage in that area. The scope of this work is on revising the relation triples in these documents. However, there

may be improper content within the Wikipedia articles themselves. The authors of this paper are not responsible for ethical issues arising from such improper Wikipedia content.

References

- Christoph Alt, Aleksandra Gabryszak, and Leonhard Hennig. 2020. [TACRED revisited: A thorough evaluation of the TACRED relation extraction task](#). In *Proceedings of ACL*.
- Ateret Anaby-Tavor, Boaz Carmeli, Esther Goldbraich, Amir Kantor, George Kour, Segev Shlomov, Naama Tepper, and Naama Zwerdling. 2020. [Do not have enough data? deep learning to the rescue!](#) In *Proceedings of AAAI*.
- Lidong Bing, Sneha Chaudhari, Richard Wang, and William Cohen. 2015. [Improving distant supervision for information extraction using label propagation through lists](#). In *Proceedings of EMNLP*.
- Lidong Bing, Wai Lam, and Tak-Lam Wong. 2013. [Wikipedia entity expansion and attribute extraction from the web using semi-supervised learning](#). In *Proceedings of WSDM*.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. [Freebase: a collaboratively created graph database for structuring human knowledge](#). In *Proceedings of ACM SIGMOD*.
- 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](#). In *Proceedings of NIPS*.
- Christian Buck, Kenneth Heafield, and Bas Van Ooyen. 2014. [N-gram counts and language models from the common crawl](#). In *Proceedings of LREC*.
- Jih-wei Chen, Tsu-Jui Fu, Chen-Kang Lee, and Wei-Yun Ma. 2021. [H-FND: hierarchical false-negative denoising for distant supervision relation extraction](#). In *Findings of ACL*.
- Qiao Cheng, Juntao Liu, Xiaoye Qu, Jin Zhao, Jiaqing Liang, Zhefeng Wang, Baoxing Huai, Nicholas Jing Yuan, and Yanghua Xiao. 2021. [HacRED: A large-scale relation extraction dataset toward hard cases in practical applications](#). In *Findings of ACL*.
- Yew Ken Chia, Lidong Bing, Soujanya Poria, and Luo Si. 2022. [RelationPrompt: leveraging prompts to generate synthetic data for zero-shot relation triplet extraction](#). *Findings of ACL*.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of ACL*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of NAACL*.
- Markus Eberts and Adrian Ulges. 2021. [An end-to-end model for entity-level relation extraction using multi-instance learning](#). In *Proceedings of EACL*.
- Joseph L Fleiss. 1971. [Measuring nominal scale agreement among many raters](#). *Psychological Bulletin*.
- Tianyu Gao, Xu Han, Yuzhuo Bai, Keyue Qiu, Zhiyu Xie, Yankai Lin, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. 2021. [Manual evaluation matters: Reviewing test protocols of distantly supervised relation extraction](#). In *Findings of ACL*.
- Tianyu Gao, Xu Han, Hao Zhu, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. 2019. [FewRel 2.0: Towards more challenging few-shot relation classification](#). In *Proceedings of EMNLP*.
- Xu Han, Tianyu Gao, Yankai Lin, Hao Peng, Yaoliang Yang, Chaojun Xiao, Zhiyuan Liu, Peng Li, Jie Zhou, and Maosong Sun. 2020. [More data, more relations, more context and more openness: A review and outlook for relation extraction](#). In *Proceedings of AACL*.
- Xu Han, Hao Zhu, Pengfei Yu, Ziyun Wang, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2018. [Fewrel: A large-scale supervised few-shot relation classification dataset with state-of-the-art evaluation](#). In *Proceedings of EMNLP*.
- Kailong Hao, Botao Yu, and Wei Hu. 2021. [Knowing false negatives: An adversarial training method for distantly supervised relation extraction](#). In *Proceedings of EMNLP*.
- Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Sebastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. 2010. [Semeval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation*.
- Quzhe Huang, Shibo Hao, Yuan Ye, Shengqi Zhu, Yansong Feng, and Dongyan Zhao. 2022. [Does recommend-revise produce reliable annotations? an analysis on missing instances in DocRED](#). In *Proceedings of ACL*.
- Varun Kumar, Ashutosh Choudhary, and Eunah Cho. 2020. [Data augmentation using pre-trained transformer models](#). In *Proceedings of the 2nd Workshop on Life-long Learning for Spoken Language Systems*.
- Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick Van Kleef, Sören Auer, et al. 2015. [Dbpedia – a large-scale, multilingual knowledge base extracted from wikipedia](#). *Semantic web*.

- Yangming Li, Lemao Liu, and Shuming Shi. 2021. [Empirical analysis of unlabeled entity problem in named entity recognition](#). In *Proceedings of ICLR*.
- Linlin Liu, Bosheng Ding, Lidong Bing, Shafiq Joty, Luo Si, and Chunyan Miao. 2021. [MulDA: A multilingual data augmentation framework for low-resource cross-lingual NER](#). In *Proceedings of ACL*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *arXiv preprint arXiv:1907.11692*.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. [Distant supervision for relation extraction without labeled data](#). In *Proceedings of ACL*.
- Guoshun Nan, Zhijiang Guo, Ivan Sekulic, and Wei Lu. 2020. [Reasoning with latent structure refinement for document-level relation extraction](#). In *Proceedings of ACL*.
- Tapas Nayak and Hwee Tou Ng. 2019. [Effective attention modeling for neural relation extraction](#). In *Proceedings of CoNLL*.
- Tapas Nayak and Hwee Tou Ng. 2020. [Effective modeling of encoder-decoder architecture for joint entity and relation extraction](#). In *Proceedings of AAAI*.
- Sebastian Riedel, Limin Yao, Andrew McCallum, and Benjamin M. Marlin. 2013. [Relation extraction with matrix factorization and universal schemas](#). In *Proceedings of NAACL*.
- Evan Sandhaus. 2008. [The New York Times annotated corpus](#). *Linguistic Data Consortium*.
- Timo Schick and Hinrich Schütze. 2021. [Generating datasets with pretrained language models](#). In *Proceedings of EMNLP*.
- George Stoica, Emmanouil Antonios Platanios, and Barnabás Póczos. 2021. [Re-tacred: addressing shortcomings of the tacred dataset](#). In *Proceedings of AAAI*.
- Qingyu Tan, Ruidan He, Lidong Bing, and Hwee Tou Ng. 2022. [Document-level relation extraction with adaptive focal loss and knowledge distillation](#). In *Findings of ACL*.
- Denny Vrandečić and Markus Krötzsch. 2014. [Wiki-data: a free collaborative knowledgebase](#). *Communications of the ACM*.
- Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. [ACE 2005 multilingual training corpus](#). *Linguistic Data Consortium*.
- Zihao Wang, Kwunping Lai, Piji Li, Lidong Bing, and Wai Lam. 2019. [Tackling long-tailed relations and uncommon entities in knowledge graph completion](#). In *Proceedings of EMNLP*.
- Peter West, Chandra Bhagavatula, Jack Hessel, Jena D Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck, and Yejin Choi. 2021. [Symbolic knowledge distillation: from general language models to commonsense models](#). *arXiv preprint arXiv:2110.07178*.
- Chaojun Xiao, Yuan Yao, Ruobing Xie, Xu Han, Zhiyuan Liu, Maosong Sun, Fen Lin, and Leyu Lin. 2020. [Denoising relation extraction from document-level distant supervision](#). In *Proceedings EMNLP*.
- Yiben Yang, Chaitanya Malaviya, Jared Fernandez, Swabha Swayamdipta, Ronan Le Bras, Ji-Ping Wang, Chandra Bhagavatula, Yejin Choi, and Doug Downey. 2020. [Generative data augmentation for commonsense reasoning](#). In *Findings of EMNLP*.
- 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](#). In *Proceedings of ACL*.
- Ningyu Zhang, Xiang Chen, Xin Xie, Shumin Deng, Chuanqi Tan, Mosha Chen, Fei Huang, Luo Si, and Huajun Chen. 2021. [Document-level relation extraction as semantic segmentation](#). *Proceedings of IJCAI*.
- 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 EMNLP*.
- Ran Zhou, Xin Li, Ruidan He, Lidong Bing, Erik Cambria, Luo Si, and Chunyan Miao. 2022. [MELM: Data augmentation with masked entity language modeling for low-resource NER](#). In *Proceedings of ACL*.
- Wenxuan Zhou and Muhao Chen. 2021. [An improved baseline for sentence-level relation extraction](#). *arXiv preprint arXiv:2102.01373*.
- Wenxuan Zhou, Kevin Huang, Tengyu Ma, and Jing Huang. 2021. [Document-level relation extraction with adaptive thresholding and localized context pooling](#). In *Proceedings of AAAI*.

Example 1 *Error Type 1: Misunderstanding of Definition: creator vs. architect*

Error Cause: Annotator's Misunderstanding of Wikidata Relation Definition

Error Analysis: There are over 9,000 relation types in the Wikidata knowledge base, whereas DocRED only contains 84 of them. The relation between buildings and its designer is architect (P84). But this relation is not in the DocRED's label space. In this case, the architect relation shall be excluded from DocRED. However, annotators from Huang et al. (2022) uses creator to describe this kind of relation, which is not precise.

Document: Sir [David Alan Chipperfield](#) (born 18 December 1953) is an English architect . He established David Chipperfield Architects in 1985 . His major works include the [River and Rowing Museum](#) in Henley - on - Thames , Oxfordshire (1989, 1998) ; the [Museum of Modern Literature in Marbach](#) , Germany ; the [Des Moines Public Library](#) , Iowa (2002, 2006) ; the [Neues Museum](#) , Berlin (1997 , 2009) ; The [Hepworth Wakefield gallery](#) in Wakefield , UK (2003, 2011) , the [Saint Louis Art Museum](#) , Missouri (2005, 2013) ; and the [Museo Jumex](#) in Mexico City (2009, 2013) . Rowan Moore , the architecture critic of the Guardian of London , described his work as serious , solid , not flamboyant or radical , but comfortable with the history and culture of its setting . " He deals in dignity , in gravitas , in memory and in art . " David Chipperfield Architects is a global architectural practice with offices in London , Berlin , Milan , and Shanghai .

Wrong Triples: ([Museo Jumex](#), *creator*, [David Alan Chipperfield](#)), ([River and Rowing Museum](#), *creator*, [David Alan Chipperfield](#)), ([Museum of Modern Literature](#), *creator*, [David Alan Chipperfield](#)), ([Saint Louis Art Museum](#), *creator*, [David Alan Chipperfield](#)), ([Hepworth Wakefield gallery](#), *creator*, [David Alan Chipperfield](#)), ([Neues Museum](#), *creator*, [David Alan Chipperfield](#))

Example 2 *Error Type 1: Misunderstanding of Definition: series vs. part of*

Error Cause: Annotator's Misunderstanding of Wikidata Relation Definition

Document: [Chapman Square](#) is the debut studio album released by four piece British band Lawson . The album was released on 19 October 2012 via Polydor Records . The album includes their three top ten singles " When She Was Mine " , " [Taking Over Me](#) " and " [Standing in the Dark](#) " . The album was mainly produced by John Shanks with Duck Blackwell , Paddy Dalton , Ki Fitzgerald , Carl Falk , and Rami Yacoub . The album was re - released in the autumn of 2013 as [Chapman Square Chapter II](#) , with the lead single from the re - release being " Brokenhearted " , which features American rapper B.o . B. As of July 2016 , the album has sold 169,812 copies .

Wrong Triples: ([Taking Over Me](#), *series*, [Chapman Square Chapter II](#)), ([Brokenhearted](#), *series*, [Chapman Square Chapter II](#)), ([When She Was Mine](#), *series*, [Chapman Square](#)), ([Standing in the Dark](#), *series*, [Chapman Square](#))

Example 3 *Error Type 2: Commonsense Bias*

Document: Alecu Russo (born in March 17 , 1819 , near Chişinău , died on February 5 , 1859 , in [Iaşi](#)) , was a [Moldavian](#) Romanian writer , literary critic and publicist . Russo is credited with having discovered one of the most elaborate forms of the Romanian national folk ballad Mioriţa . He was also a contributor to the Iaşi periodical Zimbrul , in which he published one of his best - known works , Studie Moldovană (" Moldovan Studies ") , in 1851 - 1852 . He also wrote Iaşii şi locuitorii lui în 1840 " Iaşi and its inhabitants in 1840 " – a glimpse into Moldavian society during the Organic Statute administration , and two travel accounts (better described as folklore studies) , Piatra Teiului and Stâncă Corbului . Russo is also notable for his Amintiri (" Recollections ") , a memoir .

Wrong Triples: ([Iaşi](#), *located in administrative territorial entity*, [Moldavian](#)), ([Moldavian](#), *contains administrative territorial entity*, [Iaşi](#))

Example 4 *Error Type 3: Slippery slope*

Error Cause: Improper reasoning based on punctuation for judgement of date of birth/death.

Document: South Wigston High School was founded in 1938 and is a school serving the local community of South Wigston .Today the school is an 11 – 16 yrs Academy. The main feeder primary schools are Glen Hills , Fairfield and Parkland . The school also attracts students from many areas of the city of Leicester and the county of Leicestershire . The school is oversubscribed and is growing year on year . South Wigston is known for its wide range of extra - curricular opportunities and for being a school that is inclusive and at the heart of the community . The school has extensive grounds and a purpose build sports centre opened by Gary Lineker . [Notable alumni](#) include [Sue Townsend \(1946 - 1950 \)](#) , author ; [Louis Deacon \(1991 - 1995 \)](#) , Rugby Player for Leicester Tigers and England ; and [Brett Deacon \(1992 - 1996 \)](#) , Rugby Player for Leicester Tigers and England . In 2016 BBC2 produced a documentary entitled , The Secret Life of Sue Townsend Aged 68. Much of the documentary was filmed at the school and current students participated .

Wrong Triples: ([Louis Deacon](#), *date of birth*, [1991](#)), ([Louis Deacon](#), *date of death*, [1995](#)), ([Brett Deacon](#), *date of death*, [1996](#)), ([Brett Deacon](#), *date of birth*, [1992](#)), ([Sue Townsend](#), *date of birth*, [1946](#)), ([Sue Townsend](#), *date of death*, [1950](#))

Table 10: Examples for the common error types by Huang et al. (2022). We use blue to color the **entities** and green to color the **relations**.

A Our Machine-Guided Annotation vs. Annotation from Scratch

In this section, we compare our Re-DocRED dataset with a concurrent work (Huang et al., 2022) on revising the DocRED dataset. Our work uses machine-guided annotation methods, whereas their work asks the annotators to annotate from scratch (denoted as “From Scratch”). As mentioned in Huang et al. (2022), annotating from scratch is an extremely challenging task. This is mainly due to the quadratic complexity of the document-level RE task. Suppose there are N entities in one document and the label space of interest contains R relations. The search space for human annotation is $N * (N - 1) * R$. In particular, for an average case of DocRED ($N=20$, $R=96$), an annotator will need to make 36,480 classification decisions for one document. In contrast, for our machine-guided annotation, annotators will only need to make decisions on the recommended candidates, averaging only 25.5 decisions per document. This is primarily due to the pattern recognition capability of deep neural models, which significantly reduces the search space for human annotators. The size of our re-annotated dataset is much larger and our dataset contains all relations in the DocRED’s label space, while the Scratch dataset only contains 91 out of the 96 relations in DocRED. Besides, we have also conducted significantly more experimental analysis.

It is also mentioned in Huang et al. (2022) that re-annotation from scratch by two human experts may still not be the ground truth due to natural error rate. Therefore, we examined and analyzed all the annotated 96 documents of “From Scratch”. We found that there are several types of systematic errors in Huang et al. (2022) and we show the error types in Table 11. Firstly, annotation from scratch is susceptible to annotators’ misunderstanding of relation definition (25.5%). For the first example in Table 10, the relation between an architect and his designed building is *architect*³, whereas annotators in Huang et al. (2022) deem such relation as *creator*⁴. This is imprecise as the *architect* relation was not in the label space of DocRED. Therefore, it is not possible to find this relation between the architect and its design. Secondly, annotating from scratch is susceptible to human commonsense bias (19.0%). This is primarily due to human’s memo-

³<https://www.wikidata.org/wiki/Property:P84>

⁴<https://www.wikidata.org/wiki/Property:P170>

Error Types	Percentage
Misunderstanding of definition	25.5%
Commonsense bias	19.0%
Slippery slope reasoning	22.1%
Others	33.4%

Table 11: Common error types of Huang et al. (2022).

	Added Triples	Errors	Precision
From Scratch	2,057	293	85.8
Re-DocRED	2,067	76	96.3

Table 12: Error rates of Re-DocRED and the From Scratch dataset (Huang et al., 2022) based on examining all 96 documents. We observe that Re-DocRED has higher precision.

rization of popular entities, such as countries and geographical locations (example 3 in Table 10). The third major error type is due to slippery slope reasoning, as shown by example 4 in Table 10. The numbers within brackets were falsely identified as the *date of birth* and *date of death*, whereas the passage is about a renowned high school. It can be inferred that the numbers behind the alumni names are indicating the time periods that they were in this school. This error arises because most *date of birth* and *date of death* are described by brackets and numbers. However, such a pattern does not necessarily mean all numbers in brackets are indicating such relations.

It is worth noting that the annotators in Huang et al. (2022) are already experts in English and the annotators went through discussion after annotation. However, there are still a considerable number of errors from their dataset. We believe that this is due to the complex nature of this annotation task. In addition, we also conducted human evaluation of our Re-DocRED dataset and compared the precision of the two datasets in Table 12. We can see that our Re-DocRED dataset has significantly higher precision for the added triples. Moreover, we compare the unit price and unit time for different annotation strategies in Table 13. We can see that annotating from scratch costs three times more than our machine-guided annotation.

Hence, by comparing the two approaches for annotating document-level relation extraction datasets, we conclude that:

1. Even though “From Scratch” annotation is conducted by human experts, there are still miss-

	Unit Price	Unit Time
Annotating from Scratch	48 CNY	40 mins
One Round of Revision	7.8 CNY	10 mins
Two Rounds of Revision	15.8 CNY	15 mins

Table 13: Cost and unit time required for different annotation strategies.

ing triples in the annotated 96 documents. That is, annotating from scratch does not completely eliminate the incompleteness problem when the number of entities N and relation types R are large.

2. Annotation from scratch is not as precise as recommend-revise. As Table 10 shows, human annotation of Huang et al. (2022) contains several types of systematic errors.

3. Annotating from scratch is hard to scale. According to Huang et al. (2022) and feedback from our annotators, it takes more than 30 minutes by experts to annotate one document. Then the two experts will still spend extra time discussing and resolving the conflicts.

4. The recommend-revise scheme is able to mitigate the false negative problem and is easier to scale up.

B Coreferential Error Annotation

Coreferential Errors Besides the major problem of incomplete annotation, coreferential errors are also detrimental to the evaluation of DocRED. Note that an entity in the DocRED dataset can have multiple mention appearances in a document. If some mentions that are referring to the same entity are not included in the entity cluster, a redundant entity cluster will be formed. This kind of coreferential errors will affect the relation predictions involving the redundant entity. Since the complexity of the DocRE task is quadratic in the number of entities, it is important to make sure that the coreferential annotations are correct. Errors in coreferential annotation can be propagated and amplified during relation extraction.

Coreferential Annotation There are a considerable number of entities of DocRED that have the same surface names but refer to different entities. In such cases, we examined all entities that contain the same surface name and entity types in the DocRED dataset. For these entities, annotators will need to decide whether: (1) the two entities are coreferential to each other, (2) the overlapping

mentions are wrongly grouped to a certain entity cluster, and (3) the two mentions are indeed referring to two separate entities. As a result, we merged 102 coreferential entity pairs in the evaluation documents and 122 pairs in the training documents. Therefore, the average number of entities per document of Re-DocRED is slightly lower than the original DocRED.

C Details on Negative Sampling

This section describes the details of the experiments with negative sampling. Given a text T and a set of n entities $\{e_1, \dots, e_n\}$, the document-level RE task requires making predictions on all possible entity pairs (e_i, e_j) for $i, j \in 1, \dots, n, i \neq j$. As we can see in Table 1, over 90% of the entity pairs are negative instances. Instead of using all negative instances, we sample a fraction of negative instances for training. Figure 3 shows the performance with respect to the sampling rate when the model is trained on the training set of DocRED. Figure 4 shows the performance when the training set of Re-DocRED is used. Note that evaluation is conducted on the development set of Re-DocRED. From both figures, we observe that precision is positively correlated with the negative sampling percentage and recall is negatively correlated with this percentage. Besides, when training on highly incomplete data (DocRED), sampling 10% of the negative instances improves the F1 score by 14.72 (62.21 vs 47.49) compared to using 100% of the negative instances. However, the performance improvement is not that significant when training on Re-DocRED (79.13 vs 78.47), and the best sampling rate in this scenario is 70%.




Figure 3: Performance with respect to sampling rate. The model is trained on the training set of DocRED and is evaluated on the development set of Re-DocRED.

D Common Model Errors

In this section, we show examples of the mistakes predicted by the best model (KD-DocRE model

Example 1 Error Type 1: Extraneous Prediction (MR)**Error Cause:** Popular pattern bias**Document:** "Lookin Ass" is a song by American rapper and singer [Nicki Minaj](#). It was produced by Detail and Choppa Boi. It was recorded by [Minaj](#) for the [Young Money Entertainment](#) compilation album (2014). The music video for the track was released on February 14, 2014. ... On March 11, 2014, "Lookin Ass" was serviced to urban contemporary radio in the United States as Young Money: Rise of an Empires third official single. It was sent to US rhythmic radio stations on March 18, 2014, two weeks after its predecessor, "[Trophies](#)".**Extraneous Triples:** ([Young Money Entertainment](#), *performer*, [Nicki Minaj](#)), ([Trophies](#), *performer*, [Nicki Minaj](#))

Example 2 Error Type 1: Extraneous Prediction (MR)**Error Cause:** Popular pattern bias**Document:** The Portland Golf Club is a private golf club in the northwest United States, in suburban [Portland](#), Oregon. It is located in the unincorporated Raleigh Hills area of eastern [Washington County](#), southwest of downtown [Portland](#) and east of Beaverton. PGC was established in the winter of 1914, when a group of nine businessmen assembled to form a new club after leaving their respective clubs. The present site was chosen due to its relation to the Spokane, Portland and Seattle Railway's interurban railroad line with frequent passenger service to the site because automobiles and roads were few. ...**Extraneous Triples:** ([Portland](#), *located in the administrative territorial entity*, [Washington County](#)); ([Washington County](#), *contains administrative territorial entity*, [Portland](#))

Example 3 Error Type 2: Missing Triples (MS)**Error Cause:** Failed in coreferential reasoning**Document:** [Kurt Tucholsky](#) (; 9 January 1890 – 21 December 1935) was a German - Jewish journalist, satirist, and writer. He also wrote under the pseudonyms [Kaspar Hauser](#) (after the historical figure), [Peter Panter](#), [Theobald Tiger](#) and [Ignaz Wrobel](#). Born in Berlin - Moabit, he moved to Paris in 1924 and then to Sweden in 1929. [Tucholsky](#) was one of the most important journalists of the Weimar Republic. As a politically engaged journalist and temporary co - editor of the weekly magazine [Die Weltbühne](#) he proved himself to be a social critic in the tradition of [Heinrich Heine](#).**Missing Triples:** ([Kurt Tucholsky](#), *employer*, [Die Weltbühne](#))

Example 4 Error Type 2: Missing Triples (MS)**Error Cause:** Fail to find long-tail relations**Document:** CBBC (short for Children's BBC) is a British children's television strand owned by the BBC and aimed for children aged from 6 to 12. BBC programming aimed at under six year old children is broadcast on the CBeebies channel. CBBC broadcasts from 7 am to 9 pm on the digital CBBC Channel, available on most UK digital platforms. The CBBC brand was used for the broadcast of children's programmes on BBC One on weekday afternoons and on BBC Two mornings until these strands were phased out in 2012 and 2013 respectively, as part of the BBC's "Delivering Quality First" cost-cutting initiative. CBBC programmes were also broadcast in high definition alongside other BBC content on [BBC HD](#), generally at afternoons on weekends, unless the channel was covering other events. This ended when [BBC HD](#) closed on [26 March 2013](#), but CBBC HD launched on 10 December 2013.**Missing Triples:** ([BBC HD](#), *dissolved, abolished or demolished*, [26 March 2013](#))

Table 14: Examples of the two most common error types. We use blue and green to color the **entities** and **relations**, respectively.




Figure 4: Performance with respect to sampling rate. The model is trained and evaluated on Re-DocRED.

with distant supervision pretraining) in our experiments. Similar to [Tan et al. \(2022\)](#), we split the union of ground-truth triples and predicted triples into four categories: (1) **Correct (C)**, where predicted triples are in the ground truth. (2) **Wrong**

Ground Truth			
Predictions	$r \in \mathbf{R}$	NA	
	$r \in \mathbf{R}$	C: 13,610 (71.59%) W: 580 (3.05%)	MR: 1,884 (9.91%)
	NA	MS: 2,935 (15.44%)	

Table 15: Statistics of our error distribution on the development set of Re-DocRED based on KD-DocRE. The final evaluation score is evaluated on $r \in \mathbf{R}$ triples, hence the correct predictions of **NA** are ignored when calculating the final scores.

(**W**), where the output relation type is wrong but the predicted head and tail entities are in the ground truth. (3) **Missed (MS)**, where the model predicts no relation for a pair of head and tail entities but there is some relation in the ground truth. (4) **More (MR)**, where the model predicts an extraneous re-

lation for a pair of head and tail entities that is not in the ground truth. The performance breakdown is shown in Table 15. We observe that the majority of the errors were in the **MS** and **MR** categories. We found that predictions on popular relations tend to fall under the **MR** category. We further show this popularity bias pattern with examples in Table 14. We can see that popular relation patterns tend to be under the **MR** category. From the first example, “Nicki Minaj” is a popular artiste. But the song “Trophies” was not performed by “Nicki Minaj”. In contrast, the model wrongly predicts that this song is also performed by her. On the other hand, in example 4, the cue that “BBC HD” dissolved on “Mar 26, 2013” is obvious, whereas the model failed to find this relation.

E Details of Logical Rules

In this section, we show the logical rules that we used. After examining the DocRED dataset (Yao et al., 2019), we found that there are two types of logical inconsistencies. The first type is the incompleteness of inverse relations, and the second is the inconsistency in co-occurring relations. The inverse relations are logical relations that can be implied by reversing the direction of relation triples. For example, if entity 1 is the *participant* of an event (entity 2), this event should have a *participant* relation with entity 1. We show all the inverse relation pairs that we used in Table 17. Besides inverse relations, we also added triples by co-occurring rules. This is mainly because we found that these relations are logically correlated and their co-occurrence is inconsistent in the original DocRED dataset. For example, when describing the relation between entities and wars, there are two involved relations: relation *conflict*⁵ and *participant* of⁶. For such cases, the two relations are considered to be present when *conflict* is present. Similarly, if a triple (**entity1**, *country*, **entity2**) is present and **entity1** is of type *location* or *organization*, (**entity1**, *located in*, **entity2**) is considered to be present as well. We show the list of co-occurring relations in Table 16.

The logical rules automatically add relation triples to the Re-DocRED dataset, but it is possible that erroneous triples are also added due to corner cases. However, according to the human evaluation described in Appendix A, the precision of our Re-

DocRED dataset is still higher than (Huang et al., 2022).

Relation	Co-occurring relation
country	located in
conflict	participant of

Table 16: List of co-occurring relations.

Relation	Inverse relation
author	notable work
performer	notable work
producer	notable work
composer	notable work
director	notable work
lyrics by	notable work
participant	participant of
participant of	participant
has part	part of
sibling	sibling
series	has part
spouse	spouse
characters	present in work
conflict	participant
parent organization	subsidiary
subsidiary	parent organization
follows	followed by
followed by	follows
father	child
replaced by	replaces
head of government	applies to jurisdiction
replaces	replaced by
legislative body	applies to jurisdiction
head of state	applies to jurisdiction
mother	child
part of	has part
sister city	sister city
capital	capital of

Table 17: List of inverse relations.

⁵<https://www.wikidata.org/wiki/Property:P607>

⁶<https://www.wikidata.org/wiki/Property:P1344>