

Universal Information Extraction with Meta-Pretrained Self-Retrieval

Xin Cong^{1,2} Bowen Yu^{1,2} Mengcheng Fang³ Tingwen Liu^{1,2*} Haiyang Yu³
Zhongkai Hu³ Fei Huang³ Yongbin Li^{3*} and Bin Wang⁴

¹Institute of Information Engineering, Chinese Academy of Sciences. Beijing, China

²School of Cyber Security, University of Chinese Academy of Sciences. Beijing, China

³DAMO Academy, Alibaba Group

⁴Xiaomi AI Lab, Xiaomi Inc., Beijing, China

{congxin, liutingwen}@iie.ac.cn

{yubowen.ybw, yifei.yhy, zhongkai.hzk, f.huang, shuide.lyb}@alibaba-inc.com

fmc1121@zju.edu.cn

wangbin11@xiaomi.com

Abstract

Universal Information Extraction (Universal IE) aims to solve different extraction tasks in a uniform text-to-structure generation manner. Such a generation procedure tends to struggle when there exist complex information structures to be extracted. Retrieving knowledge from external knowledge bases may help models to overcome this problem but it is impossible to construct a knowledge base suitable for various IE tasks. Inspired by the fact that large amount of knowledge are stored in the pretrained language models (PLM) and can be retrieved explicitly, in this paper, we propose MetaRetriever to retrieve task-specific knowledge from PLMs to enhance universal IE. As different IE tasks need different knowledge, we further propose a Meta-Pretraining Algorithm which allows MetaRetriever to quickly achieve maximum task-specific retrieval performance when fine-tuning on downstream IE tasks. Experimental results show that MetaRetriever achieves the new state-of-the-art on 4 IE tasks, 12 datasets under fully-supervised, low-resource and few-shot scenarios.

1 Introduction

Information extraction (IE) is the task of extracting specific information structures of interest from unstructured text (Andersen et al., 1992; Grishman, 2019). The various IE tasks at hand can be highly heterogeneous, as each task may have different extraction targets, such as entities (Lample et al., 2016), relations (Zheng et al., 2017), and events (Lin et al., 2018). Traditional IE methods tend to be task-specific, resulting in specialized and isolated approaches for different tasks.

In recent years, researchers have made efforts towards universal information extraction, or Universal IE (Lu et al., 2022), which aims to develop

a unified method capable of solving a range of IE tasks. One proposed approach involves the use of a Structure Extraction Language to express different extraction targets in a unified structure form, and pretrains a Transformer model to generate this unified structure through a sequence-to-sequence process. However, this structure generation process can be limited in its ability to fully utilize contextual semantic correlation between the extracted information, as the Transformer generation process is unidirectional. In cases where the information structure to be extracted is complex, the universal IE model may struggle to generate accurate results.

Previous studies on retrieval-augmented generation (Lewis et al., 2020b; Wang et al., 2021; Cai et al., 2022; Geng et al., 2022) have shown that utilizing external knowledge bases to retrieve task-specific information can improve the ability of models to generate complex sequences. However, this approach is not practical for Universal IE, as it is almost impossible to build a knowledge base suitable for various IE tasks, and it is also the ultimate goal of IE. Recently, research on knowledge probing (Petroni et al., 2019; Roberts et al., 2020; Zhang et al., 2022; Yu et al., 2022) has demonstrated that large amounts of knowledge are stored in pretrained language models (PLMs) and can be retrieved explicitly. Utilizing this knowledge to enhance IE models could avoid the need for extensive resource construction. In light of these findings, the question arises: *can PLMs be used as knowledge bases to retrieve knowledge and improve universal IE models?* If so, universal IE models would be able to generate more accurate results.

Based on this idea, in this paper, we propose MetaRetriever, a pretrained model for universal IE. Unlike existing method which generates results in a single step, MetaRetriever utilizes a retrieve-then-extract approach: first, it retrieves

*Corresponding Author

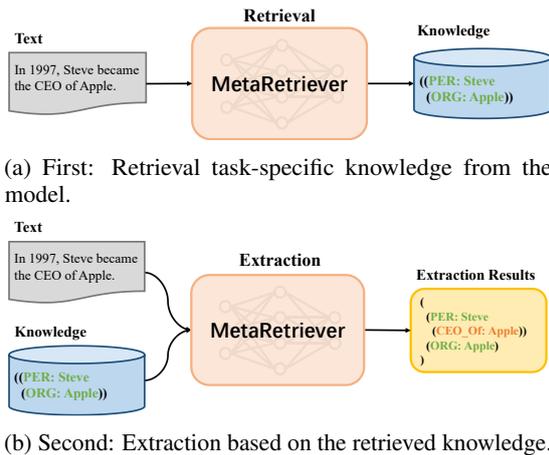


Figure 1: Illustrations of our proposed MetaRetriever with the retrieve-then-extract manner.

task-specific knowledge from itself, and then it uses this knowledge as additional input to extract information. However, different IE tasks require different knowledge, and simply retrieving information from PLM can lead to irrelevant information being included, which can negatively impact performance. To address this, incorporating meta-learning techniques (Finn et al., 2017), we develop a Meta-Pretraining Algorithm (MPA) which allows the model to quickly learn the semantics of extraction targets of different tasks to retrieve the relevant knowledge. Specifically, MPA pretrains the model with a bi-level optimization, in which the model are optimized to achieve maximal task-specific retrieval performance given the extraction targets after the parameters have been updated through a small number of gradient steps on a new IE task. Thus, in the fine-tuning phase, MetaRetriever can learn downstream tasks quickly to retrieve relevant knowledge and disregard irrelevant information.

We conduct experiments on 12 datasets of 4 main IE tasks under fully-supervised, few-shot and low-resource scenarios. Experiments show that MetaRetriever significantly surpasses previous works and achieves the new state-of-the-art. On fully-supervised scenario, MetaRetriever achieves 0.7% gains on average over UIE (Lu et al., 2022). On few-shot and low-resource scenarios, our model outperforms UIE with over 2.0% improvement.

Our contributions can be summarized as follows: (1) We introduce a retrieve-then-extract perspective into universal IE. It retrieves task-specific knowledge from the pretrained language model to improve performance. (2) We propose a meta-pretraining algorithm to make models fast-adapt

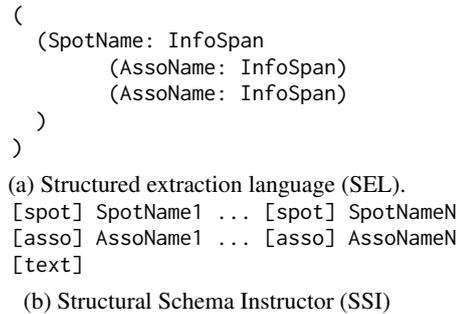


Figure 2: Illustrations of SEL and SSI.

into various downstream IE tasks by retrieving task-specific knowledge. (3) We construct large-scale pretraining corpus for universal IE and it can promote future research¹.

2 Preliminaries

2.1 Structured Extraction Language

The UIE paper (Lu et al., 2022) presents the concept of the Structured Extraction Language (SEL) for expressing the diverse extraction targets of various IE tasks in a uniform representation. As illustrated in Figure 2a, SEL is a hierarchical key-value structure that comprises three components:

- (1) SpotName: the targeted class name of the extracted information pieces in a specific IE task, e.g., “PERSON” in named entity recognition;
- (2) AssoName: the targeted class name of the relationship between different information, e.g., “Works For” in relation extraction;
- (3) InfoSpan: the text span corresponding to the specific SpotName or AssoName.

Additionally, the colon (“:”) symbol denotes the mapping of the text span to its class name or relationship, and the hierarchical structure is indicated by the use of parentheses (“(” and “)”).

2.2 Structural Schema Instructor

To make the model be aware of which extraction targets the IE task focus on, UIE further proposes Structural Schema Instructor (SSI), a schema-based prompt. As shown in Figure 2b, SSI first uses several special symbols ([spot], [asso]) to denote SpotNames and AssoNames. Then these SpotNames and AssoNames are concatenated before each text sequence with a “[text]” symbol as a prompt. UIE takes the prompted text as input to generate the linearized SEL to achieve information extraction.

¹Our code and constructed corpus are released at <https://github.com/AlibabaResearch/DAMO-ConvAI/tree/main/metaretriever>

3 Methodology

Our proposed MetaRetriever achieves universal IE via a retrieve-then-extract manner. It first retrieves task-specific knowledge from itself according to the input text and then generates the extraction results based on the retrieved knowledge. To make MetaRetriever only retrieve task-specific knowledge and neglect other distracting knowledge, we further proposed a meta-pretraining algorithm based on meta-learning technique. In this section, we first introduce how MetaRetriever works via the retrieve-then-extract manner and then detail the meta-pretraining algorithm.

3.1 Retrieve-then-Extract

Given that providing task-specific knowledge can be beneficial for IE and that pretrained models have acquired a significant amount of knowledge through training on large corpora (Petroni et al., 2019; Roberts et al., 2020; Zhang et al., 2022; Yu et al., 2022), our proposed system, MetaRetriever, utilizes this pretrained knowledge by first retrieving relevant information from the model itself. This retrieved knowledge then serves as an additional input, along with the original input, to produce the final extraction results. The design of MetaRetriever is depicted in Figure 1.

3.1.1 Retrieval Procedure

We use SEL (Lu et al., 2022) to express knowledge for knowledge retrieval and knowledge understanding. To specify which extraction targets MetaRetriever should retrieve, we adopt SSI as a prefix that controls which kinds of knowledge need to be retrieved. Thus, the input is in the form of:

$$s \oplus x = [s_1, s_2, \dots, s_{|s|}, x_1, x_2, \dots, x_{|x|}] \quad (1)$$

where $s = [s_1, s_2, \dots, s_{|s|}]$ is SSI and $x = [x_1, x_2, \dots, x_{|x|}]$ is the original text. Defining a unified format of task-specific knowledge for various IE tasks is a challenging issue. In this paper, we adopt a simple strategy: using the ground truth linearized SEL sequence of the corresponding input text as the knowledge we wish to retrieve. Taking SSI (s) and the text sequence (x) as input, MetaRetriever generates the linearized SEL (k) to retrieve task-specific knowledge. The retrieval procedure can be represented as follows:

$$k = \text{MetaRetriever}(s \oplus x) \quad (2)$$

where $k = [k_1, \dots, k_{|k|}]$ is the retrieved knowledge represented by SEL.

Algorithm 1 Meta-Pretraining Algorithm

Require: α, β : inner/outer step size; \mathcal{D} : pretraining corpus

- 1: Initialize parameters θ of the model
- 2: Construct support-query pair based on \mathcal{D}
- 3: // Outer Loop
- 4: **while** not converged **do**
- 5: Sample the support set (x^S, y^S)
 and the query set (x^Q, y^Q)
- 6: Compute record loss $\mathcal{L}_{\text{record}}(y; \theta)$
- 7: Corrupt raw text: $x', x'' \leftarrow \text{Corrupt}(x)$
- 8: Compute LM loss $\mathcal{L}_{\text{LM}}(x', x''; \theta)$
- 9: // Inner Loop
- 10: **for** j in number of inner-loop updates J **do**
- 11: Compute retrieval loss $\mathcal{L}_{\text{retrv}}(x^S, y^S; \theta_j)$
- 12: Compute extraction loss $\mathcal{L}_{\text{ext}}(x^S, k^S, y^S; \theta_j)$
- 13: Compute inner loss $\mathcal{L}_{\text{inner}} = \mathcal{L}_{\text{retrv}} + \mathcal{L}_{\text{ext}}$
- 14: Compute adapted parameters
 $\theta_{j+1} \leftarrow \theta_j - \alpha \nabla_{\theta_j} \mathcal{L}_{\text{inner}}$
- 15: **end for**
- 16: Compute retrieval loss $\mathcal{L}_{\text{retrv}}(x^Q, y^Q; \theta_J)$
- 17: Compute extraction loss $\mathcal{L}_{\text{ext}}(x^Q, k^Q, y^Q; \theta_J)$
- 18: Compute outer loss
 $\mathcal{L}_{\text{outer}} = \mathcal{L}_{\text{retrv}} + \mathcal{L}_{\text{ext}} + \mathcal{L}_{\text{record}} + \mathcal{L}_{\text{LM}}$
- 19: Update model parameters $\theta \leftarrow \theta - \beta \nabla_{\theta} \mathcal{L}_{\text{outer}}$
- 20: **end while**

3.1.2 Extraction Procedure

After retrieval of task-specific knowledge, MetaRetriever incorporates this information as an additional input to generate the extraction results. To ensure that MetaRetriever utilizes the retrieved knowledge in generating the results, we concatenate it with the original input by appending it as a suffix:

$$s \oplus x \oplus k = [s_1, \dots, s_{|s|}, x_1, \dots, x_{|x|}, k_1, \dots, k_{|k|}] \quad (3)$$

Then, MetaRetriever will take $s \oplus x \oplus k$ as the new input to generate the final extraction results y :

$$y = \text{MetaRetriever}(s \oplus x \oplus k) \quad (4)$$

3.2 Meta-Pretraining Algorithm

As the knowledge contained in the pretrained model may not always be relevant to the downstream IE tasks, we propose a Meta-Pretraining Algorithm (MPA) that utilizes a meta-learning technique (Finn et al., 2017) to enable MetaRetriever to quickly learn the semantics of the extraction targets of downstream IE tasks, and thus, retrieve task-specific knowledge only. The algorithm is detailed in Algorithm 1.

To make MetaRetriever retrieve task-specific knowledge, MPA meta-pretrain MetaRetriever with a bi-level optimization (Finn et al., 2017) to optimize MetaRetriever to achieve maximum task-specific retrieval performance after the parameters have been updated through a small number of gradient steps on a new IE task. In the bi-level optimization, models will experience two optimization

trials: the inner loop and the outer loop. The inner loop aims to mimic the downstream fine-tuning procedure to compute a intermediate parameters of the model via SGD (Saad, 1998; Sutskever et al., 2013). The outer loop aims to mimic the downstream evaluation procedure to optimize the model parameters based on the intermediate parameters computed in the inner loop. The inner loop and the outer loop are nested and thus, it calculates high-order gradients to update parameters. Optimized by such a bi-level optimization, MetaRetriever would learn that after finetuning a small number of steps (the inner loop), MetaRetriever can achieve lower retrieval loss (the outer loop), a.k.a. learning a fast-adaptation ability for task-specific retrieval.

To accomplish this, MetaRetriever should be pre-trained on a set of simulated “IE tasks” \mathcal{T} in the pre-training phase. Each “IE task” contains a support set \mathcal{S} and a query set \mathcal{Q} . The support set is to mimic the training set used in the inner loop and the query set is to mimic the test set used in the outer loop. To construct various “IE tasks”, a intuitive way is to follow the widely-used *episodic training* technique in conventional meta-learning methods (Vinyals et al., 2016) to randomly sample instances by category. However, such a random-sampling operation requires random access to all data, which is infeasible when processing large-scale corpora, as such corpora cannot be stored in memory and must be accessed from external storage instead, resulting in intolerable long input/output (IO) time. To address this limitation, we design a support-query pairing method based on the graph maximum-weighted matching algorithm (Galil, 1986), which significantly reduces IO time (see in Section 5.5). Details of our method are introduced in Appendix B.

Inner Loop Since the inner loop simulates the downstream fine-tuning process, MetaRetriever computes the intermediate parameters by considering two inner losses: the retrieval loss $\mathcal{L}_{\text{retrv}}$ and the extraction loss \mathcal{L}_{ext} . Let MetaRetriever be represented by a parametrized function f_θ with parameters θ . When learning a mimic IE task $\mathcal{T}_i = \{\mathcal{S}_i, \mathcal{Q}_i\}$, the retrieval loss is computed based on the support set $\mathcal{S}_i = (x^{\mathcal{S}_i}, y^{\mathcal{S}_i})$.

$$\mathcal{L}_{\text{retrv}} = \text{CrossEntropy}(f_\theta(x^{\mathcal{S}_i}), y^{\mathcal{S}_i}) \quad (5)$$

It takes the text with SSI as the input and calculates the cross-entropy loss between the output sequence and the ground truth retrieved knowledge SEL se-

quence. Then, we compute the extraction loss:

$$\mathcal{L}_{\text{ext}} = \text{CrossEntropy}(f_\theta(x^{\mathcal{S}_i}, k^{\mathcal{S}_i}), y^{\mathcal{S}_i}) \quad (6)$$

where MetaRetriever takes the retrieved knowledge $k^{\mathcal{S}_i}$ with original input and generates the final predicted linearized SEL expression. We compute the cross-entropy loss between the predicted linearized SEL and ground truth as the extraction loss. Thus, the total inner loss is calculated as follows:

$$\mathcal{L}_{\text{inner}} = \mathcal{L}_{\text{retrv}} + \mathcal{L}_{\text{ext}} \quad (7)$$

The inner loop may consist of several update steps, resulting in multiple updates of the intermediate parameters. Let the number of steps in the inner loop be denoted as J , then the model parameters θ are updated iteratively to θ_J based on the inner loss $\mathcal{L}_{\text{inner}}$:

$$\theta_{j+1} = \theta_j - \alpha \nabla_{\theta_j} \mathcal{L}_{\text{inner}}(f_{\theta_j}) \quad (8)$$

where α is the inner step size.

Outer Loop In the outer loop, the model parameters are updated by optimizing the performance of f_{θ_J} with respect to θ . As the outer loop simulates the downstream evaluation process, given the query set $\mathcal{Q}_i = (x^{\mathcal{Q}_i}, y^{\mathcal{Q}_i})$, the retrieval loss is calculated to evaluate the retrieval performance:

$$\mathcal{L}_{\text{retrv}} = \text{CrossEntropy}(f_{\theta_J}(x^{\mathcal{Q}_i}), y^{\mathcal{Q}_i}) \quad (9)$$

The extraction loss is computed to assess the extraction results based on the retrieved knowledge:

$$\mathcal{L}_{\text{ext}} = \text{CrossEntropy}(f_{\theta_J}(x^{\mathcal{Q}_i}, k^{\mathcal{Q}_i}), y^{\mathcal{Q}_i}) \quad (10)$$

Additionally, to maintain the ability to understand language, we include a T5-style language modeling loss \mathcal{L}_{LM} (Raffel et al., 2020a). And to acquire the capability of generating valid linearized SEL structures, we include a record loss $\mathcal{L}_{\text{record}}$ as proposed in UIE (Lu et al., 2022). The overall objective in the outer loop is calculated using these four losses:

$$\mathcal{L}_{\text{outer}} = \mathcal{L}_{\text{retrv}} + \mathcal{L}_{\text{ext}} + \mathcal{L}_{\text{LM}} + \mathcal{L}_{\text{record}} \quad (11)$$

It is worth noting that the language modeling loss and the record loss are not the primary learning objectives in the downstream IE tasks, thus these losses are calculated based on the original parameters θ rather than the intermediate parameters θ_J .

Eventually, the model parameters θ are updated as follows:

$$\theta \leftarrow \theta - \beta \nabla_{\theta} \sum_{\mathcal{T}_i \in \mathcal{T}} \mathcal{L}_{\text{outer}}(f_{\theta_J}) \quad (12)$$

where β is the outer step size. The bi-level optimization is applied to the model parameters θ , but the objective is computed using the updated model parameters θ_J , which calculates high-order gradients to update parameters. Essentially, our proposed method optimizes the model parameters such that one or a small number of gradient steps on a new task will yield optimal performance on that task.

4 Experiments

4.1 Pretraining Datasets

To pretrain MetaRetriever, we collect a large-scale corpus using CROCODILE (Cabot and Navigli, 2021), an automatic relation extraction dataset construction tool. CROCODILE is based on the distant supervision technique (Mintz et al., 2009) which aligns texts from English Wikipedia² and knowledge base from Wikidata³ to create distant labeled relation extraction data. Since distant supervision inevitably introduces noise (Mintz et al., 2009; Cabot and Navigli, 2021), CROCODILE further uses Natural Language Inference (NLI) to filter data with low confidence. Finally, we construct a 71 million distant labeled pretraining corpus and obtained 6.9 million high-quality pretraining corpus after NLI filtering. Detailed statistics of our pretraining corpus are listed in Appendix A.

Our analysis experiment (Section 5.3) indicates that the pretraining corpus we constructed can yield comparable results in the trained model as the data constructed by UIE. As UIE’s pretraining data is not open-sourced, **our data will be the first 10-million-level corpus to be open-sourced, and will be released after the acceptance of the paper.**

4.2 Evaluation Datasets

Following previous work (Lu et al., 2022), we conduct experiments across 4 IE tasks: named entity recognition (NER), relational triple extraction (RTE), event extraction (EE), structured sentiment extraction (Senti). Specifically, we use 12 IE benchmarks in total for these 4 IE tasks⁴: ACE04 (Mitchell et al., 2005), ACE05 (Walker et al., 2006); CoNLL03 (Tjong Kim Sang and De Meulder, 2003), CoNLL04 (Roth and Yih, 2004), SciERC (Luan et al., 2018), CASIE (Satya-

²<https://www.wikipedia.org/>

³<https://www.wikidata.org/>

⁴As NYT datasets overlaps with pre-training data, we didn’t conduct on it for fair comparison.

Task	Dataset	UIE	MetaRetriever
NER	ACE04	85.69	86.10 (0.41 \uparrow)
	ACE05-Ent	83.88	84.01 (0.13 \uparrow)
	CoNLL03	91.94	92.38 (0.44 \uparrow)
RTE	ACE05-Rel	62.73	64.37 (1.64 \uparrow)
	CoNLL04	73.48	73.66 (0.18 \uparrow)
	SciERC	35.35	35.77 (0.42 \uparrow)
EE Trg.	ACE05-Evt	71.33	72.38 (1.05 \uparrow)
	CASIE	69.14	69.76 (0.62 \uparrow)
EE Arg.	ACE05-Evt	50.62	52.62 (2.00 \uparrow)
	CASIE	58.56	60.37 (1.81 \uparrow)
Senti	14-res	72.86	73.41 (0.55 \uparrow)
	14-lap	62.68	62.83 (0.15 \uparrow)
	15-res	65.51	65.85 (0.34 \uparrow)
	16-res	73.26	73.55 (0.29 \uparrow)

Table 1: Overall results on 12 datasets in the fully supervised settings. The performance of UIE are reported based on the official UIE-en-base model.

panich et al., 2020), SemEval-14 (Pontiki et al., 2014), SemEval-15 (Pontiki et al., 2015), SemEval-16 (Pontiki et al., 2016). The detailed statistics of these datasets can be seen in Appendix D.

4.3 Evaluation Metrics

We use span-based offset Micro-F1 as the primary metric to evaluate the model for different IE tasks:

- (1) **NER**: an entity is correct if its offsets and type are correct.
- (2) **RTE**: a relation is correct if its relation type is correct and the offsets and entity types of the related head/tail entity are correct.
- (3) **EE Trg.**: an event trigger is correct if its offsets and event type are correct.
- (4) **EE Arg.**: an event argument is correct if its offsets, role type, and event type are correct.
- (5) **Senti**: a sentiment triple is correct if the offsets boundary of the target, the offsets boundary of the opinion span, and the target sentiment polarity are correct.

After generation, we reconstruct the offset of predicted information pieces by finding the matched utterance in the input sequence one by one.

To validate the effectiveness of our proposed method, we compare MetaRetriever with the state-of-the-art model UIE (Lu et al., 2022) in fully-supervised, few-shot and low-resource scenarios.

Task	Model	Few-Shot			Low-Resource		
		1-Shot	5-Shot	10-Shot	1%	5%	10%
NER (CoNLL03)	UIE	46.43	67.09	73.90	82.84	88.34	89.63
	MetaRetriever	49.44	69.88	74.19	83.45	89.11	90.42
	Δ	3.01 \uparrow	2.79 \uparrow	0.29 \uparrow	0.61 \uparrow	0.77 \uparrow	0.79 \uparrow
RTE (CoNLL04)	UIE	22.05	45.41	52.39	30.77	51.72	59.18
	MetaRetriever	29.90	47.02	53.95	36.31	52.59	59.45
	Δ	7.85 \uparrow	1.61 \uparrow	1.56 \uparrow	5.54 \uparrow	0.87 \uparrow	0.27 \uparrow
Event Trigger (ACE05-Evt)	UIE	38.14	51.21	53.23	41.53	55.70	60.29
	MetaRetriever	39.85	49.43	53.58	43.98	58.35	63.19
	Δ	1.71 \uparrow	1.78 \downarrow	0.35 \uparrow	2.45 \uparrow	2.65 \uparrow	2.90 \uparrow
Event Argument (ACE05-Evt)	UIE	11.88	27.44	33.64	12.80	30.43	36.28
	MetaRetriever	13.30	27.70	32.31	14.86	31.85	37.85
	Δ	1.42 \uparrow	0.26 \uparrow	1.33 \downarrow	2.06 \uparrow	1.42 \uparrow	1.57 \uparrow
Sentiment (16res)	UIE	10.50	26.24	39.11	24.24	49.31	57.61
	MetaRetriever	18.80	34.14	43.53	34.47	51.78	58.79
	Δ	8.30 \uparrow	7.90 \uparrow	4.42 \uparrow	10.23 \uparrow	2.47 \uparrow	1.18 \uparrow

Table 2: Few-Shot and Low-resource results. Results of UIE are reported from original paper (Lu et al., 2022).

4.4 Main Results

Table 1 and Table 2 summarize the results of our proposed MetaRetriever compared to the previous work, UIE, on 4 types of IE tasks. From the experimental results, we have several observations (1) **Equipped with task-specific knowledge, MetaRetriever outperforms UIE on all fully-supervised, few-shot, and low-resource scenarios.** MetaRetriever surpasses UIE with an average of 0.7%, 2.5%, and 2.3% in F1 score respectively in the fully-supervised, few-shot, and low-resource settings. This strongly demonstrates that MetaRetriever can utilize the retrieved knowledge to improve IE performance. (2) **Empowered by meta-pretraining, MetaRetriever can effectively learn downstream task-specific extraction targets.** Despite the fact that event extraction and sentiment extraction were not present in the pretraining phase, MetaRetriever still outperforms UIE in these tasks, validating that MetaRetriever can quickly learn downstream IE tasks through the meta-pretraining algorithm to support task-specific extraction. (3) **Benefiting from task-specific knowledge and meta-pretraining, MetaRetriever exhibits significant superior performance in data-scarce scenarios.** MetaRetriever achieves large gains in the few-shot and low-resource settings. This is attributed to two factors: first, the meta-pretraining algorithm can

quickly learn downstream IE tasks, which reduces the requirement for training data, and second, the retrieved task-specific knowledge provides crucial clues for the model to make correct predictions, thus MetaRetriever can achieve impressive performance in data-scarce scenarios.

5 Analysis

5.1 Effect of Meta-Pretraining Algorithm

	CoNLL03			CoNLL04		
	Sup.	Few.	Low.	Sup.	Few.	Low.
MetaRetriever	92.38	49.44	83.45	73.66	29.90	36.31
SimpleRetriever	92.14	43.72	82.66	73.33	24.47	32.43

Table 3: Effect of Meta-Pretraining Algorithm.

To examine the contribution of our Meta-Pretraining Algorithm, we remove the inner loop in Algorithm 1 and pretrain our model using only the outer loop. Without the inner loop, the bi-level optimization is reduced to single-level optimization. We name this pretrained model SimpleRetriever and compare its performance with MetaRetriever on CoNLL03 and CoNLL04 datasets in fully-supervised, few-shot(1-shot) and low-resource(1%) settings. The results are presented in Table 3 and it can be observed that: Compared to MetaRetriever, the performance of SimpleRetriever declines on all datasets in all settings. Notably, in data-scarce

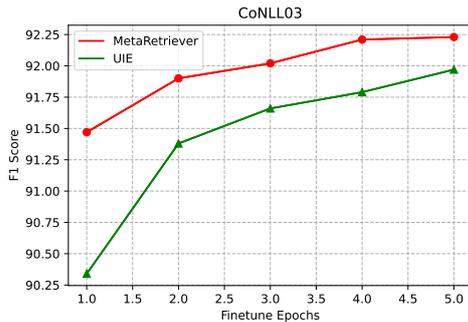


Figure 3: Effect of Finetune Epoch.

scenarios, SimpleRetriever underperforms MetaRetriever by 3.92% on average. This demonstrates that the meta-pretraining algorithm enables the model to quickly learn downstream IE tasks and thus improve performance.

5.2 Effect of Finetune Epoch

In order to demonstrate the efficiency of our MetaRetriever in rapidly learning downstream information extraction tasks, we conducted a series of experiments. Specifically, we compare the performance of MetaRetriever to that of UIE across a range of finetuning epochs, from 1 to 5 on CoNLL03 dataset. The results of these experiments, presented in Figure 3, clearly indicate that MetaRetriever outperforms UIE from the very first finetuning epoch. Furthermore, as the number of finetuning epochs increases, the performance of MetaRetriever consistently surpasses that of UIE, with the greatest difference observed when comparing MetaRetriever’s performance in 2 finetuning epochs to UIE’s performance in 5 finetuning epochs. These results strongly suggest that MetaRetriever’s ability to rapidly learn downstream tasks is significantly enhanced by its use of meta-pretraining. Additionally, we also conducted experiments on the CoNLL04 datasets, the results of which are provided in Appendix E.

5.3 Effect of Structure Complexity

In order to determine the effectiveness of MetaRetriever in dealing with complex structures when compared to UIE, we conduct an experiment using the test set of CoNLL03, grouping the test instances according to the number of entities present. We report the performance gain of MetaRetriever compared with UIE in different group in Figure 4. We can observe that, as the number of entities present in a test instance increases, MetaRetriever’s per-

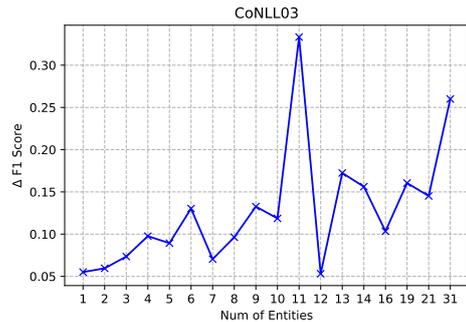


Figure 4: Effect of Structure Complexity.

formance advantage over UIE also increases. This demonstrates that MetaRetriever is able to effectively leverage task-specific knowledge to handle complex structures and improve information extraction performance. We also conduct further experiments on the CoNLL04 dataset, the results of which can be found in Appendix F.

5.4 Effect of Pretraining Corpus Collection

	UIE	UIE-ours
CoNLL03	91.94	91.96
CoNLL04	73.48	71.60
ACE05 Trg.	71.33	73.39
ACE05 Arg.	50.62	51.87
16res	73.26	72.16
AVG	72.13	72.20

Table 4: Effect of Pretraining Corpus Collection.

To evaluate the quality of the corpus we have collected for pretraining, we use this corpus to pre-train a variant of UIE which we refer to as UIE-ours. We then compared UIE-ours to the original UIE on four representative datasets for each IE task in fully-supervised settings. The results, presented in Table 4, show that UIE-ours, which is pretrained on our collected corpus, is comparable in performance to the original UIE (with an average gain of 0.07%). This indicates that our collected corpus is of similar quality to that used for training the original UIE and we plan to make it publicly available to promote future research. Additionally, this also demonstrates that the improvement of MetaRetriever is due to our proposed method rather than the corpus used for pretraining, further validating the effectiveness of our proposed method.

<i>(1) Gale-force winds were reported in California late Wednesday, where gusts reached 69 mph at Ontario and 57 mph at El Toro.</i>	
UIE	((LOC: California)(LOC: Ontario)(LOC: El Toro))
MetaRetriever	((LOC: California)(LOC: Ontario(Located_In: California))(LOC: El Toro(Located_In: California)))
RETRIEVED KNOWLEDGE	((LOC: California)(LOC: Ontario)(LOC: El Toro(Located_In: California)))
<i>(2) Alekseev, along with others in the Soviet diplomatic corps, were not receiving accurate information from Moscow.</i>	
UIE	(PER: Alekseev (Live_In: Soviet)) (LOC: Soviet) (LOC: Moscow)
MetaRetriever	(PER: Alekseev (Live_In: Soviet)) (LOC: Soviet) (LOC: Moscow(Located_In: Soviet))
RETRIEVED KNOWLEDGE	(PER: Alekseev (Live_In: Soviet)) (LOC: Soviet) (LOC: Moscow(Located_In: Soviet))

Table 5: Case Study.

5.5 Effect of Support-Query Construction for Meta-Pretraining

Conventional meta-learning methods typically employ the episodic training technique to construct "tasks" for meta-training (Vinyals et al., 2016; Finn et al., 2017). This technique creates "tasks" by randomly sampling instances based on their category, which requires random access to all data. However, since large-scale corpora cannot be stored in memory and must be accessed from external storage instead, random-sampling operation will cause long IO time. To address this, we design a support-query pairing method which can significantly reduce IO times. To illustrate the efficiency of our method, we compared the time cost of our method with episodic training on 10K instances. Episodic training required 100.42s to construct "IE tasks" while our method only required 5.57s, demonstrating the significant efficiency of our method.

5.6 Case Study

In this study, we compare the performance of our MetaRetriever model with that of the UIE model using a selection of cases from the CoNLL04 dataset, as illustrated in Table 5.

In the first case, our MetaRetriever, aided by the retrieved task-specific knowledge, "(LOC: El Toro (Located_In: California))", correctly predict "(LOC: Ontario (Located_In: California))". In contrast, the UIE model was unable to make a precise prediction without the benefit of this task-specific knowledge. This demonstrates that models can significantly benefit from the incorporation of retrieved task-specific knowledge to generate more accurate linearized SEL expressions.

In the second case, we observe that MetaRetriever makes a redundant prediction, "(LOC: Moscow (Located_In: Soviet))". Though this relational triple is not labeled in the CoNLL04 dataset, it is a common sense knowledge that Moscow is located in Russia, which is former So-

viet Union. MetaRetriever was able to predict this based on the background knowledge acquired from the pretraining corpus.

6 Related Work

Information Extraction (IE) aims at extracting text spans or a tuple of text spans of interests from plain texts. There exist many specific IE tasks: Named Entity Recognition (Lample et al., 2016; Lin et al., 2019), Relational Triple Extraction (Zheng et al., 2017; Levy et al., 2017), Event Extraction (Lin et al., 2018; Wadden et al., 2019; Du and Cardie, 2020; Li et al., 2022c,b,a), etc. For a long time, researchers are devoted to propose task-customized and isolated methods to achieve these different IE tasks. In the recent years, with the rising of pre-training technique, pretraining a universal model for several NLP tasks has attracted a lot of attention (Devlin et al., 2019; Liu et al., 2019; Lewis et al., 2020a; Raffel et al., 2020b; Xue et al., 2021). Following this trend, several works make attempts to unify diverse IE tasks. Yan et al. (2021b) first propose a span-offset-based generation to solve various NER tasks in a universal manner. Yan et al. (2021a) use a similar way to solve different aspect-based sentiment analysis tasks. Lu et al. (2022) propose UIE which designs Structured Extraction Language to formulate all IE tasks in a unified form but it will experience poor performance when there exists complex structure to be extracted. Several works have shown that retrieving knowledge from external knowledge bases can empower models to generate complex sequences (Lewis et al., 2020b; Wang et al., 2021; Cai et al., 2022; Geng et al., 2022) and a large number of knowledge stored in pretrained language models can be "retrieved" explicitly (Petroni et al., 2019; Roberts et al., 2020; Zhang et al., 2022; Yu et al., 2022). In this paper, we propose MetaRetriever to retrieve task-specific knowledge from PLMs to enhance universal IE.

Meta-learning aims to learn better learning algorithms from a series of tasks, i.e., learning to learn.

In meta-learning, there exist diverse branches focusing on different aspects of the learning algorithm: learning-to-initialize (Finn et al., 2017, 2018), learning-to-compare (Vinyals et al., 2016; Snell et al., 2017), learning-to-optimize (Ravi and Larochelle, 2017), etc. Existing works are most devoted to solve the few-shot problem, cross-domain problem, etc. We make the first attempt to adapt meta-learning into universal IE.

7 Conclusion

In this work, we present MetaRetriever, a novel approach for enhancing universal information extraction through the retrieval of task-specific knowledge from within the model itself. To ensure that the retrieved knowledge is truly task-specific and to minimize the impact of irrelevant or "noise" knowledge, we also propose a meta-pretraining algorithm that enables MetaRetriever to quickly adapt to downstream information extraction tasks. Empirical results demonstrate the effectiveness of our approach, with MetaRetriever achieving new state-of-the-art performance on a range of information extraction tasks across 12 datasets in fully-supervised, low-resource, and few-shot scenarios. Further analysis experiments also validate the contributions and effectiveness of MetaRetriever.

8 Limitations

While our MetaRetriever has demonstrated its superior performance on 4 IE tasks, 12 datasets, in fully-supervised, few-shot and low-resource scenarios, it still has several limitations. First, since MetaRetriever will first retrieve task-specific knowledge and then make predictions in the inference phase, such a retrieve-then-extract manner will take longer inference time than non-retrieve methods unavoidably. Second, our proposed meta-pretraining algorithm is based on bi-level optimization. In the pretraining phase, it needs to calculate high-order gradients to optimize parameters and calculating high-order gradient requires more time, which will cause longer time to pretrain MetaRetriever. Detailed discussions are shown in Appendix G.

9 Ethical Considerations

This paper constructs a new dataset for universal information extraction, and we discuss some related ethical considerations here. (1) **Intellectual property.** The Wikipedia corpus is shared under

the CC BY-SA 3.0 license⁵ and Wikidata is shared under the CC0 1.0 license⁶. CROCODILE is licensed under the CC BY-SA-NC 4.0 license⁷. Our data source and the construct tool are all free for research use. (2) **Controlling Potential Risks.** Since the texts in Wikipedia do not involve private information and annotating entities and relations does not require many judgments about social issues, we believe our collected dataset does not create additional risks. To ensure it, we manually checked some randomly sampled data and did not note risky issues. (3) **Worker Treatments.** Since the dataset construction procedure is automatic without any manual annotation, it does not create potential risk about the worker treatments.

Acknowledgements

This work is supported by the National Key Research and Development Program of China (grant No.2021YFB3100600), the Strategic Priority Research Program of Chinese Academy of Sciences (grant No.XDC02040400) and the Youth Innovation Promotion Association of CAS (Grant No. 2021153).

References

- Peggy M. Andersen, Philip J. Hayes, Steven P. Weinstein, Alison K. Huettner, Linda M. Schmandt, and Irene B. Nirenburg. 1992. [Automatic extraction of facts from press releases to generate news stories](#). In *Third Conference on Applied Natural Language Processing*, pages 170–177, Trento, Italy. Association for Computational Linguistics.
- Pere-Lluís Hugué Cabot and Roberto Navigli. 2021. [REBEL: relation extraction by end-to-end language generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*.
- Deng Cai, Yan Wang, Lemao Liu, and Shuming Shi. 2022. Recent advances in retrieval-augmented text generation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 3417–3419.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for*

⁵<https://creativecommons.org/licenses/by-sa/3.0/>

⁶<https://creativecommons.org/publicdomain/zero/1.0/>

⁷<https://creativecommons.org/licenses/by-nc-sa/4.0/>

- Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xinya Du and Claire Cardie. 2020. [Event extraction by answering \(almost\) natural questions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 671–683, Online. Association for Computational Linguistics.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. [Model-agnostic meta-learning for fast adaptation of deep networks](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*.
- Chelsea Finn, Kelvin Xu, and Sergey Levine. 2018. [Probabilistic model-agnostic meta-learning](#). In *Conference on Neural Information Processing Systems*.
- Zvi Galil. 1986. [Efficient algorithms for finding maximum matching in graphs](#). *ACM Comput. Surv.*, 18(1):23–38.
- Zhichao Geng, Hang Yan, Zhangyue Yin, Chenxin An, and Xipeng Qiu. 2022. [TURNER: the uncertainty-based retrieval framework for chinese NER](#). *CoRR*, abs/2202.09022.
- Ralph Grishman. 2019. [Twenty-five years of information extraction](#). *Natural Language Engineering*, 25(6):677–692.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *The Third International Conference on Learning Representations*, San Diego.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. [Neural architectures for named entity recognition](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270, San Diego, California. Association for Computational Linguistics.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. 2017. [Zero-shot relation extraction via reading comprehension](#). In *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*, pages 333–342, Vancouver, Canada. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020a. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020b. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Qian Li, Jianxin Li, Jiawei Sheng, Shiyao Cui, Jia Wu, Yiming Hei, Hao Peng, Shu Guo, Lihong Wang, Amin Beheshti, et al. 2022a. [A survey on deep learning event extraction: Approaches and applications](#). *IEEE Transactions on Neural Networks and Learning Systems*.
- Qian Li, Jianxin Li, Lihong Wang, Cheng Ji, Yiming Hei, Jiawei Sheng, Qingyun Sun, Shan Xue, and Pengtao Xie. 2022b. [Type information utilized event detection via multi-channel gms in electrical power systems](#). *CoRR*, abs/2211.08168.
- Qian Li, Hao Peng, Jianxin Li, Jia Wu, Yuanxing Ning, Lihong Wang, Philip S. Yu, and Zheng Wang. 2022c. [Reinforcement learning-based dialogue guided event extraction to exploit argument relations](#). *IEEE ACM Trans. Audio Speech Lang. Process.*
- Hongyu Lin, Yaojie Lu, Xianpei Han, and Le Sun. 2018. [Nugget proposal networks for Chinese event detection](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1565–1574, Melbourne, Australia. Association for Computational Linguistics.
- Hongyu Lin, Yaojie Lu, Xianpei Han, and Le Sun. 2019. [Sequence-to-nuggets: Nested entity mention detection via anchor-region networks](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5182–5192, Florence, Italy. Association for Computational Linguistics.
- 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](#). *CoRR*, abs/1907.11692.
- Yaojie Lu, Qing Liu, Dai Dai, Xinyan Xiao, Hongyu Lin, Xianpei Han, Le Sun, and Hua Wu. 2022. [Unified structure generation for universal information extraction](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2022, Dublin, Ireland, May 22-27, 2022.
- Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. [Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3219–3232, Brussels, Belgium. Association for Computational Linguistics.
- Mike Mintz, Steven Bills, Rion Snow, and Daniel Jurafsky. 2009. [Distant supervision for relation extraction without labeled data](#). In *ACL 2009, Proceedings of*

- the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing of the AFNLP, 2-7 August 2009, Singapore.*
- Alexis Mitchell, Stephanie Strassel, Shudong Huang, and Ramez Zakhary. 2005. [Ace 2004 multilingual training corpus](#).
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick S. H. Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander H. Miller. 2019. [Language models as knowledge bases?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Ion Androutsopoulos, Suresh Manandhar, Mohammad AL-Smadi, Mahmoud Al-Ayyoub, Yanyan Zhao, Bing Qin, Orphée De Clercq, Véronique Hoste, Marianna Apidianaki, Xavier Tannier, Natalia Loukachevitch, Evgeniy Kotelnikov, Nuria Bel, Salud María Jiménez-Zafra, and Gülşen Eryiğit. 2016. [SemEval-2016 task 5: Aspect based sentiment analysis](#). In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, pages 19–30, San Diego, California. Association for Computational Linguistics.
- Maria Pontiki, Dimitris Galanis, Haris Papageorgiou, Suresh Manandhar, and Ion Androutsopoulos. 2015. [SemEval-2015 task 12: Aspect based sentiment analysis](#). In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pages 486–495, Denver, Colorado. Association for Computational Linguistics.
- Maria Pontiki, Dimitris Galanis, John Pavlopoulos, Haris Papageorgiou, Ion Androutsopoulos, and Suresh Manandhar. 2014. [SemEval-2014 task 4: Aspect based sentiment analysis](#). In *Proceedings of the 8th International Workshop on Semantic Evaluation (SemEval 2014)*, pages 27–35, Dublin, Ireland. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020a. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020b. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Sachin Ravi and Hugo Larochelle. 2017. [Optimization as a model for few-shot learning](#). In *International Conference on Learning Representations*.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. [How much knowledge can you pack into the parameters of a language model?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*.
- Dan Roth and Wen-tau Yih. 2004. [A linear programming formulation for global inference in natural language tasks](#). In *Proceedings of the Eighth Conference on Computational Natural Language Learning (CoNLL-2004) at HLT-NAACL 2004*, pages 1–8, Boston, Massachusetts, USA. Association for Computational Linguistics.
- David Saad. 1998. Online algorithms and stochastic approximations. *Online Learning*, 5(3):6.
- Taneeya Satyapanich, Francis Ferraro, and Tim Finin. 2020. [Casie: Extracting cybersecurity event information from text](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8749–8757.
- Jake Snell, Kevin Swersky, and Richard S. Zemel. 2017. [Prototypical networks for few-shot learning](#). In *Conference on Neural Information Processing Systems*.
- Ilya Sutskever, James Martens, George E. Dahl, and Geoffrey E. Hinton. 2013. [On the importance of initialization and momentum in deep learning](#). In *Proceedings of the 30th International Conference on Machine Learning*.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the conll-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of CoNLL-2003*, pages 142–147. Edmonton, Canada.
- Oriol Vinyals, Charles Blundell, Tim Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. 2016. [Matching networks for one shot learning](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*.
- David Wadden, Ulme Wennberg, Yi Luan, and Hananeh Hajishirzi. 2019. [Entity, relation, and event extraction with contextualized span representations](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5784–5789, Hong Kong, China. Association for Computational Linguistics.
- Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. 2006. [Ace 2005 multilingual training corpus](#).
- Xinyu Wang, Yong Jiang, Nguyen Bach, Tao Wang, Zhongqiang Huang, Fei Huang, and Kewei Tu. 2021. [Improving named entity recognition by external context retrieving and cooperative learning](#). In *Proceedings of the 59th Annual Meeting of the Association for*

Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021.

Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.

Hang Yan, Junqi Dai, Tuo Ji, Xipeng Qiu, and Zheng Zhang. 2021a. [A unified generative framework for aspect-based sentiment analysis](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021.*

Hang Yan, Tao Gui, Junqi Dai, Qipeng Guo, Zheng Zhang, and Xipeng Qiu. 2021b. [A unified generative framework for various NER subtasks](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021.*

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*.

Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2022. Automatic chain of thought prompting in large language models. *arXiv preprint arXiv:2210.03493*.

Suncong Zheng, Feng Wang, Hongyun Bao, Yuexing Hao, Peng Zhou, and Bo Xu. 2017. [Joint extraction of entities and relations based on a novel tagging scheme](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1227–1236, Vancouver, Canada. Association for Computational Linguistics.

A Details of Pretraining Corpus

We apply CROCODILE⁸ (Cabot and Navigli, 2021), which is a automatic distant supervised labeling tool developed for relation extraction, to collect our pretraining corpus. CROCODILE utilizes widely-used distant supervision technique (Mintz et al., 2009) to build relation extraction dataset by

⁸<https://github.com/Babelscape/crocodile>

# Sentence	71,947,589
# Entity	208,261,778
# Triple	27,011,527
# Entity Types	45,062
# Relation Types	1,465
Average Length	33.43

Table 6: Statistics of Raw Pretraining Corpus.

# Sentence	6,951,695
# Entity	48,352,090
# Triple	12,497,480
# Entity Types	39,511
# Relation Types	1,255
Average Length	47.47

Table 7: Statistics of Filtered Pretraining Corpus.

aligning English Wikipedia⁹ texts and Wikidata¹⁰ knowledge base. Additionally, CROCODILE applies RoBERTa NLI model¹¹ to filter distant labeled data with low-confidence. Finally, we construct 71M original pretraining corpus without NLI filtering and obtain 6.9M high-quality pretraining corpus with NLI filtering. The statistics of them are shown in Table 6 and Table 7 respectively. Since our collected corpus contains the 10-million-level sentences covering about 40K entity types and over 1K relation types, we believe the collected corpus is diverse enough in genre and topic, which exhibits a good coverage for general domain. Both of these pretraining corpus will be public after acceptance.

B Details of Meta-Pretraining Dataset Construction Procedure

The goal of our meta-pretraining is to train MetaRetriever that can quickly learn the extraction target semantics of various downstream IE tasks for retrieval using only a few training epochs. To accomplish this, MetaRetriever is pretrained on a set of “IE tasks” based on the meta-learning technique (Finn et al., 2017), such that the trained model can quickly learn new IE tasks using only a small number of trials. To construct various “IE tasks”, a intuitive way is to follow the *episodic training* technique (Vinyals et al., 2016) which are widely used in conventional meta-learning methods. In *episodic training*, it randomly samples N classes

⁹<https://www.wikipedia.org/>

¹⁰<https://www.wikidata.org/>

¹¹[xlm-roberta-large-xnli](https://huggingface.co/roberta-large-xnli)

from original class set as a N -classification “task”. For each N -classification task, it will randomly samples K instances as the support set to mimic the training set for this task and then use remain instances as the query set to mimic the test set, i.e., the N -way- K -shot setting. In each training epoch, it will sample different N -way- K -shot tasks to train models. However, when pretraining a model, it is time-consuming and memory-consuming to sample classes from a large-scale corpus. To meta-pretrain MetaRetriever, we design Support-Query Pairing Algorithm based on the graph maximum-weighted matching algorithm (Galil, 1986). Our Support-Query Pairing Algorithm can be divided into three steps: (1) Partition by class; (2) Deduplication; (3) Graph Matching.

Partition by class Different from *episodic training* which treats each a support set and a query set as a task, we construct a “task” with only two instances. One acts as the support instance and the other as the query instance, i.e., the size of our support set and query are both 1. To make the support instance “support” the query instance, their classes should be as nearly identical as possible so the model can learn the semantics of the target classes from the support instance to make predictions on the query instance. To achieve this goal, we partition all pretraining instances based on the class. Specifically, we denote the raw training dataset $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ where x_i refers to the raw text, y_i refers to the golden linearized SEL expression, and N refers to the number of the pretraining instances. Given the training corpus \mathcal{D} , we can get the class set \mathbb{C} which contains all SpotNames and AssoNames of \mathcal{D} . Then, we collect instances for each classes \mathcal{D}_c , $c \in \mathbb{C}$ where \mathcal{D}_c refers to all instances which contains c -type target information piece. Therefore, we can get the partitioned dataset $\mathcal{D}_{\mathbb{C}}$.

Deduplication Getting the partitioned dataset $\mathcal{D}_{\mathbb{C}}$, there exists serious duplication phenomenon because one instance will contain several classes and thus it will be partitioned into multiple subset \mathcal{D}_c . Such a duplication will cause a severe data imbalance since if one instance contains m classes, it will be duplicated m times, hurting the model performance. To alleviate this problem, we should deduplicate the partitioned dataset $\mathcal{D}_{\mathbb{C}}$. Specifically, we sort the class set \mathbb{C} by the number of instances of each class. From less to more, we re-

move these instances from the large subset which have existed in the small subset. After that, we get the deduplicated partitioned dataset $\hat{\mathcal{D}}_{\mathbb{C}}$.

Graph Matching Since we treat a support-query instance pair as a task, we want to find a optimal pairing solution that the classes of the support instance could cover the classes of the query instance as many as possible globally. To achieve this goal, we first define a pairing score function to evaluate a support-query pair. Given a support instance s and a query instance q , we define the pairing score function as follows,

$$\rho(s, q) = \frac{|F_{\mathbb{C}}(s) \cap F_{\mathbb{C}}(q)|}{|F_{\mathbb{C}}(s)|} + \frac{1}{|F_{\mathbb{C}}(s)|} \quad (13)$$

where $F_{\mathbb{C}}(\cdot)$ means the class set contained in the given instance. In this score function, the first term evaluates the ratio how the support instance covers the query instance. The second term evaluates the number of the classes of the support instance. We wish the support instance should cover the query instance exactly so the number of the classes of the support instance should be as small as possible besides the class coverage. Then, we define a matching score function to calculate the pairing degree given any two instances x and y :

$$\varphi(x, y) = \max \{\rho(x, y), \rho(y, x)\} \quad (14)$$

Given a subset \mathcal{D}_c , $c \in \mathbb{C}$, we can construct a weighted undirected complete graph \mathcal{G}_c . In \mathcal{G}_c , each instance in \mathcal{D}_c acts as a node and the arbitrary two nodes x and y have a edge with $\varphi(x, y)$ weight. Next, we run the maximum-weighted graph matching algorithm (Galil, 1986) to get a graph matching result \mathcal{M}_c . In \mathcal{M}_c , each node matching pair represents a support-query pair. We run the aforementioned matching procedure for each class $c \in \mathbb{C}$ to get all matching pairs $\mathcal{M}_{\mathbb{C}} = \{\mathcal{M}_c\}$, $c \in \mathbb{C}$. We use $\mathcal{M}_{\mathbb{C}}$ to meta-pretrain our MetaRetriever.

Through the procedure introduced above, the I/O times of support-query pairing are 2. First, it reads all data to get the categories of each instance to run the support-query pairing algorithm. Second, given the pairing results, it takes the second I/O time to re-range all data to construct support-query pair. For the episodic training, every random-sampling will lead to a I/O time while our method can reduce I/O times significantly.

C Details of Model Implementation

Since MetaRetriever works in a sequence-to-sequence manner, we use the encoder-decoder-style

	Entl	Rel	Evtl	#Train	#Val	#Test
ACE04	7	-	-	6,202	745	812
ACE05-Ent	7	-	-	7,299	971	1,060
CoNLL03	4	-	-	14,041	3,250	3,453
ACE05-Rel	7	6	-	10,051	2,420	2,050
CoNLL04	4	5	-	922	231	288
SciERC	6	7	-	1,861	275	551
ACE05-Evt	-	-	33	19,216	901	676
CASIE	21	-	5	11,189	1,778	3,208
14res	2	3	-	1,266	310	492
14lap	2	3	-	906	219	328
15res	2	3	-	605	148	322
16res	2	3	-	857	210	326

Table 8: Detailed datasets statistics of all downstream IE tasks. |*| indicates the number of categories, and # is the number of sentences in the specific subset.

architecture to achieve this procedure. We choose T5-base architecture (Raffel et al., 2020a) as our model with 220M parameters. We pretrain our model on our constructed pretraining dataset. We use Adam (Kingma and Ba, 2015) as the optimizer with learning rate $1e-4$ with the linear scheduling with a warming up proportion 6%. As previous work (Lu et al., 2022), we randomly sample 10 SpotName and AssoName to preserve the generalization ability of the model when calculating $\mathcal{L}_{\text{retrv}}$ and \mathcal{L}_{gen} . For \mathcal{L}_{LM} , we set the corruption rate as 15% and the average corrupting span length as 3 as T5 (Raffel et al., 2020b). We truncate the length of the input sequence to 128 during pretraining. In the pretraining phase, we set the step of the inner loop as 1 and the inner step size α and the outer step size β are set as $1e-4$. The inner-loop update step J is set as 1. We adopt the filtered version of our pretraining corpus (6.9M instances) to pretrain MetaRetriever. We pretrain our model for both 4 epochs with batch size 512 on 8 NVIDIA A100 80G GPUs with 56 hours. For the downstream fine-tuning, we apply the same grid searching strategy for hyper-parameters as UIE (Lu et al., 2022) to find the optimal results. In the inference phase, we use the greedy search strategy to decode the final predicted SEL sequence.

D Details of Downstream Task Datasets

We conduct evaluation experiments on 4 IE tasks, 12 datasets, and the detailed statistic of each dataset is shown in Table 8.

E Details of Effect of Finetune Epoch

To verify that our MetaRetriever can fast learn the downstream IE tasks, we conduct experiments on both CoNLL03 and CoNLL04. Experimental results are shown in Figure 5. Results of CoNLL04 show the same trend as CoNLL03, which can prove the effect of our proposed Meta-Pretraining Algorithm.

F Details of Effect of Structure Complexity

To validate that MetaRetriever own the superiority when there exist complicated structures needed to be extracted. We group the test set of CoNLL03 and CoNLL04 according to the number of entities and the number of relational triples and calculate the performance difference between MetaRetriever and UIE in different group. Experimental results are shown in Figure 6. From it, we can learn that when there exist more entities and relational triples, MetaRetriever performs better.

G Discussion of Efficiency

While our MetaRetriever has achieved the new state-of-the-art on 4 IE tasks, 12 datasets, in fully-supervised, few-shot and low-resource scenarios, it still has several limitations on training efficiency and inference efficiency.

	UIE	MetaRetriever
CoNLL04	0.83s	1.46s

Table 9: Inference time of MetaRetriever and UIE on the test set of CoNLL04 dataset.

First, since MetaRetriever will first retrieve task-specific knowledge and then make predictions in the inference phase, such a retrieve-then-extract manner will take longer inference time than non-retrieve methods unavoidably. We conduct experiments on the test set of CoNLL04 dataset to compare overall inference time of MetaRetriever with UIE. All hyper-parameters are set to be the same for a fair comparison. Experimental results are shown in Table 9 and we can find that MetaRetriever cost nearly twice time as UIE to make predictions. As MetaRetriever works in a retrieve-than-extract manner, such a time cost is reasonable.

Second, our proposed meta-pretraining algorithm is based on bi-level optimization. In the

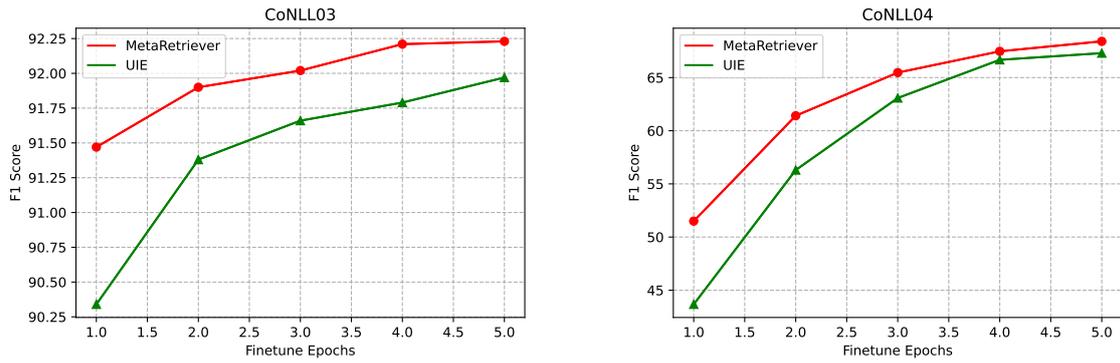


Figure 5: Effect of Finetune Epoch.

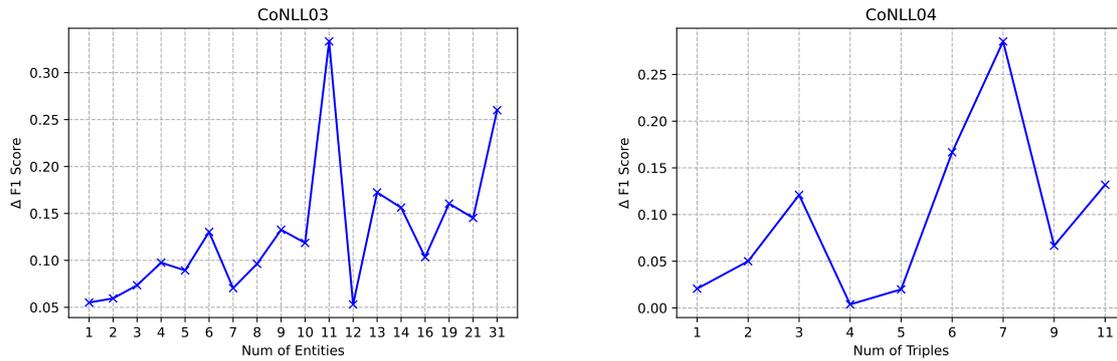


Figure 6: Effect of Structure Complexity.

	SimpleRetriever	MetaRetriever
10K	2min	2min25s

Table 10: Pretraining time of MetaRetriever and SimpleRetriever on 10K instances for one epoch.

pretraining phase, it needs to calculate high-order gradients to optimize parameters and calculating high-order gradient requires more time. Therefore, it takes longer time to pretrain MetaRetriever. To illustrate the time cost, we conduct experiments on 10K instances to compare the pretraining time of MetaRetriever with SimpleRetriever which is pretrained without meta-pretraining algorithm. All hyper-parameters are set to be the same for a fair comparison. Table 10 gives the results. From it, we can observe that compared with SimpleRetriever, MetaRetriever takes about 1/4 longer time than SimpleRetriever. Finally, We spent 56 hours to pre-train MetaRetriever on filtered pretraining corpus (6.9M instances).

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section 8
- A2. Did you discuss any potential risks of your work?
Section 9
- A3. Do the abstract and introduction summarize the paper’s main claims?
Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 4 and Appendix A

- B1. Did you cite the creators of artifacts you used?
Section 4 and Appendix A
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Section 9
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Not applicable. Section 9 and Appendix A
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Section 9
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Appendix A
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Appendix A

C Did you run computational experiments?

Appendix C

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Appendix C

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Appendix C

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Appendix C

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Section 4 and Appendix A

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.