

# Towards Human-Like Machine Comprehension: Few-Shot Relational Learning in Visually-Rich Documents

Hao Wang<sup>1</sup>, Tang Li<sup>1</sup>, Chenhui Chu<sup>2</sup>, Nengjun Zhu<sup>1</sup>, Rui Wang<sup>3,\*</sup> and Pinpin Zhu<sup>1</sup>

<sup>1</sup> School of Computer Engineering and Science, Shanghai University, China

<sup>2</sup> Graduate School of Informatics, Kyoto University, Japan

<sup>3</sup> Department of Computer Science and Engineering, Shanghai Jiao Tong University, China  
{wang-hao, lttttt, zhu\_nj, zhupp}@shu.edu.cn, chu@i.kyoto-u.ac.jp, wangrui12@sjtu.edu.cn

## Abstract

Key-value relations are prevalent in Visually-Rich Documents (VRDs), often depicted in distinct spatial regions accompanied by specific color and font styles. These non-textual cues serve as important indicators that greatly enhance human comprehension and acquisition of such relation triplets. However, current document AI approaches often fail to consider this valuable prior information related to visual and spatial features, resulting in suboptimal performance, particularly when dealing with limited examples. To address this limitation, our research focuses on few-shot relational learning, specifically targeting the extraction of key-value relation triplets in VRDs. Given the absence of a suitable dataset for this task, we introduce two new few-shot benchmarks built upon existing supervised benchmark datasets. Furthermore, we propose a variational approach that incorporates relational 2D-spatial priors and prototypical rectification techniques. This approach aims to generate relation representations that are more aware of the spatial context and unseen relation in a manner similar to human perception. Experimental results demonstrate the effectiveness of our proposed method by showcasing its ability to outperform existing methods. This study also opens up new possibilities for practical applications.

**Keywords:** few-shot learning, visually-rich document, variational inference

## 1. Introduction

Relational learning, also known as key-value relation extraction, is a fundamental task in comprehending Visually-Rich Documents (VRDs) (Dengel and Klein, 2002; Liu et al., 2019; Li et al., 2022). It involves automatically identifying and extracting key and value entities, as well as classifying the relations between these entities in a text based on a predefined schema. This task is typically performed on scanned or digitally generated documents, such as invoices, receipts, and business forms.

Recent document AI models pre-trained on large-scale scanned document datasets (Xu et al., 2020; Hong et al., 2021; Garncarek et al., 2021; Huang et al., 2022) have demonstrated promising performance by effectively leveraging multi-modal information. However, when it comes to real-world applications, the diverse layout formats and form styles found in VRDs present a significant challenge for adapting these AI models (Li et al., 2022). These approaches lack the inherent ability to automatically detect and identify new types of entities and relations in unfamiliar domains, especially in the absence of annotated data. In contrast, humans possess the remarkable ability to swiftly comprehend key-value patterns in VRD by analyzing just a few lines on the page.

Despite recent advances (Cheng et al., 2020;

Wang and Shang, 2022; Wang et al., 2022b), there is currently no established systematic framework suitable for addressing the more realistic task of few-shot relational learning in VRDs. This research area remains relatively unexplored and presents several unresolved challenges. One of the main challenges is effectively leveraging layout features in few-shot scenarios. In VRDs, key and value entities often have fixed positions and arrangements, such as being located at specific positions and arranged either up-down or left-right. Therefore, incorporating 2D spatial features can provide strong complementary supervision signals when extracting key-value triplets from VRDs. Another challenge lies in the optimal design of the multimodal fusion mechanism for few-shot learning. While it has been demonstrated that multimodal information can greatly enhance supervised learning, determining how to align and aggregate features across different modalities using only a small number of instances remains unclear. These challenges underscore the necessity for further exploration and development in this research direction.

Since there is currently no suitable dataset for the few-shot relational learning task, we first create two benchmark datasets, Few-CORD and Few-SEAB, based on the supervised VRD understanding benchmark datasets CORD (Park et al., 2019) and SEAB (Zhang et al., 2022). They allow us to evaluate the model's ability to transfer knowl-

---

\* Corresponding author.

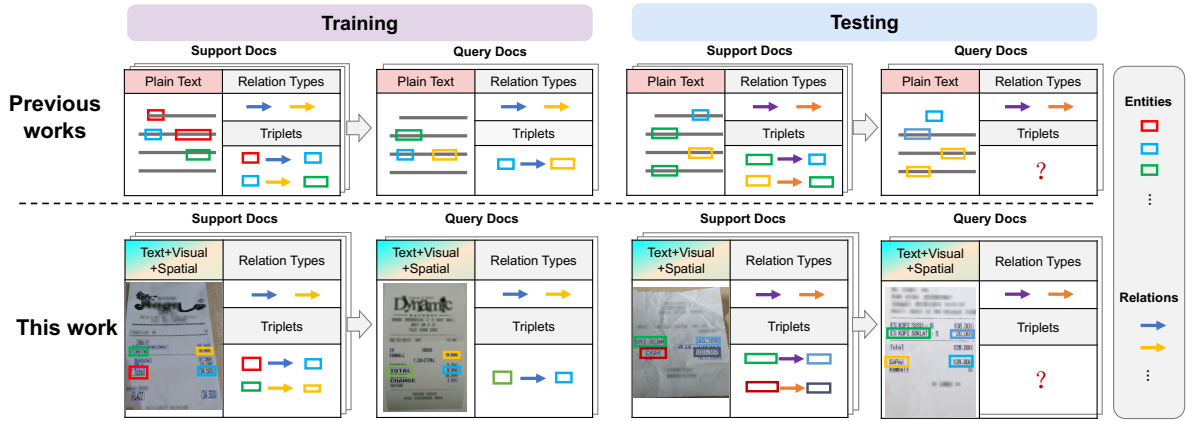


Figure 1: Illustration of the distinction between our work and previous works (Popovic and Färber, 2022) during an episode in the few-shot relational learning scenario. In the testing task, we aim to extract triplets that consist of entities and relation types for a given query document. Notably, this task involves a different set of relation types compared to the training task and is performed on a novel collection of documents. Conventional approaches typically rely on off-the-shelf OCR engines to extract text from original document images and solely rely on text features for extracting relational triplets. In contrast, our work takes a human-like perspective and leverages multimodal information to effectively extract the relational triplets. While we use simple receipts with well-aligned layouts to illustrate the idea, it is crucial to acknowledge that real-world scenarios are considerably more complex and challenging.

edge to new classes. Then, we propose a novel variational approach for few-shot relational learning in VRDs to address the challenges mentioned above. This approach incorporates spatial priors and effectively captures multimodal representations, resulting in improved performance on downstream tasks. We utilize a spatial prior encoder to leverage explicit information from Region of Interest (ROI) windows. These windows help us better understand the layout styles in documents and exploit the inherent spatial relationships between entities. To overcome the issue of biased prototypes generated by existing models due to the limited number of  $K$ -shot instances per category in the training data, we introduce prototypical rectification as a way to access optimal prototypes. In summary, our contributions are as follows:

- To the best of our knowledge, this work is the first to tackle the challenge of few-shot relational learning in VRDs. In order to facilitate research in this area, we introduce two new benchmark datasets specifically designed for the few-shot learning setting.
- We propose a novel variational approach that not only allows for the incorporation of spatial priors but also enables the extraction of robust relation-agnostic features, thereby alleviating the sensitivity of prototypes more accurately.
- Our method achieves new state-of-the-art performance for few-shot relational learning in VRDs, which has been extensively evaluated

on the constructed datasets to demonstrate its effectiveness.

## 2. Task Description

As shown in Figure 1, this task follows the few-shot settings in previous works (Han et al., 2018; Baldini Soares et al., 2019) and aims to extract key and value entities involving new relation classes, by training on known relation classes with a small number of examples. The testing set examples often include relation types that were not present in the training set. For practicality and simplicity, we consider only the key and value entities from a single relation within a given document as a training/test example. Then, we adopt a single collapse sequential labeling model (Wang and Blunsom, 2013; Islam and Foulds, 2019) to jointly extract sets of key and value entities ( $\mathcal{E}_k$  or  $\mathcal{E}_v$ ) given their relation types  $\mathcal{R}$  from the document  $\mathcal{D}$ . This is achieved by extending the label space of the previous named entity recognition (NER) scheme for entity types, e.g., “Cash” (in CORD) and “Consignee” (in SEAB), with additional key or value identifiers (“-Key”, “-Value”) into the combination labels “Cash-Key/Cash-Value” and “Consignee-Key/Consignee-Value”. In fact, we perform token-level classification to extract entity mentions from the document using a standard “BIO” tagging scheme - “Begin, Inside, Other”. For example, “Consignee-Key-B” indicates the first token in the entity mention of “Consignee-Key”, while “O” represents “Other” that none of the above applies.

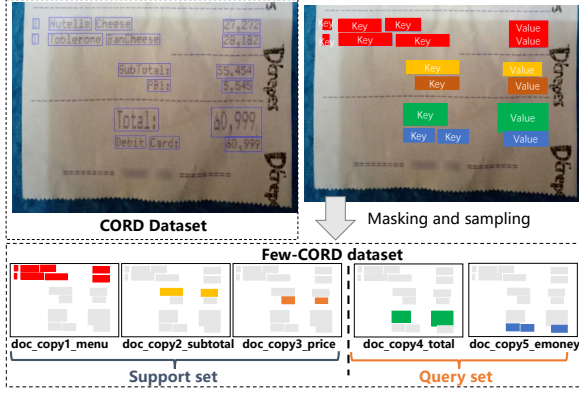


Figure 2: Copying, masking and sampling.

### 3. Few-Shot Datasets

#### 3.1. Relation-Wise Sampling Strategy

During sampling, as shown in Figure 2, we should carefully select and organize instances from the documents to ensure that the  $N$ -way setting is maintained while still capturing the necessary contextual information. By implementing an appropriate sampling strategy, our approach aims to strike a balance between incorporating document-level context and adhering to the constraints of the  $N$ -way setting, enabling effective few-shot learning for sequential labeling tasks. We use the extended dataset consisting of multiple copies of the original documents. We sample from this extended dataset, adding one more converted document to the support or query set after each sampling step until the support or query set reaches the desired number of entity classes ( $N$ -way) and instances per class ( $K$ -shot). The overall procedure of the hierarchical sampling method is summarized in Algorithm 1.

In our approach, we employ  $N$ -way<sup>1</sup>  $K$ -shot learning to train our few-shot relational learning system. This involves constructing training or testing episodes iteratively. While considering the document-level context is important for few-shot learning tasks, sampling at the document level can pose challenges within the  $N$ -way setting. This is because documents often contain multiple relations/entities, which may exceed the limits of the  $N$ -way setting. To maintain the  $N$ -way setting, as shown in Figure 2, we hereby make relation-wise copies for each document which means each copy contains only entities involved in a particular relation type, and other bounding boxes are simply re-labeled as `Other` type. The mathematical expression is:  $\mathcal{D}_{\mathcal{R}} = \mathcal{M}(\mathcal{D}, \bar{\mathcal{R}})$ , where  $\mathcal{D}_{\mathcal{R}}$  represents the copy in which irrelevant entities to  $\mathcal{R}$  are masked. The masking operation  $\mathcal{M}$  selectively retains only

<sup>1</sup>Note that  $N/2$  denotes the number of relation types thus  $N$  denotes the number of entity types (key+value).

#### Algorithm 1 Relation-wise $N$ -way $K$ -shot Sampling

**Input:** Documents  $\hat{\mathcal{D}}, N, K, K'$ ;  
**Output:** Support Set  $\mathcal{S}$ , Query Set  $\mathcal{Q}$ ;

- 1:  $\mathcal{S} \leftarrow [], \mathcal{Q} \leftarrow []$ ; ▷ initialization
- 2: **for**  $j \leftarrow 1$  **to**  $N/2$  **do**
- 3:  $\mathcal{S}[j] \leftarrow \{\}$ ; ▷  $N/2$  is the number of relation types
- 4:  $\mathcal{Q}[j] \leftarrow \{\}$ ; ▷  $N$  is the number of entity types
- 5: **end for**
- 6: **repeat**
- 7: randomly sample  $(\mathcal{D}^{(i)}, \mathcal{R}^{(i)})$  from the extended masked document dataset; ▷  $\mathcal{D}^{(i)} \in \hat{\mathcal{D}}$
- 8: **if**  $|\mathcal{S}| < N/2$  &  $|\mathcal{S}[j]| < K$  **then**
- 9:  $\mathcal{S}[j] \leftarrow \mathcal{S}[j] \cup (\mathcal{D}^{(i)}, \mathcal{R}^{(i)})$ ; ▷ add to the support set
- 10: **end if**
- 11: **if**  $|\mathcal{Q}| < N/2$  &  $|\mathcal{Q}[j]| < K'$  and  $(\mathcal{D}^{(i)}, \mathcal{R}^{(i)}) \notin \mathcal{S}[j]$  **then**
- 12:  $\mathcal{Q}[j] \leftarrow \mathcal{Q}[j] \cup (\mathcal{D}^{(i)}, \mathcal{R}^{(i)})$ ; ▷ add to the query set
- 13: **end if**
- 14: **until**  $|\mathcal{S}| = |\mathcal{Q}| = N/2$  and  $\{\forall j \mid |\mathcal{S}[j]| = K$  and  $|\mathcal{Q}[j]| = K'\}$
- 15: **return**  $\mathcal{S}, \mathcal{Q}$ ;

the entities related to  $\mathcal{R}$  and suppresses or eliminates the rest, denoted as  $\bar{\mathcal{R}}$ .

For each episode, we randomly choose  $N$  classes ( $N$ -way) and sample  $K$  examples ( $K$ -shot) from the extended dataset containing multiple copies for each relation class to build support set  $\mathcal{S}_{train} = \{\mathcal{D}^{(i)}, \mathcal{R}^{(i)}\}_{i=1}^{N \times K}$  and  $K'$  examples for a query set  $\mathcal{Q}_{train} = \{\mathcal{D}^{(j)}, \mathcal{R}^{(j)}\}_{j=1}^{N \times K'}$ , ensuring that the support set  $\mathcal{S}$  and query set  $\mathcal{Q}$  do not overlap ( $\mathcal{S} \cap \mathcal{Q} = \emptyset$ ). In addition, when adding one or more examples to the support or query set after each sampling step until the support or query set reaches the desired number of entity classes ( $N$ -way) and instances per class ( $K$ -shot).

In the training phase, we train the few-shot learning system using the support and query sets  $(\mathcal{S}_{train}, \mathcal{Q}_{train})$ , where the supervisions for both sets are visible. In the testing phase, we predict the new classes in the query set  $\mathcal{Q}_{test}$  and evaluate the performance given the ground truths.

#### 3.2. Dataset Details

This section provides an overview of the Few-CORD and Few-SEAB datasets. We construct these two benchmark datasets on the top of the CORD (Park et al., 2019) and SEAB (Zhang et al., 2022) datasets and serve for few-shot relational learning. These datasets consist of diverse and realistic examples, enabling us to effectively measure the effectiveness of relational learning from VRDs. Table 1 provides the details of the newly




Figure 3: The double-humped distribution of specific key-value types on a document page, including: “Shipper” and “Weight” in the SEAB dataset and “Menu” and “Total” in the CORD dataset.

Dataset	Train			Test		
	#Doc	# BD	#Types	#Doc	#BD	#Types
CORD	800	18,915	32/16	200	4,466	32/16
SEAB	3,562	249,255	44/22	953	73,873	44/22
Few-CORD	1,211	32,160	18/9	702	15928	14/7
Few-SEAB	20,831	575,391	24/12	4,048	146,312	20/10

Table 1: Statistics of the supervised and few-shot datasets, including the numbers of documents (Doc), bounding boxes (BD), and types of entities/relations (Types A/B means A: base classes in training; B: novel classes in testing).

created few-shot datasets.

## 4. Human-Like Approaches

### 4.1. Motivation

People possess the remarkable ability to quickly grasp relation patterns with minimal exposure to instances. They can effectively utilize layout information across various document images, given the explicit spatial relationships between key and value entities. These two-dimensional spatial priors serve as valuable cues for learning of the relations in VRDs. However, existing approaches struggle to effectively leverage these features. Moreover, people have the capacity to infer insights from unseen examples after seeing only a few instances. They can learn high-dimensional class-agnostic features that transcend linguistic boundaries (Wang et al., 2022a) and contextual limitations (Han et al., 2023) by establishing connections across different classes, irrespective of their novelty or familiarity. This unique capability allow them to comprehend and generalize knowledge within diverse contexts to identify new class relations and entities.

Therefore, in this study, our objective is to introduce a novel approach that emulates human-like behavior in learning relations. Our approach focuses on incorporating relational 2D-spatial priors and aims to learn high-dimensional class features, irrespective of their novelty or familiarity. By combining these two aspects, we aim to bridge the gap between human cognition and existing mod-

els, thereby enabling more effective few-shot relational learning in VRDs.

### 4.2. Architecture

Figure 4 illustrates the overall architecture of our method, which comprises three stages: 1) **ROI Regression**, which learn regions of interest (ROIs) using carefully curated golden window instances; 2) **Prototypical Rectification**, which tackles the issues associated with biased prototype representations; and 3) **Proximity-based classification**, which predicts the final token labels. Figure 4 visually illustrates the comprehensive architecture of our method, which involves three key stages.

### 4.3. ROI Regression

Our preliminary studies have revealed a strong correlation between specific key and value entities in the two-dimensional layout space of VRDs. In general, keys and values tend to exhibit a distinctive double-hump distribution within this spatial arrangement (see Figure 3), which can be seen as a de facto two-dimensional spatial prior.

To leverage this distributional regularity, we introduce the concept of a “golden window” that encompasses both the keys and values belonging to a particular type. During training, this golden window serves as an explicit supervision signal, guiding the model to focus on relevant local regions of interest (ROIs). We first employ LayoutLM (Xu et al., 2020) and LayoutLMv2 (Xu et al., 2021) as our multimodal encoder to extract token-level features  $\mathbf{H}$  from multimodal channels, including text, spatial (1D position and 2D layout), and visual. For simplicity, we call them LLM. Then, we utilize a Variational Autoencoder (VAE) (Kingma and Welling, 2013) to model the distributions of local ROIs. Specifically, the output of the decoder in VAE corresponds to four dimensions  $(x_1, y_1, x_2, y_2)$  used for reconstructing the golden window. This reconstruction enables improved accuracy in predicting key-value associations by explicitly considering their geometric layout within visually-rich documents.

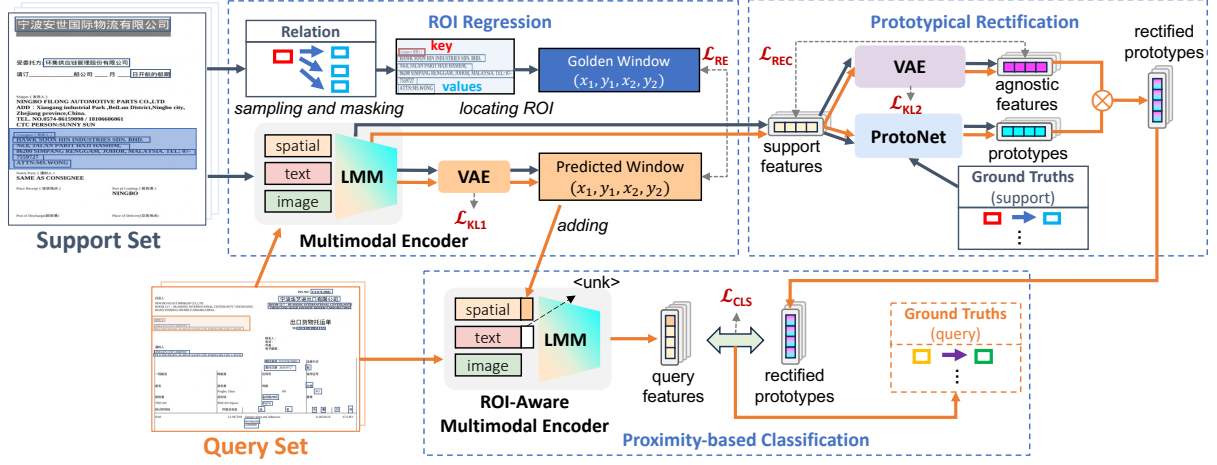


Figure 4: Our model architecture comprises three key components: ROI regression, prototypical rectification, and proximity-based classification. It can generate more robust representations that encompass multiple modalities by directing its attention to relevant regions (by predicting the explicit ROI windows) and learn high-dimensional relation-agnostic features using prototypical rectification, helping adaptation to new relation classes.

#### 4.4. Prototypical Rectification

Given the support features extracted by the multimodal encoder, we utilize the prototypical network (ProtoNet) paradigm (Snell et al., 2017) to compute an  $M$ -dimensional representation  $\mathbf{p}_c \in \mathbb{R}^M$  for each entity class. We take the average of the token embeddings to obtain prototypes.

However, ProtoNet encodes the support instances as a single feature vector, but this does not provide an accurate estimate of the class center (Yang et al., 2020; Gao et al., 2019a; Liu et al., 2022), especially when the data is sparse, and the examples are highly variable, making it difficult to represent the distribution of the class.

Building on recent advancements in variational feature learning (Han et al., 2023; Xia et al., 2020; Tran and Nguyen, 2018), we introduce a variational rectification mechanism to incorporate category distribution information into the prototypes. This mechanism enables us to transform support features into class-wise representations. We expect the rectified prototypes  $\tilde{\mathbf{z}}_c$  to capture more generic features of the class that are robust to the variance of support instances. Here,  $\mathbf{p}_c$  represents the prototype feature of class  $k$ . We approximate the class distribution  $N(\mu_p, \Sigma_p)$  and sample a variational feature  $\tilde{\mathbf{z}}_p$  from this distribution. Then, we combine the prototypes and the variational feature using the following equation:

$$\tilde{\mathbf{p}}_c = \mathcal{A}(\mathbf{p}_c, \tilde{\mathbf{z}}_p) = \mathbf{p}_c \otimes \text{sigmoid}(\tilde{\mathbf{z}}_p), c \in C \quad (1)$$

#### 4.5. Proximity-based Classification

To incorporate the ROI information, we enhance the multimodal input sequence by appending the coordinates of the predicted window  $b_{\text{pred}}$  to the

end of the sequence, specifically for the spatial modality. In order to maintain alignment between the multimodal inputs, we introduce a special token  $\langle \text{UNK} \rangle$  to the corresponding text modality. Consequently, the resulting incorporated sequence is extended from the initial  $n$  tokens to  $n+1$  tokens, denoted as  $[t_1, t_2, \dots, t_n, \langle \text{UNK} \rangle]$ . We take the last hidden layer outputs of the ROI-aware multimodal encoder  $\mathbf{H}$  as the final query features. This embedding captures sufficient multimodal information while incorporating the ROI information for better relational representations.

$$\tilde{\mathbf{H}} = \text{LMM} \left( \begin{bmatrix} t_1 & t_2 & \dots & t_n & \langle \text{UNK} \rangle \\ b_1 & b_2 & \dots & b_n & b_{\text{pred}} \\ \mathcal{I}_{b_1} & \mathcal{I}_{b_2} & \dots & \mathcal{I}_{b_n} & \mathcal{I}_{b_{\text{pred}}} \end{bmatrix} \right) \quad (2)$$

For each query instance  $Q_{\text{test}}$ , we perform a proximity-based classification by computing the 2-norm distance (i.e., the Euclidean distance) between the query embeddings  $\tilde{\mathbf{h}}_i$  and rectified prototypes  $\tilde{\mathbf{p}}_c$ .

$$\mathbf{d}_c = d_{l_2\text{-norm}}(\mathbf{p}_c, \tilde{\mathbf{h}}_i). \quad (3)$$

Then, we normalize the prediction probability of  $x$  over all classes using the softmax function. The model predicts the label as the nearest prototype to the input token. The cross-entropy loss is used as the loss function for multi-classification task.

#### 4.6. Training Objective

We train our few-shot relation learning model in an end-to-end manner with the following multi-task loss,

$$\mathcal{L}_{\text{final}} = \mathcal{L}_{\text{REC}} + \alpha \mathcal{L}_{\text{KL1}} + \mathcal{L}_{\text{RE}} + \beta \mathcal{L}_{\text{KL2}} + \mathcal{L}_{\text{CLS}}, \quad (4)$$

Method	LLM	Proto	VAE	Few-CORD					Few-SEAB				
				1-shot	2-shot	3-shot	4-shot	5-shot	1-shot	2-shot	3-shot	4-shot	5-shot
ProtoNet	BERT	✓	✗	32.32	<u>35.14</u>	<u>38.87</u>	<u>40.08</u>	<u>42.95</u>	26.88	28.59	<u>30.08</u>	<u>31.74</u>	<u>34.11</u>
NNShot	BERT	✓	✗	29.19	32.34	35.16	36.28	38.04	25.16	26.10	27.92	28.80	29.95
StructShot	BERT	✓	✗	<b>33.54</b>	34.95	37.41	38.31	40.38	<u>27.30</u>	<u>28.63</u>	29.20	30.19	31.75
VFA	BERT	✗	✓	30.08	31.34	32.86	34.89	37.14	24.83	25.61	27.11	28.65	29.06
<b>ProtoRec</b>	BERT	✓	✓	<u>33.30</u>	<b>35.40</b>	<b>39.01</b>	<b>40.53</b>	<b>43.14</b>	<b>27.45</b>	<b>29.04</b>	<b>30.35</b>	<b>32.15</b>	<b>34.28</b>
ProtoNet	LayoutLM	✓	✗	70.25	74.10	77.02	79.31	80.40	60.95	64.02	66.31	69.53	73.12
NNShot	LayoutLM	✓	✗	68.20	72.70	73.76	75.24	76.67	58.80	61.75	62.27	64.40	66.89
StructShot	LayoutLM	✓	✗	<u>71.38</u>	73.88	74.52	77.24	77.83	61.15	63.14	63.50	65.28	68.10
VFA	LayoutLM	✗	✓	68.39	69.79	71.83	73.18	74.98	58.14	60.23	63.14	64.36	66.12
VFA+ROI	LayoutLM	✗	✓	68.71	70.07	72.31	73.91	75.65	58.76	60.61	63.91	65.08	66.79
<b>ROI-Aware</b>	LayoutLM	✓	✓	71.33	<u>74.96</u>	<u>77.84</u>	<u>80.78</u>	<u>81.22</u>	<u>61.79</u>	<u>64.92</u>	<u>66.45</u>	<u>69.84</u>	<u>73.19</u>
<b>ProtoRec+ROI</b>	LayoutLM	✓	✓	<b>73.21</b>	<b>76.19</b>	<b>78.42</b>	<b>81.35</b>	<b>81.54</b>	<b>62.77</b>	<b>65.54</b>	<b>66.59</b>	<b>69.95</b>	<b>73.28</b>
ProtoNet	LayoutLMv2	✓	✗	70.30	74.22	77.16	79.35	80.52	61.18	64.10	66.43	69.80	73.37
NNShot	LayoutLMv2	✓	✗	68.17	72.80	73.88	75.39	76.84	59.09	61.81	62.35	64.49	67.14
StructShot	LayoutLMv2	✓	✗	71.45	73.95	74.62	77.50	78.11	61.30	63.34	63.57	65.37	68.18
VFA	LayoutLMv2	✗	✓	68.87	70.16	72.05	73.63	75.56	58.77	60.59	63.70	65.03	67.18
VFA+ROI	LayoutLMv2	✗	✓	69.21	71.39	72.63	74.15	76.01	59.20	61.57	63.97	65.28	68.26
<b>ROI-Aware</b>	LayoutLMv2	✓	✓	<u>71.59</u>	<u>75.83</u>	<u>77.87</u>	<u>80.92</u>	<u>81.40</u>	<u>61.94</u>	<u>65.02</u>	<u>66.48</u>	<u>69.86</u>	<u>73.52</u>
<b>ProtoRec+ROI</b>	LayoutLMv2	✓	✓	<b>73.37</b>	<b>76.90</b>	<b>78.54</b>	<b>81.46</b>	<b>81.88</b>	<b>63.05</b>	<b>65.59</b>	<b>66.83</b>	<b>70.18</b>	<b>73.70</b>
<b>ProtoRec+ROI+CF</b>	LayoutLMv2	✓	✓	<u>73.32</u>	<u>76.96</u>	<u>78.63</u>	<u>81.40</u>	<u>81.85</u>	<u>62.98</u>	<u>65.47</u>	<u>66.80</u>	<u>70.21</u>	<u>73.65</u>
<b>ProtoRec+ROI+SW</b>	LayoutLMv2	✓	✓	<b>73.56</b>	<b>77.39</b>	<b>78.80</b>	<b>81.88</b>	<b>82.18</b>	<b>63.43</b>	<b>65.87</b>	<b>67.30</b>	<b>70.56</b>	<b>73.96</b>

Table 2: The averaged f1-scores on the Few-CORD and Few-SEAB datasets. **Bold** and underline indicate the best and second-best scores in each group. **ROI-Aware** and **ProtoRec** are two main components of the variational model proposed in this work. **+CF** denotes filling the windows with a light color. **+SW** refers to adversarial learning by shrinking windows.

where the weight coefficients  $\alpha = \beta = 2.5 \times 10^{-4}$ . In the training phase, parameters are updated in each episode. In the testing phase, we directly deploy the model to predict novel types of entities without computing the loss or updating the model parameters.

## 5. Implementation Details

### 5.1. Backbones

We implement our model based on the approach proposed by Ding et al. (2021). For the backbone architecture, we utilize pre-trained document intelligence models such as LayoutLM (Xu et al., 2020) and LayoutLMv2 (Xu et al., 2021). These models serve as the foundation for our method, providing robust and effective feature extraction capabilities. In addition, we also evaluate the performance of the baseline model using Bert (Devlin et al., 2019) solely for the text modality.

### 5.2. Hyper-parameters

We optimize the model with AdamW optimizer on a dual NVIDIA 3090 GPUs machine, and the learning rate is  $1 \times 10^{-5}$ . We fine-tune the model with 10,000 iterations in training and evaluate the performance using the averaged scores over 500 testing iterations.

### 5.3. Evaluation Scheme

Following (Han et al., 2018; Gao et al., 2019b; Ding et al., 2021), we adopt episode evaluation. We

compute the micro-F1 score over a number of test episodes. Each episode contains a  $K$ -shot support set with manually annotated labels and a  $K'$ -shot query set without any annotated labels.

### 5.4. Baselines

**ProtoNet**: is a baseline prototype system that assigns each token representation to the nearest label representation by learning the examples in the training set using a prototypical network (Snell et al., 2017; Garcia and Bruna, 2018).

**NNShot** and **StructShot** (Yang and Katiyar, 2020): are state-of-the-art methods based on token-level nearest neighbor classification. In contrast to the ProtoNet, NNShot determines the tag of a query based on token-level distance. On the other hand, StructShot incorporates an additional Viterbi decoding during the inference phase to improve the overall performance.

**VFA**: Variational Feature Aggregation (VFA) (Han et al., 2023) is originally proposed for few-shot target detection in the CV domain. We re-implement and adapt it to our specific task.

## 6. Experiment Results

### 6.1. Overall Evaluation

Table 2 shows the f1-scores on the Few-CORD and Few-SEAB datasets. In general, we have observed that the BERT backbone performs much worse than LayoutLM and LayoutLMv2. This

is primarily due to BERT’s limitation in capturing non-textual information. In contrast, LayoutLM and LayoutLMv2, which includes visual information, proves to be more effective. However, since LayoutLM itself does not inherently model visual features, we utilize ResNet as its visual encoder. Nonetheless, the misalignment between visual modalities and other modalities affects the effectiveness of LayoutLM as compared to LayoutLMv2.

From a horizontal perspective, it is generally observed that as the number of shots, denoted as  $K$ , increases, the performance of all methods tends to improve. This improvement can be attributed to the fact that as  $K$  increases, the prototypical representations become closer to the actual data points, leading to better performance. Overall, our prototype rectification (**ProtoRec**) and **ROI-aware** methods consistently yield stable improvements regardless of the backbone model used. It is worth mentioning that the combination of ProtoRec and ROI-Aware (**ProtoRec+ROI-Aware**) achieves the best performance among the methods mentioned. It is important to note that adding a Layout-aware Encoder to BERT’s backbone is not possible due to BERT’s lack of 2D position information, so we only show the ProtoRec result in the BERT group.

## 6.2. Semantic Similarity

Prototype sharing (Rymarczyk et al., 2021) can help identify prototypical similarities between classes. In this study, we aim to explore the relevance of relations and semantic similarity between entities, conducting an empirical study on entity-class similarity. For this purpose, we employ BERT and LayoutLMv2, having pre-trained on the training set to generate embeddings for all entities mentioned in the test set. We randomly selected 100 instances of entity embeddings for each fine-grained type and averaged them. Then, we calculate the dot product between the central representations of each entity type to measure their similarity. By comparing the results obtained from BERT with those from LayoutLMv2, we gain insights into the similarity. Figure 5 depicts that entity types sharing the same coarse-grained type tend to exhibit greater similarity, thereby facilitating knowledge transfer.

## 6.3. Transfer Learning Capacity

Evaluating the model’s transfer learning capacity is crucial for few-shot learning. To conduct the experiment on transfer learning, following (Ding et al., 2021), we reorganize the SEAB dataset and construct two novel datasets, Few-SEAB (Inter) and Few-SEAB (Intra), based on it by adopting different splitting strategies. We divide the entire set of entities into 4 coarse-grained

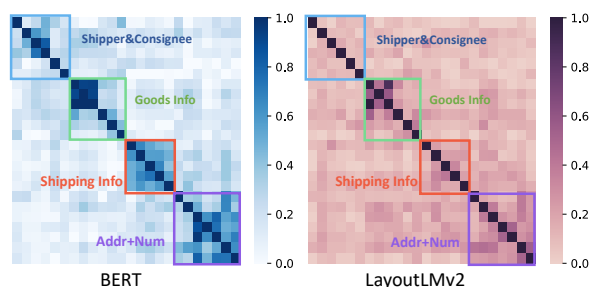


Figure 5: Semantic similarity heatmap for the entity representations generated by using BERT and LayoutLMv2.

disjoint subsets, e.g., “Shipper and Consignee”, “Goods Information”, “Shipping Information”, and “Address+Numbers”. In other words, we construct the Few-SEAB (Intra) dataset  $\langle \hat{D}_{\text{train}}, \hat{D}_{\text{test}} \rangle$  according to the coarse-grained types with the principle that the entities in different sets belong to different coarse-grained types. For example, “Shipper and Consignee”, “Goods Information” only appear in  $\hat{D}_{\text{train}}$ , and “Shipping Information”, “Address+Numbers” only appear in  $\hat{D}_{\text{test}}$ . This characteristic ensures that the training and test sets have minimal overlap in terms of shared knowledge. In contrast, while the fine-grained entity types in the Few-SEAB (Inter) dataset are mutually disjoint in  $\hat{D}_{\text{train}}$  and  $\hat{D}_{\text{test}}$ , the coarse-grained types are shared, which means both  $\hat{D}_{\text{train}}$  and  $\hat{D}_{\text{test}}$  should contain all four fine-grained types “Shipper and Consignee”, “Goods Information”, “Shipping Information”, and “Address+Numbers”. Table 3 gives the comparison of the results under the Intra and Inter transfer learning settings.

## 6.4. Effects of Prototypical Rectification

In the testing phase, the model needs to refer to the mean feature of  $K$ -shot examples as the class prototype. As shown in Figure 6, our estimated class prototypes are more robust and accurate than the baseline. We get real class centers for all entity types via supervised tasks. We use all few-shot learning models to compute the relative distance from each model-generated prototype to the real center. The distances to real class centers do not increase too much as the shot decreases because our method can fully leverage base classes’ distributions to estimate novel classes’ distributions. The prototypes sampled from distributions are robust to the variance of support examples. While the baseline is sensitive to the number of support examples.

Method	LLM	Proto	VAE	Inter					Intra				
				1-shot	2-shot	3-shot	4-shot	5-shot	1-shot	2-shot	3-shot	4-shot	5-shot
ProtoNet	BERT	✓	✗	33.03	36.73	37.08	38.36	39.54	24.39	26.68	27.86	28.35	29.87
NNShot	BERT	✓	✗	34.35	34.68	35.75	35.93	36.47	23.46	24.43	25.10	25.57	26.20
StructShot	BERT	✓	✗	<b>36.45</b>	36.63	36.83	37.10	37.29	<b>25.28</b>	25.87	26.19	26.86	27.22
VFA	BERT	✗	✓	31.65	33.87	35.48	36.11	36.69	21.90	23.31	24.58	25.42	25.94
<b>ProtoRec</b>	BERT	✓	✓	<u>34.80</u>	<u>37.10</u>	<u>37.45</u>	<u>38.71</u>	<u>39.86</u>	<u>25.06</u>	<u>27.18</u>	<u>28.51</u>	<u>28.92</u>	<u>30.09</u>
ProtoNet	LayoutLM	✓	✗	66.75	69.31	70.86	74.88	77.84	55.67	58.23	60.92	64.20	65.89
NNShot	LayoutLM	✓	✗	63.83	65.50	66.32	68.43	70.88	53.74	56.37	58.83	62.13	63.48
StructShot	LayoutLM	✓	✗	67.69	69.35	69.64	71.64	72.35	55.53	57.90	58.96	62.41	64.81
VFA	LayoutLM	✗	✓	64.45	65.89	67.34	70.63	72.90	51.48	54.67	55.86	59.34	61.22
VFA+ROI	LayoutLM	✗	✓	65.32	66.17	67.81	71.23	73.12	51.96	55.04	56.12	59.76	61.46
<b>ROI-Aware</b>	LayoutLM	✓	✓	<u>68.80</u>	<u>70.39</u>	<u>71.48</u>	<u>75.60</u>	<u>78.17</u>	<u>56.80</u>	<u>58.31</u>	<u>61.20</u>	<u>64.38</u>	<u>66.95</u>
<b>ProtoRec+ROI</b>	LayoutLM	✓	✓	<b>69.67</b>	<b>71.18</b>	<b>71.88</b>	<b>75.83</b>	<b>78.64</b>	<b>57.94</b>	<b>59.13</b>	<b>61.79</b>	<b>64.85</b>	<b>67.05</b>
ProtoNet	LayoutLMv2	✓	✗	67.21	69.55	71.23	75.14	78.22	56.16	58.52	61.65	64.96	66.31
NNShot	LayoutLMv2	✓	✗	63.87	66.20	66.87	68.87	71.26	54.70	57.23	58.93	62.27	63.02
StructShot	LayoutLMv2	✓	✗	68.16	69.73	70.75	71.63	74.27	56.82	58.24	59.21	62.54	63.27
VFA	LayoutLMv2	✗	✓	64.89	66.46	68.05	70.98	73.59	52.42	55.61	56.74	60.38	61.69
VFA+ROI	LayoutLMv2	✗	✓	65.43	67.04	68.78	71.61	74.04	52.95	56.38	57.61	60.84	61.93
<b>ROI-Aware</b>	LayoutLMv2	✓	✓	<u>70.11</u>	<u>71.19</u>	<u>72.64</u>	<u>75.80</u>	<u>78.67</u>	<u>57.69</u>	<u>59.80</u>	<u>61.91</u>	64.87	67.14
<b>ProtoRec+ROI</b>	LayoutLMv2	✓	✓	<b>71.59</b>	<b>72.76</b>	<b>73.32</b>	<b>75.97</b>	<b>79.59</b>	<b>58.79</b>	<b>60.44</b>	<b>62.37</b>	<b>65.13</b>	<b>67.53</b>

Table 3: The averaged f1-score on the Few-SEAB (Inter) and Few-SEAB (Intra) datasets. Few-SEAB (Inter) requires the entity types to be disjoint in the training set and the test set. Few-SEAB (Intra) requires the entity types to be disjoint in the training set and the test set, while Few-SEAB (Inter) shares the coarse-grained types in the training and test sets.

	Train			Test		
	#Doc	#BD	#Types	#Doc	#BD	#Types
Inter	18,966	529,378	24/12	4,652	167,087	20/10
Intra	20,187	575,687	24/12	4,554	163,647	20/10

Table 4: Statistics of the inter- and intra-learning settings in the Few-SEAB dataset, including the numbers of documents (Doc), bounding boxes (BD), and types of entities/relations (Types).

## 7. Related Work

### 7.1. VRD Understanding

Researchers have been actively working on addressing the task of understanding VRDs by incorporating multiple modalities, including text, layout, and visual features. In recent years, a variety of approaches have emerged in this field. Grid-based methods, including Chargrid (Katti et al., 2018), BertGrid (Denk and Reisswig, 2019), and ViBERTgrid (Lin et al., 2021), utilize 2D feature maps for document representation. Graph neural network (GNN)-based methods, such as GraphIE (Liu et al., 2019), MatchVIE (Tang et al., 2021), and DualVIE (Zhang et al., 2022), FormNet (Lee et al., 2022), FormNetv2 (Lee et al., 2023) model associations among text segments using graph structures. By leveraging the connectivity between segments, these models can better understand relationships within a document. Transformer-based methods such as LayoutLM (Xu et al., 2020) incorporate two-dimensional relative position information based on BERT’s architecture (Devlin et al.,

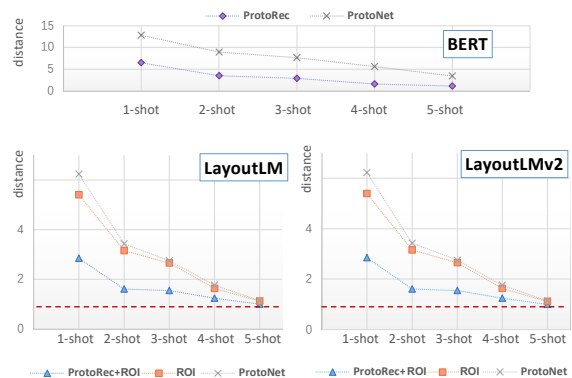


Figure 6: Average distance from estimated prototype to class centroid for  $K$ -shot instances. Distance from the estimated prototype of the  $K$ -shot instance to the real centroid of the class. We define the distance between the prototype estimated from 5-shots and the actual centroid of the class as 1 and use this as a reference. We report the average distance for all new classes.

2019). This enables the model to perceive the positions of text segments within a document. Later works like LayoutLMv2 (Xu et al., 2021), StrucText (Li et al., 2021), StrucTextv2 (Yu et al., 2023), LiLT (Wang et al., 2022a), and LayoutLMv3 (Huang et al., 2022) further integrate visual channel input into a unified multimodal Transformer framework, leading to better alignments and representations across different modalities. These approaches have demonstrated promising progress



in improving the performance of understanding VRDs. While Wang et al. (2021) and Wang et al. (2023) study the effect of reading order on the task of towards human-like machine reading, they did not discuss the local order involving in key-value relationships in the document layout.

## 7.2. Few-shot Relational Learning

Most recent works on few-shot relational learning in NLP focus on relation extraction or classification using only the text modality (Han et al., 2018; Baldini Soares et al., 2019). These works typically follow the  $N$ -way  $K$ -shot setting (Snell et al., 2017), where a relation instance needs to be assigned to one of  $N$  classes based on only  $K$  examples per class. Evaluation is done on few-shot relation extraction benchmarks such as FewRel (Han et al., 2018), FewRel2.0 (Han et al., 2018), and Few-Shot TACRED (Sabo et al., 2021). While some models have achieved performance surpassing human performance on these tasks, it has been argued by Brody et al. (2021) and Sabo et al. (2021) that existing benchmarks are far from real-world applications. They emphasize that more challenges remain unsolved in this direction, including reducing reliance on entity type information and focusing more on relations for realistic scenarios. Another challenge is transitioning to document-level relational learning (Popovic and Färber, 2022), which requires different architectures compared to sentence-level approaches. Given that the goal of few-shot relation learning is to quickly adapt to unseen relation classes with limited samples by leveraging training on known relation classes, metric learning serves as the primary paradigm for this task. Prominent methods include Prototypical network (Snell et al., 2017; Garcia and Bruna, 2018) or its variants like prototype rectification (Liu et al., 2022). These methods learn a prototype for each class and classify items based on their similarities to the prototypes. Additionally, efforts have been made to develop frameworks for document-level relational representation learning in visually-rich documents (Li et al., 2022) and enhance prototypes using relation information through prototype rectification modules (Liu et al., 2022), which have also shown promising results.

## 8. Conclusion

In this paper, we have addressed the research topic of few-shot relational learning in visually-rich documents. Given the limited availability of datasets in this domain, we have reorganized existing supervised benchmark datasets and developed a sampling algorithm specifically tailored for few-shot learning settings. Inspired by human-like cognition, we also propose a novel variational approach to incorporate 2D-spatial pri-

ors and relation-agnostic features to improve the model’s performance on the few-shot relational learning task. The 2D-spatial priors models the ROI window which guides the model’s attention towards relevant regions for a given relation within the document image. Additionally, we have introduced a prototypical rectification mechanism to enhance its ability to generalize and adapt to new instances despite limited training data. Through extensive experiments conducted on our newly created datasets, we have demonstrated the effectiveness of both our ROI-aware and prototypical rectification techniques in improving performance on few-shot relational learning tasks for VRDs. These advancements significantly contribute to the progress of research in few-shot relational learning for VRDs and pave the way for further exploration in this direction.

## 9. Acknowledgements

This work was supported by National Natural Science Foundation of China (Young Program: 62306173, General Program: 62176153), JSPS KAKENHI Program (JP23H03454), Shanghai Sailing Program (21YF1413900), Shanghai Pujiang Program (21PJ1406800), Shanghai Municipal Science and Technology Major Project (2021SHZDZX0102), the Alibaba-AIR Program (22088682), and the Tencent AI Lab Fund (RBFR2023012).

## 10. Bibliographical References

- Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. 2019. [Matching the blanks: Distributional similarity for relation learning](#). In *The 57th Annual Meeting of the Association for Computational Linguistics*, pages 2895–2905.
- Yujia Bao, Menghua Wu, Shiyu Chang, and Regina Barzilay. 2020. Few-shot text classification with distributional signatures. In *International Conference on Learning Representations*.
- Sam Brody, Sichao Wu, and Adrian Benton. 2021. [Towards realistic few-shot relation extraction](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5338–5345, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mengli Cheng, Minghui Qiu, Xing Shi, Jun Huang, and Wei Lin. 2020. [One-shot text field labeling using attention and belief propagation for structure information extraction](#). In *Proceedings of*

- the 28th ACM International Conference on Multimedia*, page 340–348, New York, NY, USA. Association for Computing Machinery.
- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. 2021. InfoXLM: An information-theoretic framework for cross-lingual language model pre-training. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3576–3588.
- 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 the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451.
- Andreas Dengel and Bertin Klein. 2002. smartfix: A requirements-driven system for document analysis and understanding. In *Proceedings of the 5th International Workshop on Document Analysis Systems V, DAS '02*, page 433–444, Berlin, Heidelberg. Springer-Verlag.
- Timo Denk and Christian Reisswig. 2019. Bertgrid: Contextualized embedding for 2d document representation and understanding. In *Workshop on Document Intelligence at NeurIPS*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. *the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*.
- Ning Ding, Guangwei Xu, Yulin Chen, Xiaobin Wang, Xu Han, Pengjun Xie, Haitao Zheng, and Zhiyuan Liu. 2021. Few-NERD: A few-shot named entity recognition dataset. In *The Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP 2021)*, pages 3198–3213.
- Tianyu Gao, Xu Han, Zhiyuan Liu, and Maosong Sun. 2019a. Hybrid attention-based prototypical networks for noisy few-shot relation classification. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence*.
- Tianyu Gao, Xu Han, Ruobing Xie, Zhiyuan Liu, Fen Lin, Leyu Lin, and Maosong Sun. 2020. Neural snowball for few-shot relation learning. In *Proceedings of AAAI Conference on Artificial Intelligence*.
- Tianyu Gao, Xu Han, Hao Zhu, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. 2019b. FewRel 2.0: Towards more challenging few-shot relation classification. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, pages 6250–6255.
- Victor Garcia and Joan Bruna. 2018. Few-shot learning with graph neural networks. In *Proceedings of the 2018 Conference on Learning Representations*.
- Łukasz Garncarek, Rafał Powalski, Tomasz Stanisławek, Bartosz Topolski, Piotr Halama, Michał Turski, and Filip Graliński. 2021. Lambert: Layout-aware language modeling for information extraction. In *Document Analysis and Recognition – ICDAR 2021*, pages 532–547.
- Jiaming Han, Yuqiang Ren, Jian Ding, Ke Yan, and Gui-Song Xia. 2023. Few-shot object detection via variational feature aggregation. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence, AAAI'23/IAAI'23/EAAI'23*. AAAI Press.
- 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 *the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4803–4809.
- Teakgyu Hong, Donghyun Kim, Mingi Ji, Won-seok Hwang, Daehyun Nam, and Sungrae Park. 2021. Bros: A pre-trained language model focusing on text and layout for better key information extraction from documents. In *AAAI Conference on Artificial Intelligence*.
- Jiaxin Huang, Chunyuan Li, Krishan Subudhi, Damien Jose, Shobana Balakrishnan, Weizhu Chen, Baolin Peng, Jianfeng Gao, and Jiawei Han. 2020. Few-shot named entity recognition: A comprehensive study. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*.
- Yupan Huang, Tengchao Lv, Lei Cui, Yutong Lu, and Furu Wei. 2022. Layoutlmv3: Pre-training for document ai with unified text and image

- masking. In *Proceedings of the 30th ACM International Conference on Multimedia*, page 4083–4091, New York, NY, USA. Association for Computing Machinery.
- Rashidul Islam and James Foulds. 2019. [Scalable collapsed inference for high-dimensional topic models](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2836–2845, Minneapolis, Minnesota. Association for Computational Linguistics.
- Anoop R Katti, Christian Reisswig, Cordula Guder, Sebastian Brarda, Steffen Bickel, Johannes Höhne, and Jean Baptiste Faddoul. 2018. [Char-grid: Towards understanding 2D documents](#). In *the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4459–4469.
- Diederik P Kingma and Max Welling. 2013. [Auto-encoding variational bayes](#). In *International Conference on Learning Representations*.
- Gregory Koch, Richard Zemel, Ruslan Salakhutdinov, et al. 2015. [Siamese neural networks for one-shot image recognition](#). In *ICML*.
- Shantanu Kumar. 2017. [A survey of deep learning methods for relation extraction](#). *ACM Comput. Surv.*
- Chen-Yu Lee, Chun-Liang Li, Timothy Dozat, Vincent Perot, Guolong Su, Nan Hua, Joshua Ainslie, Renshen Wang, Yasuhisa Fujii, and Tomas Pfister. 2022. [FormNet: Structural encoding beyond sequential modeling in form document information extraction](#). In *the Annual Meeting of the Association for Computational Linguistics*, pages 3735–3754.
- Chen-Yu Lee, Chun-Liang Li, Hao Zhang, Timothy Dozat, Vincent Perot, Guolong Su, Xiang Zhang, Kihyuk Sohn, Nikolay Glushnev, Renshen Wang, Joshua Ainslie, Shangbang Long, Siyang Qin, Yasuhisa Fujii, Nan Hua, and Tomas Pfister. 2023. [FormNetV2: Multimodal graph contrastive learning for form document information extraction](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9011–9026, Toronto, Canada. Association for Computational Linguistics.
- Xin Li, Yan Zheng, Yiqing Hu, Haoyu Cao, Yunfei Wu, Deqiang Jiang, Yinsong Liu, and Bo Ren. 2022. [Relational representation learning in visually-rich documents](#). In *ACM MM*, page 4614–4624.
- Yulin Li, Yuxi Qian, Yuechen Yu, Xiameng Qin, Chengquan Zhang, Yan Liu, Kun Yao, Junyu Han, Jingtuo Liu, and Errui Ding. 2021. [Structext: Structured text understanding with multi-modal transformers](#). In *ACM MM*, page 1912–1920.
- Weihong Lin, Qifang Gao, Lei Sun, Zhuoyao Zhong, Kai Hu, Qin Ren, and Qiang Huo. 2021. [Vibertgrid: A jointly trained multi-modal 2d document representation for key information extraction from documents](#). In *ICDAR*, pages 548–563.
- Xiaojing Liu, Feiyu Gao, Qiong Zhang, and Huasha Zhao. 2019. [Graph convolution for multimodal information extraction from visually rich documents](#). In *NAACL*, pages 32–39.
- Yang Liu, Jinpeng Hu, Xiang Wan, and Tsung-Hui Chang. 2022. [Learn from relation information: Towards prototype representation rectification for few-shot relation extraction](#). In *Findings of NAACL*, pages 1822–1831.
- Bin Lu, Xiaoying Gan, Lina Yang, Weinan Zhang, Luoyi Fu, and Xinbing Wang. 2022. [Geometer: Graph few-shot class-incremental learning via prototype representation](#). In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, page 1152–1161.
- Seunghyun Park, Seung Shin, Bado Lee, Junyeop Lee, Jaeheung Surh, Minjoon Seo, and Hwalsuk Lee. 2019. [Cord: A consolidated receipt dataset for post-ocr parsing](#). In *Document Intelligence Workshop at NeurIPS*.
- Nicholas Popovic and Michael Färber. 2022. [Few-shot document-level relation extraction](#). In *North American Chapter of the Association for Computational Linguistics*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#).
- Dawid Rymarczyk, Łukasz Struski, Jacek Tabor, and Bartosz Zieliński. 2021. [Protopshare: Prototypical parts sharing for similarity discovery in interpretable image classification](#). In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, page 1420–1430.
- Ofer Sabo, Yanai Elazar, Yoav Goldberg, and Ido Dagan. 2021. [Revisiting few-shot relation classification: Evaluation data and classification schemes](#). *Transactions of the Association for Computational Linguistics*, 9:691–706.

- Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, page 4080–4090, Red Hook, NY, USA.
- Shengli Sun, Qingfeng Sun, Kevin Zhou, and Tengchao Lv. 2019. Hierarchical attention prototypical networks for few-shot text classification. In *the 2019 Conference on Empirical Methods in Natural Language Processing*, pages 476–485.
- Guozhi Tang, Lele Xie, Lianwen Jin, Jiapeng Wang, Jingdong Chen, Zhen Xu, Qianying Wang, Yaqiang Wu, and Hui Li. 2021. [Matchvie: Exploiting match relevancy between entities for visual information extraction](#). In *International Joint Conference on Artificial Intelligence*, pages 1039–1045.
- Van-Khanh Tran and Le-Minh Nguyen. 2018. [Adversarial domain adaptation for variational neural language generation in dialogue systems](#). In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1205–1217, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Maria Tsimpoukelli, Jacob Menick, Serkan Cabi, S. M. Ali Eslami, Oriol Vinyals, and Felix Hill. 2021. [Multimodal few-shot learning with frozen language models](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 200–212.
- Hao Wang, Qingxuan Wang, Yue Li, Changqing Wang, Chenhui Chu, and Rui Wang. 2023. [DocTrack: A visually-rich document dataset really aligned with human eye movement for machine reading](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5176–5189, Singapore. Association for Computational Linguistics.
- Jiapeng Wang, Lianwen Jin, and Kai Ding. 2022a. [LiLT: A simple yet effective language-independent layout transformer for structured document understanding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, pages 7747–7757.
- Pengyu Wang and Phil Blunsom. 2013. [Collapsed variational Bayesian inference for PCFGs](#). In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 173–182, Sofia, Bulgaria. Association for Computational Linguistics.
- Zihan Wang, Kewen Zhao, Zilong Wang, and Jingbo Shang. 2022b. [Formulating few-shot fine-tuning towards language model pre-training: A pilot study on named entity recognition](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3186–3199, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Zilong Wang and Jingbo Shang. 2022. [Towards few-shot entity recognition in document images: A label-aware sequence-to-sequence framework](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 4174–4186, Dublin, Ireland. Association for Computational Linguistics.
- Zilong Wang, Yiheng Xu, Lei Cui, Jingbo Shang, and Furu Wei. 2021. [LayoutReader: Pre-training of text and layout for reading order detection](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4735–4744, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Congying Xia, Caiming Xiong, Philip Yu, and Richard Socher. 2020. [Composed variational natural language generation for few-shot intents](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3379–3388, Online. Association for Computational Linguistics.
- Yang Xu, Yiheng Xu, Tengchao Lv, Lei Cui, Furu Wei, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, Wanxiang Che, Min Zhang, and Lidong Zhou. 2021. [LayoutLMv2: Multi-modal pre-training for visually-rich document understanding](#). In *the Annual Meeting of the Association for Computational Linguistics*.
- Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. 2020. [Layoutlm: Pre-training of text and layout for document image understanding](#). In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, page 1192–1200, New York, NY, USA. Association for Computing Machinery.
- Kaijia Yang, Nantao Zheng, Xinyu Dai, Liang He, Shujian Huang, and Jiajun Chen. 2020. [Enhance prototypical network with text descriptions for few-shot relation classification](#). In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, page 2273–2276.
- Yi Yang and Arzoo Katiyar. 2020. [Simple and effective few-shot named entity recognition with](#)

structured nearest neighbor learning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 6365–6375, Online. Association for Computational Linguistics.

Yuechen Yu, Yulin Li, Chengquan Zhang, Xiaoqiang Zhang, Zengyuan Guo, Xiameng Qin, Kun Yao, Junyu Han, Errui Ding, and Jingdong Wang. 2023. [Structextv2: Masked visual-textual prediction for document image pre-training](#). In *International Conference on Learning Representations*.

Junwei Zhang, Hao Wang, and Xiangfeng Luo. 2022. [Dual-vie: Dual-level graph attention network for visual information extraction](#). In *PRICAI 2022: Trends in Artificial Intelligence*, pages 422–434.

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](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 3999–4006. International Joint Conferences on Artificial Intelligence Organization. Main Track.

## 11. Appendix

### 11.1. Relation Types

As shown in Table 5, SEAB is manually annotated with 4 coarse-grained and 22 fine-grained relation types, and we show all the types in Figure 7. Note that the partition of relation types is shown in the Table 5. Different colors is assigned according to their specific coarse-grained types.

	# Train	# Test
SEAB	ALL	ALL
Inter	1,5,7,8,9,10,11,12,13,14,15,18	2,3,4,6,16,17,19,20,21,22
Intra	1,2,3,9,10,11,12,13,14,15,16,19	4,5,6,7,8,17, 18,20,21,22

Table 5: Partition of relation types in SEAB dataset. Each number corresponds to a relation type. Specifically, 1-22: Shipper, Consignee, Notify Party, Marker, Number of Packages, Good Description, Gross Weight, Measurement, Shipping Terms, Place of Receipt, Port of Loading, Port of Discharger, Place of Delivery, Vessel Name, Voyage no, Consignment Code, Shipping Company, HSCODE, Freight Terms, Pre-Assignment, Case Size, Remarks.

Figure 7: Realistic examples with annotations from the SEAB dataset. These examples show all entity and relation types.