

Modeling Layout Reading Order as Ordering Relations for Visually-rich Document Understanding

Chong Zhang^{1,2*}, Yi Tu³, Yixi Zhao^{1,2}, Chenshu Yuan⁴, Huan Chen³, Yue Zhang^{1,2}, Mingxu Chai^{1,2}, Ya Guo³, Huijia Zhu³, Qi Zhang^{1,2†}, Tao Gui^{2,5†}

¹School of Computer Science, Fudan University, Shanghai, China

²Shanghai Key Laboratory of Intelligent Information Processing, Shanghai, China

³Ant Tiansuan Security Lab, Ant Group, Hangzhou, China

⁴School of Statistics and Data Science, Nankai University, Tianjin, China

⁵Institute of Modern Languages and Linguistics, Fudan University, Shanghai, China

{chongzhang20, qz, tgui}@fudan.edu.cn, {qianyi.ty, guoya.gy}@antgroup.com

Abstract

Modeling and leveraging layout reading order in visually-rich documents (VrDs) is critical in document intelligence as it captures the rich structure semantics within documents. Previous works typically formulated layout reading order as a permutation of layout elements, i.e. a sequence containing all the layout elements. However, we argue that this formulation does not adequately convey the complete reading order information in the layout, which may potentially lead to performance decline in downstream VrD tasks. To address this issue, we propose to model the layout reading order as ordering relations over the set of layout elements, which have sufficient expressive capability for the complete reading order information. To enable empirical evaluation on methods towards the improved form of reading order prediction (ROP), we establish a comprehensive benchmark dataset including the reading order annotation as relations over layout elements, together with a relation-extraction-based method that outperforms previous methods. Moreover, to highlight the practical benefits of introducing the improved form of layout reading order, we propose a reading-order-relation-enhancing pipeline to improve model performance on any arbitrary VrD task by introducing additional reading order relation inputs. Comprehensive results demonstrate that the pipeline generally benefits downstream VrD tasks: (1) with utilizing the reading order information, the enhanced downstream models achieve SOTA results on both two task settings of the targeted dataset; (2) with utilizing the pseudo reading order information generated by the proposed ROP model, the performance of the enhanced models has improved across all three models and eight cross-domain VrD-IE/QA task settings without targeted optimization.

* Part of work is done during an internship at Ant Group.

† Corresponding author.

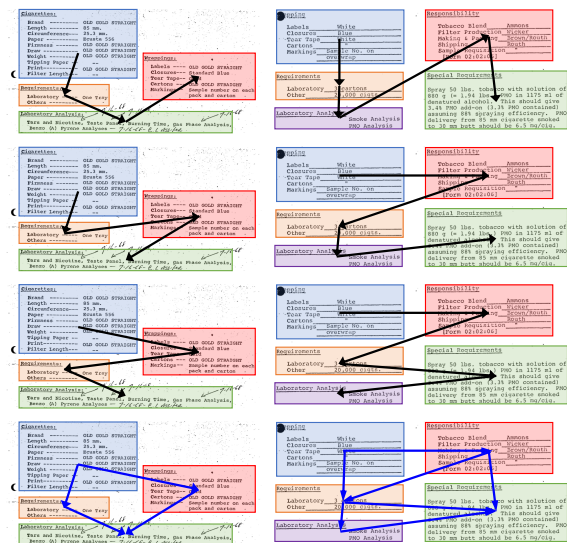


Figure 1: Motivation of reformulating layout reading order. In complex document layouts, multiple reading sequences are acceptable (displayed in the first three rows); thus the reading order information may be incomplete if represented by one single sequence. We propose to represent the relationship of *immediate succession during reading* among layout elements using a directed acyclic relation (displayed in the last row as a directed acyclic graph), ensuring that the complete layout reading order information is conveyed.

1 Introduction

The field of visually-rich documents (VrDs) has attracted widespread attention recently. Towards the growing industrial demands of automated processing of VrDs, a wide range of challenging tasks have arisen including understanding, classifying, extracting, and post-processing the information of VrDs (Harley et al., 2015; Huang et al., 2019; Wang et al., 2021, 2023e; Zhang et al., 2023b). Among those tasks, Reading Order Prediction (ROP) is a peculiar task on VrD layouts aiming to determine the human-acceptable reading order of layout elements, i.e. 2D-situated words or text regions (segments) given their bounding boxes.

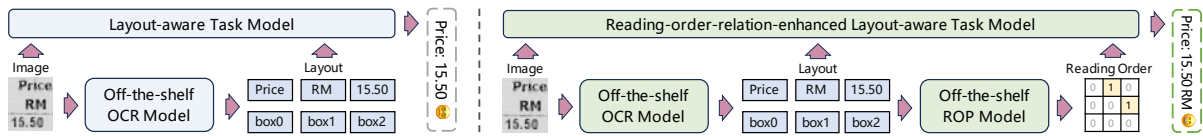


Figure 2: The reading-order-relation-enhancing pipeline (right, green) comparing with the original pipeline (left, blue) for general document processing. "RM" denotes Malaysian Ringgit.

ROP is a fundamental task for layout understanding as it reveals the influence of layouts for human reading strategies (Wright, 1999; Dyson, 2004). The task is challenging due to the diversity and complexity of layout designs among documents, and is also important in applications as previous works have demonstrated that the performance of downstream VrD tasks can be substantially improved by leveraging proper reading order information (Peng et al., 2022; Zhang et al., 2023a).

However, we argue that the prevailing task formulation of ROP is inappropriate as it does not convey the complete layout reading order information. In recent works, the layout reading order is typically represented as a permutation of layout elements, i.e. a sequence containing all the words or segments within the layout (Malerba and Ceci, 2007; Li et al., 2020). However, as shown in Fig. 1, there exists multiple permutations of the displayed layout that represents a human-acceptable reading order, and some of which conflict with other ones.¹ Therefore, **representing the layout reading order information with one single permutation is unable to convey the complete reading order information in the layout**, and could possibly inject noise in the annotations. Both issues negatively affects its application in downstream tasks. Given the problem, we propose two questions: (1) How to improve the formulation of reading order to include the complete information? (2) Would the improved formulation be truly helpful in practical applications that downstream VrD tasks could be further benefited by leveraging the improved form of reading order information?

To address the first question, we propose to model the layout reading order as ordering relations over the set of layout elements. Since there lacks a clear definition of layout reading order, we investigate its properties to conceptualize it with two terms, each of which is represented as an ordering relation: (1) Immediate Succession During Reading (ISDR) is modeled as a directed acyclic relation over layout elements; (2) Gen-

eralized Succession During Reading (GSDR) is modeled as the transitive closure of ISDR, which is a strict partial order relation over layout elements. ROP is then reformulated as a relation extraction task on layout elements. For developing automatic methods of ROP in the improved form, we also establish a new benchmark dataset annotating ISDR as relation pairs of layout segments together with a baseline relation-extraction-based method. Experiments demonstrate that the proposed method outperforms previous sequence-based methods by a large margin on ROP.

To address the second question, we propose the reading-order-relation-enhancing (RORE) pipeline which is confirmed to be effective by comprehensive experiments on downstream VrD-IE/QA tasks. For injecting reading order relation information, we propose to represent it in a $n * n$ binary matrix over n input textual tokens of PTLMs, and introduce a relation-aware attention mechanism to incorporate the reading order information. Experiment results show that with additional reading order inputs, the performance of three baseline PTLMs have shown universal improvement. Moreover, considering that ground truth annotation of reading order information is unavailable in real-world scenarios, it is expected to leverage pseudo reading order labels generated by an off-the-shelf ROP model in any arbitrary downstream tasks. The expectation is proven feasible by experiment results, since the performance of three reading-order-relation-enhanced models are comprehensively improved across eight prevailing benchmarks of VrD-IE/QA from different domains, achieving SOTA on two of them without targeted optimization.

Overall, the main contributions of this work are:

1. We identify the existing problem that the current task form of ROP fails to convey the complete layout reading order information, and suggest modeling the reading order as ordering relations over layout elements. We introduce a new comprehensive benchmark for quantitative evaluation of the improved formulation of ROP.
2. We propose a relation-extraction-based model which outperforms previous sequence-based methods on ROP. Further, the proposed model

¹In Fig. 1, annotations are drawn as block-level for better visualization, which is different from the proposed dataset.

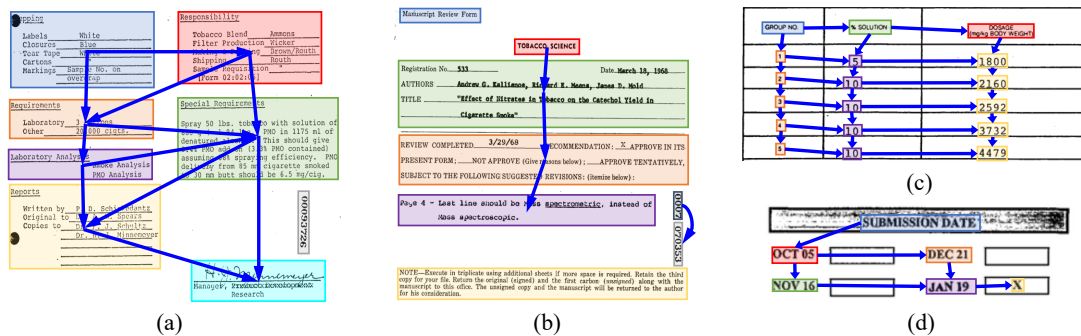


Figure 3: Several example layouts with non-linear reading order. Annotations are drawn as block-level for better visualization. (a) The complex layout includes multiple possible reading sequences (illustrated in Fig. 1); (b) The reading order of header, footer and watermark within the layout are separated from the main body; (c) The table within the layout can be read either vertically or horizontally; (d) Indirect reading order relationship is also important as relevant elements may be separated by other contents.

generate high-quality pseudo reading-order-relation labels for arbitrary inputs from various domains.

3. We introduce a pipeline to enhance arbitrary VrD tasks based on the proposed ROP model. Comprehensive results demonstrate that the performance of several prevailing methods on various downstream VrD-IE/QA tasks can be universally improved by leveraging the (probably pseudo) reading order relation information. The enhanced methods have achieved SOTA results on the targeted benchmark and also two cross-domain task settings.

The proposed dataset and code implementation of proposed models are released at <https://github.com/chongzhangFDU/ROOR>.

2 Reading Order as Ordering Relations

2.1 Conceptualization of Reading Order

This section conceptualizes the reading order of layouts by examining general human reading patterns. While layout reading order prediction is one of the most crucial task in document AI, the concept of reading order itself is rarely discussed and has not been fully clarified. We define the term **Immediate Succession During Reading (ISDR)** to precisely capture the human reading order on general textual contents.

Intuitively, the layout reading order is a manifestation of human reading strategies tackling different layout designs (Wright, 1999; Dyson, 2004). To precisely define the layout reading order, a formalized description of human reading behavior on layouts is required. For this purpose, we draw inspiration from human behavior of reading plain text. Plain text is always a verbatim sequence following its inherent order, where each word appears after the preceding one linearly. We refer

to each word as the immediate successor of its preceding word, since it is read immediately after the preceding word during reading. In layouts, when we read a layout element subsequently after the previous element, we call the layout element as the immediate succession of the previous one. Thereby, we introduce the term *immediate succession during reading* to describe the human reading order of general textual contents.

2.2 ISDR as a Directed Acyclic Relation

In this section, we illustrate that current task form of reading order prediction on layout is not suitable by examining the difference between human reading habits of plain text and layout. To address the problem, we suggest representing the ISDR relationship as a directed acyclic relation (DAR, formal definition in Appendix A) over layout elements, and reformulate the ROP task of layout as a relation extraction task.

As illustrated in §2.1, the reading order of plain text can be naturally represented as a permutation of words, i.e. the word sequence itself. Most previous works of layout reading order prediction have followed this manner and also represent the layout reading order as a permutation of layout elements. However, this formulation is unsuitable due to the following characteristics of layout: (1) Different to plain text where words are succeeded by one single subsequent word, layout elements within a document can have multiple immediate successors (Fig. 3a and 3c), and table cells can have successors to the right or below (Fig. 3c). (2) Some certain layout elements do not interact with others, and their successive relationships should also be separated from others (Fig. 3b).

To address these problems, we suggest representing the ISDR relationship as a DAR over

layout elements as an alternative formulation. Motivated by Fig. 3 which portrays ISDR as a directed graph of layout elements, we suggest conceptualizing ISDR as a binary relation over the set of layout elements. Given that ISDR represents a succession relationship, we further stipulate that this binary relation is acyclic to prevent any circular dependencies of layout elements. As a result, ISDR is formulated as a DAR over the layout elements.

2.3 Generalized Succession During Reading

In this section, we propose Generalized Succession During Reading (GSDR) as an alternative form of reading order information. GSDR focuses on the global awareness of reading order relation between layout elements, and is formulated as a strict partial order (SPO, formal definition in Appendix A) over the set of layout elements.

The ISDR relationship proposed in §2.1 conveys the complete reading order information, and it mainly focuses on the local reading order relationship between layout element. However, there are cases in reading where we focus on the global reading order information. For example, as illustrated in Fig. 3d, the question "SUBMISSION DATE" is always read before the answer "SEP 15", yet their reading order relationship is not directly represented by ISDR as the answer does not directly follow the question during reading. This motivates us to introduce the transitive closure of ISDR, namely *generalized succession during reading*, to convey the global subsequent relationship between layout segments. Specifically, an element being the generalized successor of another one indicates that this element is read after another element somehow in reading. From the order-theoretic perspective, the GSDR relationship over layout elements is a SPO, different from the GSDR relationship over words within plain text as a strict total order, indicating the different characteristics of plain text and layout. The application of GSDR is discussed in Appendix E.

2.4 Dataset Construction

In this section, we present **ROOR (Reading Order as Ordering Relations)**, a benchmark dataset featuring segment-level reading order relation annotations. This benchmark is constructed to facilitate the development of automated ROP methods of the improved task form.

The annotation of ROOR is based on EC-FUNSD (Zhang et al., 2024) due to its vari-

ety of document layouts and high-quality layout annotations. In EC-FUNSD, each sample is a single-page scanned VrD with manually-reviewed layout annotations emulating the standard output of word-and-segment annotations produced by OCR engines. Two domain experts independently annotated the ISDR relation as linking pairs of segments in each document, and resolved any conflicts between their annotations through discussion. To maintain the property of ISDR relationship as a DAR in each sample, the line-breaks within grid-lined tables were not annotated, i.e. the succession from the last cell in the previous row to the first cell in the next row. The same was applied to column-breaks. Overall, ROOR comprises 199 samples including 10,662 segments, 31,297 words and 10,967 annotated reading order linking pairs.

Statistics on ROOR indicate the prevalent occurrence of non-linear reading order within layouts. We define a segment as involved in non-linear reading order iff. both its preceding and succeeding segment uniquely exist (except for the initial and final segments which naturally lack either a predecessor or successor). Examination on ROOR shows that 23.76% of segments within the dataset are involved in non-linear reading order, demonstrating the significance to deal with this phenomenon in layout processing.

3 Methodologies

3.1 Reading Order Prediction as Relation Extraction

As we model the layout reading order as ordering relations, the reading order prediction task is then reformulated as a relation extraction task as described below: A document layout with $N_{\mathcal{D}}$ layout elements (i.e. words or segments) is represented as $\mathcal{D} = \{(w_i, \mathbf{b}_i)\}_{i=1, \dots, N_{\mathcal{D}}}$, where w_i denotes the i -th element in document and $\mathbf{b}_i = (x_i^0, y_i^0, x_i^1, y_i^1)$ denotes the position of w_i in the document layout. The coordinates (x_i^0, y_i^0) and (x_i^1, y_i^1) correspond to the bottom-left and top-right vertex of w_i 's bounding box, respectively. The set of reading order relation pair of elements from $s_{\mathcal{D}}$ is denoted as $p_{\mathcal{D}} = \{p_1, \dots, p_K\}$, where the k -th relation pair $p_k = (ks, ko)$ indicates that w_{ko} is an immediate successor of w_{ks} , ($1 \leq ks, ko \leq N_{\mathcal{D}}$). The aim of ROP is to recognize all the reading order relation pairs within the document.

3.2 A Baseline Model for ROP

After treating the task as relation extraction on VrD layouts, a baseline model is then introduced. To avoid over-engineering for the task, we straightforwardly adopt the global pointer network (Su et al., 2022), a widely used architecture for relation extraction tasks in NLP.

For a document layout \mathcal{D} with N layout elements, let L be the ground truth label of reading order relation, where $(i, j) \in L$ indicates that the j -th element is an immediate successor of the i -th element. The tokenized input of this document is denoted as $(x_1^1, \dots, x_1^{n_1}, \dots, x_N^1, \dots, x_N^{n_N})$, where x_i^j indicates the j -th token of the i -th layout element. Let \mathbf{b}_i^j be the corresponding bounding box of x_i^j , the tokenized input is then fed to the backbone encoder to obtain the layout-aware text embeddings:

$$h_1^1, \dots, h_N^{n_N} = \text{PTLM}(x_1^1, \dots, x_N^{n_N}; \mathbf{b}_1^1, \dots, \mathbf{b}_N^{n_N}) \quad (1)$$

where PTLM is the encoder of text and layout inputs; h_i^j is the layout-aware embedding of x_i^j . Layout element embeddings are then calculated by gathering their corresponding token embeddings:

$$h_i = \text{AvgPool1d}(h_i^1, \dots, h_i^{n_i}) \quad (2)$$

where h_i is the embedding of the i -th element. Then the embeddings are fed into the global pointer network for relation extraction:

$$\begin{aligned} q_i &= \mathbf{W}_q h_i + \mathbf{b}_q \\ k_i &= \mathbf{W}_k h_i + \mathbf{b}_k \\ s_{ij} &= q_i^T k_j \end{aligned} \quad (3)$$

where $\mathbf{W}_{q,k}$ and $\mathbf{b}_{q,k}$ are learnable parameters; s_{ij} is the predicted score of $(i, j) \in L$.

Model weights are optimized by a class-imbalance loss which is also introduced in (Su et al., 2022):

$$\mathcal{L} = \log \left(1 + \sum_{\substack{(i,j) \notin L \\ 1 \leq i, j \leq N}} e^{s_{ij}} \right) + \log \left(1 + \sum_{\substack{(i,j) \in L \\ 1 \leq i, j \leq N}} e^{-s_{ij}} \right) \quad (4)$$

During inference, the model identifies reading order relation pairs by filtering the predictions with a threshold:

$$\hat{L} = \{(i, j) | s_{ij} > 0\} \quad (5)$$

where \hat{L} is the set of predicted reading order relation pairs.

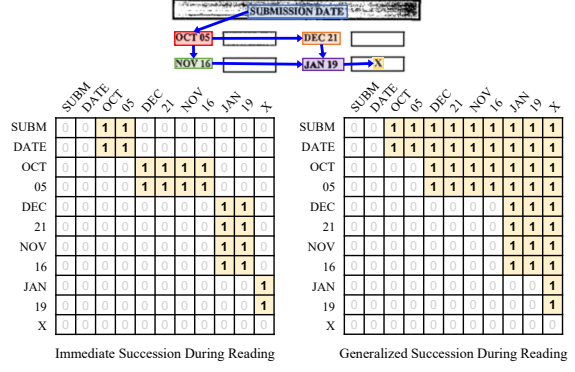


Figure 4: Reading order relation information is represented as a $n * n$ binary matrix to be leveraged in downstream VrD tasks, where n is the number of input textual tokens.

3.3 Enhancing Layout-Aware Models with Reading Order Relation

In this section, it is assumed that the reading order information for a document layout \mathcal{D} is available as a relation between its N layout elements. In applications, it could be manually annotated or automatically generated by an off-the-shelf ROP model (see Fig. 2).

To leverage the reading order information in downstream VrD tasks, we first represent the information in a $n * n$ binary matrix, where n is the number of input textual tokens. For each layout element, the pair from each token of itself to each token of its (immediate or generalized) successors is labeled 1 in the matrix, while others are labeled as 0. Fig. 4 displays an example layout whose ISDR and GSDR reading order information are represented as binary matrices. The matrix is then used in relation-aware attention modules of downstream layout-aware models.

For transformer-based models, the attention matrix within their original self-attention modules is calculated as follows:

$$a_{ij}^l = \frac{\exp(\mathbf{q}_i^{lT} \mathbf{k}_j^l / \sqrt{d_k})}{\sum_{j=1}^n \exp(\mathbf{q}_i^{lT} \mathbf{k}_j^l / \sqrt{d_k})} \quad (6)$$

where $\{a_{ij}^l\}_{1 \leq i, j \leq n}$ is the attention matrix of a self-attention module in the l -th layer; $\mathbf{q}_i^l, \mathbf{k}_i^l$ are the i -th query and key vector in the l -th layer, respectively; d_k is the dimension of the attention head.

We propose a relation-aware attention module as a replacement to the original self-attention module, leveraging the reading order relation matrix in the simplest manner. This module straightly adds the reading order relation matrix to the original

Method	Word-level	Segment-level
LR (Wang et al., 2021)	-	9.44
TPP (Zhang et al., 2023a)	-	42.96
LayoutLMv3-base	90.26	68.60
LayoutLMv3-large	93.01	82.38
Human	-	99.28

Table 1: The performance of baseline models on reading order relation prediction. Human performance indicates the annotation consistency between two annotators.

attention matrix to incorporate the information. The attention matrix within the relation-aware attention module is calculated as:

$$a_{ij}^l = \frac{\exp\left(\left(\mathbf{q}_i^{lT} \mathbf{k}_j^l + \lambda^l \rho_{ij}\right) / \sqrt{d_k}\right)}{\sum_{j=1}^n \exp\left(\left(\mathbf{q}_i^{lT} \mathbf{k}_j^l + \lambda^l \rho_{ij}\right) / \sqrt{d_k}\right)} \quad (7)$$

where $\{\rho_{ij}^l\}_{1 \leq i, j \leq n}$ is the reading order relation matrix and λ^l is a learnable scalar as the weight of reading order information in the l -th layer.

4 Experiments

4.1 Experimental Settings

Reading Order Prediction as Relation Extraction We test the proposed baseline model with word-level and segment-level reading order prediction tasks of ROOR. We take the ISDR annotations as the segment-level labels, and utilize the arranged words within segments together with the segment-level annotations to generate word-level labels. We adopt LayoutLMv3-base/large (Huang et al., 2022) as the backbone text-and-layout encoder for the proposed ROP model. In evaluation, we follow the train/validation split from EC-FUNSD and models are evaluated by F1 score.

Enhancing Downstream VrD Tasks with Reading Order Relation To examine the effectiveness of reading order relation information in enhancing downstream VrD tasks, we enhance pre-trained text-and-layout models (PTLMs) with additional reading order inputs. Two baseline models are examined: LayoutLMv3 as the most popular PTLM, and GeoLayoutLM (Luo et al., 2023) as a representative relation-aware PTLM. We assess the baseline models and their reading-order-relation-enhanced (RORE) variants on semantic entity recognition (SER) and entity linking (EL) of EC-FUNSD, leveraging the ground truth reading order information in the ROOR dataset.

Quality of Pseudo Reading Order Relation Labels for Downstream Tasks Enhancing an arbitrary VrD task with additional reading order

Method	SER	+RORE	EL	+RORE
LayoutLMv3-base	82.30	82.80 ($\uparrow 0.50$)	67.47	73.64 ($\uparrow 6.17$)
LayoutLMv3-large	83.88*	84.53 ($\uparrow 0.65$)	78.14*	79.33 ($\uparrow 1.19$)
GeoLayoutLM	83.62	84.34 ($\uparrow 0.72$)	86.18	87.42 ($\uparrow 1.24$)

Table 2: Performance of LayoutLMv3 and GeoLayoutLM and their corresponding RORE methods on EC-FUNSD. Reproduced results are marked with * (see Appendix D).

relation inputs is expected. However, it is impractical to manually annotate each sample with such information. Consequently, a robust off-the-shelf ROP model is expected to generalize well to any target domain and generate high-quality pseudo reading order relation labels for unseen datasets, since VrDs in real application could cover various domains and have high diversity in layouts.

To validate the effectiveness of the generated pseudo labels, we conduct experiments on VrD-IE/QA tasks with utilizing the pseudo reading order information generated by an off-the-shelf LayoutLMv3-large model solely trained on the ROOR dataset. For IE, we adopt three popular benchmarks for SER and EL: FUNSD (Jaume et al., 2019), CORD (Park et al., 2019) and SROIE (Huang et al., 2019). For QA, we choose four popular benchmarks that primarily featuring on layout understanding and IE, but not commonsense or complex reasoning which are less relevant to layout understanding: DocVQA (Mathew et al., 2021), InfoVQA (Mathew et al., 2022), WTQ (Borchmann et al., 2021) and TextVQA (Singh et al., 2019). A generative variant of LayoutLMv3 with a BART decoder (Lewis et al., 2019) is used as the baseline model, the performance of which is evaluated by average normalized Levenshtein similarity (ANLS). We report the performance of baseline and RORE models on these datasets. The implementation details of the above experiments are illustrated in Appendix D.

4.2 Performance of Baseline ROP Models

We compare the proposed relation-based baseline models with previous sequence-based models which predict a layout element sequence as the reading order (Wang et al., 2021; Zhang et al., 2023a). We report their performance in Tab. 1 together with human performance. Without task-specific design, the relation-based models are capable to learn the reading order relation signals, which reveals their potential to generate pseudo reading order annotations for other datasets. However, the sequence-based models struggle with incomplete supervision and demonstrate a signifi-

Method		#Params	FUNSD		CORD	SROIE	Ref. Overall Performance
			SER	EL	SER	SER	
Attention	SPADE (Hwang et al., 2021)	base	72.0	41.3	-	-	64.95
	SERA (Zhang et al., 2021)	base	-	65.96	-	-	62.91
	LiLT (Wang et al., 2022b)	base	88.41	62.76	96.07	-	85.56
	DocTr (Liao et al., 2023)	153M	84.0	73.9	98.2	-	87.87
	BROS (Hong et al., 2022)	340M	84.52	77.01	97.28	96.62	88.85
	LAGaBi (Zhu et al., 2023)	133M	91.00	-	97.05	-	87.72
	LayoutLMv3 (Huang et al., 2022)	133M	90.85	69.80	95.91*	94.80*	87.84
	LayoutLMv3 (Huang et al., 2022)	368M	91.70*	79.37*	96.98*	96.12*	91.04
Graph	Doc2Graph (Gemelli et al., 2022)	base	-	53.36	-	-	50.89
	RE ² (Ramu et al., 2024)	base	-	71.76	-	-	68.44
	MatchVIE (Tang et al., 2021)	base	81.33	-	-	96.57	83.17
	DocFormer (Appalaraju et al., 2021)	502M	84.55	-	96.99	-	84.96
	FormNet (Lee et al., 2022)	large	84.69	-	97.28	-	85.17
	mmLayout (Wang et al., 2022c)	large	86.49	-	97.38	97.91	86.34
	GraphDoc (Zhang et al., 2022)	265M	87.77	-	96.93	98.45	86.76
	DocGraphLM (Wang et al., 2023a)	base	88.77	-	96.93	-	86.71
GraphLayoutLM (Li et al., 2023)	372M	-	-	97.98*	-	86.17	
Pre-trained	StrucTexT (Li et al., 2021)	107M	83.09	44.1	-	96.88	79.62
	DocReL (Li et al., 2022)	142M	-	46.1	97.0	-	80.19
	ERNIE-Layout (Peng et al., 2022)	large	93.12	-	97.21	97.55	88.83
	Bi-VLDoc (Luo et al., 2022a)	409M	93.44	-	97.84	-	89.15
	LayoutMask (Tu et al., 2023)	404M	93.20	-	97.19	97.27	88.81
	Wukong-Reader (Bai et al., 2023)	470M	93.62	-	97.27	98.15	89.14
	GeoLayoutLM (Luo et al., 2023)	399M	91.10	88.06	98.11*	96.62*	93.47
Supervised Pre-trained	ESP (Yang et al., 2023)	50M	91.12 [†]	88.88 [†]	95.65 [†]	-	92.43
	UNER (Tu et al., 2024)	base	-	-	-	97.61 [†]	85.99
Ours	RORE-LayoutLMv3	133M+12	91.39 (\uparrow 0.54)	71.69 (\uparrow 1.89)	96.72 (\uparrow 0.81)	95.65 (\uparrow 0.85)	88.86 (\uparrow 1.02)
		368M+24	92.14 (\uparrow 0.44)	80.84 (\uparrow 1.47)	97.29 (\uparrow 0.31)	96.59 (\uparrow 0.47)	91.71 (\uparrow 0.67)
	RORE-GeoLayoutLM	399M+24	91.84 (\uparrow 0.74)	88.46 (\uparrow 0.40)	98.52 (\uparrow 0.41)	96.97 (\uparrow 0.35)	93.94 (\uparrow 0.47)

Table 3: Performance of prevailing methods on three VrD-IE benchmarks. Best results are marked bold. Reproduced results are marked with * (see Appendix D). Supervised pre-trained results (marked with †) are not directly comparable to our methods and are listed as reference (see Appendix C). "Base" and "large" models have 100~200M and 300~400M parameters, respectively. Referential overall performance is calculated by averaging on the imputed score matrix using nuclear norm minimization.

cant performance gap. We have also found that: (1) The proposed model design benefits from stronger pre-trained backbones, as the performance of the LayoutLMv3-large model is much better than the base model in Tab. 1. (2) The proposed models generally perform better when the task level and the granularity of bounding box inputs coincide, illustrated by an ablation study in Appendix E. (3) The proposed models are good at identifying non-linear cases, yet showing the tendency to over-rely on layout but not textual features, demonstrated by visualized case studies in Appendix F.

4.3 Reading Order Relation Enhanced Methods for Downstream VrD Tasks

As illustrated in §3.3, we incorporate the reading order relation information to baseline models with

Task	DocVQA	+RORE	InfoVQA	+RORE
Score	69.13	73.53 (\uparrow 4.40)	22.62	29.20 (\uparrow 6.48)
Task	WTQ	+RORE	TextVQA	+RORE
Score	39.94	41.71 (\uparrow 1.77)	45.15	45.24 (\uparrow 0.09)

Table 4: Performance of the baseline generative method and RORE variant on four VrD-QA benchmarks.

the relation-aware attention module introduced in Eqn. 7. In Tab. 2, 3 and 4, we compare the performance of RORE methods on VrD-IE/QA tasks with their corresponding baseline methods, together with other competitive methods of different categories.

We conclude three general trends from the experiment results, demonstrating the effectiveness of the proposed methods: (1) **The performance of downstream tasks are universally improved**

by proper leveraging the reading order relation information. As shown in Tab. 2, the performance of three RORE methods have shown remarkable improvements and achieved SOTA on both two tasks. Significantly, these models benefit from a +2.86 average performance improvement on EL.

(2) Our ROP model is able to generalize to unseen datasets. Tab. 3 displays that all of the four task settings over three baseline methods have substantially improved by RORE, with only leveraging the pseudo labels generated by an off-the-shelf ROP model. Tab. 4 demonstrates a similar trend on generative VrD-QA tasks. The results demonstrate that our pipeline with an off-the-shelf ROP model trained on ROOR is capable to generate usable reading order labels for any arbitrary VrD tasks, without any additional manual annotations.

(3) The proposed RORE methods are very competitive among various tasks. Although we do not pursue exceedingly good performance on cross-domain settings, the RORE methods achieve SOTA on two of them. Especially, RORE-GeoLayoutLM achieves the best comprehensive performance among all prevailing methods on VrD-IE in Tab. 3. Besides, the feasibility of using pseudo reading order labels is further verified in Appendix E.

Other notable findings according to the experiment results include: (1) EL generally benefits more than SER by leveraging the reading order relation information, according to Tab. 2 and 3. Compared with SER which focuses on capturing entity semantics, EL relies more on the modeling of logical structure and spacial relationships of the layout, which is fully modeled by the reading order relation information. (2) The enhancement brought by RORE is particularly critical in low-resource scenarios. Among the baseline methods, LayoutLMv3-base initially exhibits the worst performance, yet it achieves the greatest improvement after incorporating the reading order relation information, revealing the potential of our method in low-resource scenarios. (3) The performance improvement through RORE should be primarily attributed to the auxiliary input information, rather than the relation-aware architecture. Among the VrD-QA benchmarks, TextVQA is unique from the others with its relatively simple layout patterns. Most samples in TextVQA are book covers, within whose layouts the text regions are scarce and generally adhere to a from-top-to-bottom reading order, thus the reading order

relation inputs are less informative. According to Tab. 4, TextVQA benefits less from RORE than the other benchmarks, from which we conclude that RORE primarily benefits from the additional inputs but not the architecture. (4) ISDR and GSDR benefits the downstream tasks to a similar extent, discussed by an ablation study in Appendix E. Summarizing all the above points, we demonstrate our final conclusion: **the reading order relation information is an informative auxiliary feature, from which RORE methods learn reliable and robust human reading patterns, attributing to their effectiveness.** The conclusion is further verified by a case study in Appendix F.

5 Related Work

Modeling of Layout Reading Order The assumption that layout reading order could be represented as a permutation on layout elements can be traced back to (NAGY, 1984). Since then, rule-based methods (Ishitani, 2003; Ferilli et al., 2014), statistical-ML-based methods (Han et al., 2007; Malerba and Ceci, 2007) and deep-learning-based methods (Li et al., 2020; Zhang et al., 2023a) have been proposed. However, several previous works have noted that layout reading order information cannot be fully conveyed by one single permutation of layout elements. These works improved their utilization of reading order information according to specific application requirements (Gu et al., 2022; Tu et al., 2023). However, these works do not delve further into the underlying problem that the formulation of layout reading order should be improved.

In this work, we model the reading order as ordering relations over layout elements, inspired by previous works which modeled the non-linear structures of VrD layouts as binary relations (Aiello and Smeulders, 2003, 2004) or ordered structures (Clausner et al., 2013), and also order-theoretic modeling of linguistic objects (Vendrov et al., 2015; Athiwaratkun and Wilson, 2018; Liu et al., 2023). A more comprehensive introduction of the evolution of modeling and utilizing layout reading order is available at Appendix B.

Modeling Relation between Layout Elements Several layout-aware models have devised relation-aware modules for better modeling VrD layouts. Zhao et al. (2022); Pham et al. (2022) enhanced local awareness of layout tokens with specialized attention mechanisms, and Chen et al. (2023)

introduced a special module for better modeling the global relationships. More works attempted to better leverage the relative positional relationship between layout elements by improving attention mechanism (Xu et al., 2021a,b; Powalski et al., 2021; Hwang et al., 2021; Nguyen et al., 2021; Zhang et al., 2021; Hong et al., 2022; Zhu et al., 2023; Siyuan et al., 2024), utilizing graph structure (Liu et al., 2019; Luo et al., 2020; Carbonell et al., 2021; Gemelli et al., 2022; Luo et al., 2022b; Mohbat et al., 2023; Ramu et al., 2024; Krishnan et al., 2024), or relation-aware pre-training (Li et al., 2022; Wang and Shang, 2022; Peng et al., 2022; Luo et al., 2023; Yu et al., 2023).

6 Conclusion

In this paper, we investigate reading order prediction for visually-rich document layouts, a classic and important task of document AI. We question the traditional formulation of this task as it is unable to convey the complete reading order information within the layouts, indicating that the improper formulation would possibly affect model performance on downstream VrD tasks. To address the problem, we propose an improved formulation of this task by conceptualizing layout reading order as ordering relations. We further prove the effectiveness of utilizing the improved form of reading order information in downstream tasks by extensive experiments. The proposed reading-order-relation-enhanced methods achieve universal performance improvement on multiple task settings of VrD-IE and VrD-QA, including SOTA performance on two targeted settings and two cross-domain settings. To facilitate the development of automated methods on ROP in the improved form, we also establish a benchmark dataset with high-quality reading order relation annotations together with a baseline model, which builds a pipeline to generate usable pseudo reading order labels for any arbitrary downstream tasks.

Limitations

Our work has the following limitations:

1. There is still room for improving the design of the reading order relation prediction model. In this paper, we propose a baseline relation extraction model which is constituted by an encoding pre-trained text-and-layout representation model and a decoding global pointer network, achieving fairly good result

on the reformulated ROP. However, the optimal solution of this problem is still under explored. Appendix F have studied the strengths and weaknesses of the proposed model, which we hope would shed light on the future works attempting to improve the model design.

2. Several methods for improving VrD understanding with reading order relation information remain insufficiently explored. In §5, we have categorized the relation-aware VrD understanding methods to attention-based, graph-based and relation-aware pre-training-based methods. In this paper, we have attempted to fuse the reading order information via a relation-aware attention module, demonstrating that our method improves the performance on downstream tasks. However, other ways to leverage the reading order information are not attempted and compared with the proposed method. It is encouraged for future works to make attempts on efficiently leverage the reading order relation information with ROOR and the proposed experiment settings.

Besides, it is necessary to point out that the improved form of ROP has increased the complexity of this task, as the search space of valid answers has increased from $O(n!)$ to $O(2^{n^2})$ by introducing the reading order relation, where n denotes the number of layout elements. We regard the complexity increase as an essential budget for the precise modeling of layout reading order.

Acknowledgements

The authors wish to thank the anonymous reviewers for their helpful comments. This work was partially funded by National Natural Science Foundation of China (No.62376061,62206057), Shanghai Rising-Star Program (23QA1400200), Natural Science Foundation of Shanghai (23ZR1403500), Program of Shanghai Academic Research Leader under grant 22XD1401100, and Ant Group Research Intern Program.

References

- Marco Aiello and Arnold M.W. Smeulders. 2003. *Bidimensional Relations for Reading Order Detection*. University of Groningen, Johann Bernoulli Institute for Mathematics and Computer Science.
- Marco Aiello and Arnold MW Smeulders. 2004. Thick 2d relations for document understanding. *Information Sciences*, 167(1-4):147–176.

- James F Allen. 1983. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843.
- Srikar Appalaraju, Bhavan Jasani, Bhargava Urala Kota, Yusheng Xie, and R Manmatha. 2021. Docformer: End-to-end transformer for document understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 993–1003.
- Ben Athiwaratkun and Andrew Gordon Wilson. 2018. Hierarchical density order embeddings. In *6th International Conference on Learning Representations, ICLR 2018*.
- Haoli Bai, Zhiguang Liu, Xiaojun Meng, Li Wentao, Shuang Liu, Yifeng Luo, Nian Xie, Rongfu Zheng, Liangwei Wang, Lu Hou, et al. 2023. Wukong-reader: Multi-modal pre-training for fine-grained visual document understanding. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13386–13401.
- Łukasz Borchmann, Michał Pietruszka, Tomasz Stanisławek, Dawid Jurkiewicz, Michał Turski, Karolina Szyndler, and Filip Graliński. 2021. Due: End-to-end document understanding benchmark. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Manuel Carbonell, Pau Riba, Mauricio Villegas, Alicia Fornés, and Josep Lladós. 2021. Named entity recognition and relation extraction with graph neural networks in semi structured documents. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 9622–9627. IEEE.
- Michelangelo Ceci, Margherita Berardi, G Porcelli, and Donato Malerba. 2007. A data mining approach to reading order detection. In *Ninth International Conference on Document Analysis and Recognition (ICDAR 2007)*, volume 2, pages 924–928. IEEE.
- Xiangnan Chen, Juncheng Li, Duo Dong, Qian Xiao, Jun Lin, Xiaozhong Liu, and Siliang Tang. 2023. Global structure knowledge-guided relation extraction method for visually-rich document. *arXiv preprint arXiv:2305.13850*.
- Christian Clausner, Stefan Pletschacher, and Apostolos Antonacopoulos. 2013. The significance of reading order in document recognition and its evaluation. In *2013 12th International Conference on Document Analysis and Recognition*, pages 688–692. IEEE.
- Mary C Dyson. 2004. How physical text layout affects reading from screen. *Behaviour & information technology*, 23(6):377–393.
- Stefano Ferilli, Domenico Grieco, Domenico Redavid, and Floriana Esposito. 2014. Abstract argumentation for reading order detection. In *Proceedings of the 2014 ACM symposium on Document engineering*, pages 45–48.
- Andrea Gemelli, Sanket Biswas, Enrico Civitelli, Josep Lladós, and Simone Marinai. 2022. Doc2graph: a task agnostic document understanding framework based on graph neural networks. *arXiv preprint arXiv:2208.11168*.
- Zhangxuan Gu, Changhua Meng, Ke Wang, Jun Lan, Weiqiang Wang, Ming Gu, and Liqing Zhang. 2022. Xylayoutlm: Towards layout-aware multimodal networks for visually-rich document understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4583–4592.
- Jaekyu Ha, Robert M Haralick, and Ihsin T Phillips. 1995. Recursive xy cut using bounding boxes of connected components. In *Proceedings of 3rd International Conference on Document Analysis and Recognition*, volume 2, pages 952–955. IEEE.
- Eunjung Han, Kirak Kim, HwangKyu Yang, and Keechul Jung. 2007. Frame segmentation used mlp-based xy recursive for mobile cartoon content. In *Human-Computer Interaction. HCI Intelligent Multimodal Interaction Environments: 12th International Conference, HCI International 2007, Beijing, China, July 22-27, 2007, Proceedings, Part III 12*, pages 872–881. Springer.
- Adam W Harley, Alex Ufkes, and Konstantinos G Derpanis. 2015. Evaluation of deep convolutional nets for document image classification and retrieval. In *2015 13th International Conference on Document Analysis and Recognition (ICDAR)*, pages 991–995. IEEE.
- Teakgyu Hong, Donghyun Kim, Mingi Ji, Wonseok Hwang, Daehyun Nam, and Sungrae Park. 2022. Bros: A pre-trained language model focusing on text and layout for better key information extraction from documents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 10767–10775.
- 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 *MM '22: The 30th ACM International Conference on Multimedia, Lisboa, Portugal, October 10 - 14, 2022*, pages 4083–4091. ACM.
- Zheng Huang, Kai Chen, Jianhua He, Xiang Bai, Dimosthenis Karatzas, Shijian Lu, and CV Jawahar. 2019. Icdar2019 competition on scanned receipt ocr and information extraction. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 1516–1520. IEEE.
- Wonseok Hwang, Jinyeong Yim, Seunghyun Park, Sohee Yang, and Minjoon Seo. 2021. Spatial dependency parsing for semi-structured document information extraction. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 330–343.

- Yasuto Ishitani. 2003. Document transformation system from papers to xml data based on pivot xml document method. In *Seventh International Conference on Document Analysis and Recognition, 2003. Proceedings.*, pages 250–255. IEEE.
- Guillaume Jaume, Hazim Kemal Ekenel, and Jean-Philippe Thiran. 2019. Funsd: A dataset for form understanding in noisy scanned documents. In *2019 International Conference on Document Analysis and Recognition Workshops (ICDARW)*, volume 2, pages 1–6. IEEE.
- Samu Kovanen and Kiyoharu Aizawa. 2015. A layered method for determining manga text bubble reading order. In *2015 IEEE International Conference on Image Processing (ICIP)*, pages 4283–4287.
- Prashant Krishnan, Zilong Wang, Yangkun Wang, and Jingbo Shang. 2024. Towards few-shot entity recognition in document images: A graph neural network approach robust to image manipulation. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 16514–16526.
- Chen-Yu Lee, Chun-Liang Li, Timothy Dozat, Vincent Perot, Guolong Su, Nan Hua, Joshua Ainslie, Ren-shen Wang, Yasuhisa Fujii, and Tomas Pfister. 2022. FormNet: Structural encoding beyond sequential modeling in form document information extraction. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3735–3754, Dublin, Ireland. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Liangcheng Li, Feiyu Gao, Jiajun Bu, Yongpan Wang, Zhi Yu, and Qi Zheng. 2020. An end-to-end ocr text re-organization sequence learning for rich-text detail image comprehension. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXV 16*, pages 85–100. Springer.
- Qiwei Li, Zuchao Li, Xiantao Cai, Bo Du, and Hai Zhao. 2023. Enhancing visually-rich document understanding via layout structure modeling. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 4513–4523.
- Qiwei Li, Zuchao Li, Ping Wang, Haojun Ai, and Hai Zhao. 2024. Hypergraph based understanding for document semantic entity recognition. *arXiv preprint arXiv:2407.06904*.
- 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. *arXiv preprint arXiv:2205.02411*.
- 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 *Proceedings of the 29th ACM International Conference on Multimedia*, pages 1912–1920.
- Haofu Liao, Aruni RoyChowdhury, Weijian Li, Ankan Bansal, Yuting Zhang, Zhuowen Tu, Ravi Kumar Satzoda, R. Manmatha, and Vijay Mahadevan. 2023. Doctr: Document transformer for structured information extraction in documents. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 19584–19594.
- Tianyu Liu, Afra Amini, Mrinmaya Sachan, and Ryan Cotterell. 2023. Linear-time modeling of linguistic structure: An order-theoretic perspective. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 808–830, Singapore. Association for Computational Linguistics.
- Xiaojing Liu, Feiyu Gao, Qiong Zhang, and Huasha Zhao. 2019. Graph convolution for multimodal information extraction from visually rich documents. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Industry Papers)*, pages 32–39, Minneapolis, Minnesota. Association for Computational Linguistics.
- Chuwei Luo, Changxu Cheng, Qi Zheng, and Cong Yao. 2023. Geolayoutlm: Geometric pre-training for visual information extraction. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Chuwei Luo, Yufan Shen, Zhaoqing Zhu, Qi Zheng, Zhi Yu, and Cong Yao. 2024. Layoutllm: Layout instruction tuning with large language models for document understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15630–15640.
- Chuwei Luo, Guozhi Tang, Qi Zheng, Cong Yao, Lianwen Jin, Chenliang Li, Yang Xue, and Luo Si. 2022a. Bi-vldoc: Bidirectional vision-language modeling for visually-rich document understanding. *arXiv preprint arXiv:2206.13155*.
- Chuwei Luo, Yongpan Wang, Qi Zheng, Liangchen Li, Feiyu Gao, and Shiyu Zhang. 2020. Merge and recognize: a geometry and 2d context aware graph model for named entity recognition from visual documents. In *Proceedings of the Graph-based Methods for Natural Language Processing (TextGraphs)*, pages 24–34.
- Siwen Luo, Yihao Ding, Siqu Long, Josiah Poon, and Soyeon Caren Han. 2022b. Doc-gcn: Heterogeneous graph convolutional networks for document layout analysis. *arXiv preprint arXiv:2208.10970*.

- Donato Malerba and Michelangelo Ceci. 2007. Learning to order: A relational approach. In *International Workshop on Mining Complex Data*, pages 209–223. Springer.
- Donato Malerba, Michelangelo Ceci, and Margherita Berardi. 2008. Machine learning for reading order detection in document image understanding. In *Machine learning in document analysis and recognition*, pages 45–69. Springer.
- Minesh Mathew, Viraj Bagal, Rubèn Tito, Dimosthenis Karatzas, Ernest Valveny, and CV Jawahar. 2022. Infographicvqa. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1697–1706.
- Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. 2021. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2200–2209.
- J-L Meunier. 2005. Optimized xy-cut for determining a page reading order. In *Eighth International Conference on Document Analysis and Recognition (ICDAR'05)*, pages 347–351. IEEE.
- Fnu Mohbat, Mohammed J Zaki, Catherine Finegan-Dollak, and Ashish Verma. 2023. **GVdoc - graph-based visual DOcument classification**. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 5342–5357, Toronto, Canada. Association for Computational Linguistics.
- G NAGY. 1984. Hierarchical representation of optically scanned documents. In *Proc. 7th Int. Conf. Patt. Recogn.*, pages 347–349.
- Laura Nguyen, Thomas Scialom, Jacopo Staiano, and Benjamin Piwowarski. 2021. Skim-attention: Learning to focus via document layout. *arXiv preprint arXiv:2109.01078*.
- Seunghyun Park, Seung Shin, Bado Lee, Junyeop Lee, Jaehung Surh, Minjoon Seo, and Hwalsuk Lee. 2019. Cord: a consolidated receipt dataset for post-ocr parsing. In *Workshop on Document Intelligence at NeurIPS 2019*.
- Qiming Peng, Yinxu Pan, Wenjin Wang, Bin Luo, Zhenyu Zhang, Zhengjie Huang, Teng Hu, Weichong Yin, Yongfeng Chen, Yin Zhang, et al. 2022. Ernie-layout: Layout knowledge enhanced pre-training for visually-rich document understanding. *arXiv preprint arXiv:2210.06155*.
- Hai Pham, Guoxin Wang, Yijuan Lu, Dinei Florencio, and Cha Zhang. 2022. Understanding long documents with different position-aware attentions. *arXiv preprint arXiv:2208.08201*.
- Rafał Powalski, Łukasz Borchmann, Dawid Jurkiewicz, Tomasz Dwojak, Michał Pietruszka, and Gabriela Pałka. 2021. Going full-tilt boogie on document understanding with text-image-layout transformer. In *International Conference on Document Analysis and Recognition*, pages 732–747. Springer.
- Liang Qiao, Can Li, Zhazhan Cheng, Yunlu Xu, Yi Niu, and Xi Li. 2024. **Reading order detection in visually-rich documents with multi-modal layout-aware relation prediction**. *Pattern Recognition*, page 110314.
- Lorenzo Quirós and Enrique Vidal. 2021. **Learning to sort handwritten text lines in reading order through estimated binary order relations**. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 7661–7668.
- Lorenzo Quirós and Enrique Vidal. 2022. Reading order detection on handwritten documents. *Neural Computing and Applications*, 34(12):9593–9611.
- Pritika Ramu, Sijia Wang, Lalla Mouatadid, Joy Rimchala, and Lifu Huang. 2024. **Re2: Region-aware relation extraction from visually rich documents**. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8723–8739.
- Sarang Shrivastava, Afreen Shaikh, Shivani Shrivastava, Chung Ming Ho, Pradeep Reddy, and Vijay Saraswat. 2021. **Handling tree-structured text: parsing directory pages**.
- Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. 2019. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8317–8326.
- Huang Siyuan, Yongping Xiong, and Wu Guibin. 2024. **LayoutPointer: A spatial-context adaptive pointer network for visual information extraction**. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3737–3748, Mexico City, Mexico. Association for Computational Linguistics.
- Jianlin Su, Ahmed Murtadha, Shengfeng Pan, Jing Hou, Jun Sun, Wanwei Huang, Bo Wen, and Yunfeng Liu. 2022. Global pointer: Novel efficient span-based approach for named entity recognition. *arXiv preprint arXiv:2208.03054*.
- Ryota Tanaka, Taichi Iki, Kyosuke Nishida, Kuniko Saito, and Jun Suzuki. 2024. **Instructdoc: A dataset for zero-shot generalization of visual document understanding with instructions**. *arXiv preprint arXiv:2401.13313*.
- 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. *arXiv preprint arXiv:2106.12940*.

- Zineng Tang, Ziyi Yang, Guoxin Wang, Yuwei Fang, Yang Liu, Chenguang Zhu, Michael Zeng, Cha Zhang, and Mohit Bansal. 2023. Unifying vision, text, and layout for universal document processing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 19254–19264.
- Song Tao, Zijian Wang, Tiantian Fan, Canjie Luo, and Can Huang. 2022. Knowing where and what: Unified word block pretraining for document understanding. *arXiv preprint arXiv:2207.13979v1*.
- Yi Tu, Ya Guo, Huan Chen, and Jinyang Tang. 2023. Layoutmask: Enhance text-layout interaction in multi-modal pre-training for document understanding. *arXiv preprint arXiv:2305.18721*.
- Yi Tu, Chong Zhang, Ya Guo, Huan Chen, Jinyang Tang, Huijia Zhu, and Qi Zhang. 2024. Uner: A unified prediction head for named entity recognition in visually-rich documents. *arXiv preprint arXiv:2408.01038*.
- Ivan Vendrov, Ryan Kiros, Sanja Fidler, and Raquel Urtasun. 2015. Order-embeddings of images and language. *arXiv preprint arXiv:1511.06361*.
- Dongsheng Wang, Zhiqiang Ma, Armineh Nourbakhsh, Kang Gu, and Sameena Shah. 2023a. Docgraphlm: Documental graph language model for information extraction. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1944–1948.
- Dongsheng Wang, Natraj Raman, Mathieu Sibue, Zhiqiang Ma, Petr Babkin, Simerjot Kaur, Yulong Pei, Armineh Nourbakhsh, and Xiaomo Liu. 2023b. Docllm: A layout-aware generative language model for multimodal document understanding. *arXiv preprint arXiv:2401.00908*.
- Guoxin Wang, Yijuan Lu, Lei Cui, Tengchao Lv, Dinei Florencio, and Cha Zhang. 2022a. A simple yet effective learnable positional encoding method for improving document transformer model. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2022*, pages 453–463.
- Hao Wang, Qingxuan Wang, Yue Li, Changqing Wang, Chenhui Chu, and Rui Wang. 2023c. Doctrack: A visually-rich document dataset really aligned with human eye movement for machine reading. *arXiv preprint arXiv:2310.14802*.
- Jiapeng Wang, Lianwen Jin, and Kai Ding. 2022b. 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 (Volume 1: Long Papers)*, pages 7747–7757.
- Renshen Wang, Yasuhisa Fujii, and Alessandro Bisacco. 2023d. Text reading order in uncontrolled conditions by sparse graph segmentation. *arXiv preprint arXiv:2305.02577*.
- Wenjin Wang, Zhengjie Huang, Bin Luo, Qianglong Chen, Qiming Peng, Yinxu Pan, Weichong Yin, Shikun Feng, Yu Sun, Dianhai Yu, et al. 2022c. mmlayout: Multi-grained multimodal transformer for document understanding. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4877–4886.
- 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.
- 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.
- Zilong Wang, Yichao Zhou, Wei Wei, Chen-Yu Lee, and Sandeep Tata. 2023e. [Vrdu: A benchmark for visually-rich document understanding](#). In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '23*, page 5184–5193, New York, NY, USA. Association for Computing Machinery.
- Patricia Wright. 1999. The psychology of layout: Consequences of the visual structure of documents. *American Association for Artificial Intelligence Technical Report FS-99-04*, pages 1–9.
- Yang Xu, Yiheng Xu, Tengchao Lv, Lei Cui, Furu Wei, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, Wanxiang Che, et al. 2021a. Layoutlmv2: Multi-modal pre-training for visually-rich document understanding. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 2579–2591.
- Yiheng Xu, Tengchao Lv, Lei Cui, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, and Furu Wei. 2021b. Layoutxlm: Multimodal pre-training for multilingual visually-rich document understanding. *arXiv preprint arXiv:2104.08836*.
- Zhibo Yang, Rujiao Long, Pengfei Wang, Sibao Song, Humen Zhong, Wenqing Cheng, Xiang Bai, and Cong Yao. 2023. Modeling entities as semantic points for visual information extraction in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15358–15367.
- 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 *The Eleventh International Conference on Learning Representations*.

Chong Zhang, Ya Guo, Yi Tu, Huan Chen, Jinyang Tang, Huijia Zhu, Qi Zhang, and Tao Gui. 2023a. Reading order matters: Information extraction from visually-rich documents by token path prediction. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13716–13730.

Chong Zhang, Yixi Zhao, Chenshu Yuan, Yi Tu, Ya Guo, and Qi Zhang. 2024. Rethinking the evaluation of pre-trained text-and-layout models from an entity-centric perspective. *arXiv preprint arXiv:2402.02379*.

Yue Zhang, Zhang Bo, Rui Wang, Junjie Cao, Chen Li, and Zuyi Bao. 2021. Entity relation extraction as dependency parsing in visually rich documents. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2759–2768.

Zhenrong Zhang, Jiefeng Ma, Jun Du, Licheng Wang, and Jianshu Zhang. 2022. [Multimodal pre-training based on graph attention network for document understanding](#). *Trans. Multi.*, 25:6743–6755.

Zhiyang Zhang, Yaping Zhang, Lu Xiang, Yang Zhao, Yu Zhou, and Chengqing Zong. 2023b. A novel dataset and benchmark analysis on document image translation. In *Machine Translation*, pages 103–115, Singapore. Springer Nature Singapore.

Jun Zhao, Wayne Xin Zhao, WenYu Zhan, Tao Gui, Qi Zhang, Liang Qiao, Zhanzhan Cheng, and Shiliang Pu. 2022. Read extensively, focus smartly: A cross-document semantic enhancement method for visual documents ner. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 2034–2043.

Xi Zhu, Xue Han, Shuyuan Peng, Shuo Lei, Chao Deng, and Junlan Feng. 2023. Beyond layout embedding: Layout attention with gaussian biases for structured document understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7773–7784.

A Theoretical Foundations of Directed Acyclic Relation and Order Relations

To formalize the problem we intend to solve, necessary concepts and fundamental theorems from order theory are introduced in below, where A denotes an arbitrary set:

Definition 1 (Directed Acyclic Relation).

A binary relation $DAR \in A \times A$ is a directed acyclic relation over A iff. it is:

1. Acyclic: $\forall k \geq 1 : \bigcap_{1 \leq i \leq k} \{s_i\} \subseteq A$ s.t. $\{(s_i, s_{i+1})\}_{1 \leq i \leq k} \subseteq DAR$ where $s_{k+1} = s_1$.

Comment 1. Each directed acyclic relation R over A corresponds to a directed acyclic graph $G = (A, E)$, where the nodes are elements in A , and each edge $(s_1, s_2) \in E$ is a linked node pairs in R such that $s_1 R s_2$.

Definition 2 (Partial Order, Strict Partial Order).

A binary relation $PO \in A \times A$ is a partial order over A iff. it is:

1. Reflexive: $\forall s \in A : (s, s) \in PO$.
2. Antisymmetric: $\forall s_1, s_2 \in A : (s_1, s_2) \in PO \wedge (s_2, s_1) \in PO \Rightarrow s_1 = s_2$.
3. Transitive: $\forall s_1, s_2, s_3 \in A : (s_1, s_2) \in PO \wedge (s_2, s_3) \in PO \Rightarrow (s_1, s_3) \in PO$.

A binary relation $SPO \in A \times A$ is a strict partial order over A iff. it is antisymmetric, transitive and:

1. Irreflexive: $\forall s \in A : (s, s) \notin SPO$.

Definition 3 (Transitive Closure).

The transitive closure of a binary relation R over A is the smallest (w.r.t. \subseteq) transitive relation R^+ over A such that $R \subseteq R^+$.

Theorem 1. For any relation R over A , the transitive closure of R always uniquely exists.

Proof of Theorem 1. The proof of existence and uniqueness are required of the transitive closure of an arbitrary binary relation R over A .

1. Proof of existence: Note that the intersection of any family of transitive relations is again transitive. For an arbitrary binary relation R over A , there exists at least one transitive relation containing R , namely the trivial one: $A \times A$. The transitive closure of R is then given by the intersection of all transitive relations containing R as $R^+ = \bigcap_{R' \subseteq A \times A, R \subseteq R'} R'$.

It is trivial to see that R^+ is transitive and is the smallest among all transitive relations containing R .

2. Proof of uniqueness: Given that there exist R^1 and R^2 and both of them are the transitive closures of R , note that R^2 is a transitive

relation over A , and therefore $R^1 \subseteq R^2$. Similarly $R^2 \subseteq R^1$, therefore $R^1 = R^2$.

Theorem 2. The transitive closure of a DAR is a SPO.

The proof of Theorem 2 is based on Lemma 1.

Lemma 1. For the transitive closure R^+ of an arbitrary DAR R over set A , there do not exist $s_1, s_2 \in A$ such that $(s_1, s_2) \in R^+$ and $(s_2, s_1) \in R^+$.

Proof of Lemma 1. Assume that $(s_1, s_2) \in R^+$ and $(s_2, s_1) \in R^+$. By the definition of the transitive closure, there exist sequences of elements $(a_1, a_2, \dots, a_n), \{a_i\}_{1 \leq i \leq n} \subseteq A$ and $(b_1, b_2, \dots, b_m), \{b_j\}_{1 \leq j \leq m} \subseteq A$ such that: $(a_i, a_{i+1}) \in R$ for all $1 \leq i < n$, where $s_1 = a_1, a_n = s_2$; $(b_j, b_{j+1}) \in R$ for all $1 \leq j < m$, where $s_2 = b_1, b_m = s_1$.

However, $(a_1, a_2, \dots, a_n = b_1, b_2, \dots, b_m = a_1)$ comprises a cycle within R , which is conflict with the definition of directed acyclic relation. Therefore, the assumption does not hold, proving that there do not exist $s_1, s_2 \in A$ such that $(s_1, s_2) \in R^+$ and $(s_2, s_1) \in R^+$.

Proof of Theorem 2. The transitive closure R^+ of an arbitrary DAR R over set A is:

1. Irreflexive: When it takes $s = s_1 = s_2$ in Lemma 1, $\forall s \in A, (s, s) \notin R^+$.
2. Antisymmetric: Due to Lemma 1, the premise such that $(s_1, s_2) \in R^+ \wedge (s_2, s_1) \in R^+$ is always false for all $s_1, s_2 \in A$, so this property is always valid.
3. Transitive: Since R^+ is a transitive closure, this property is always valid.

Therefore, R^+ is a SPO over A by definition.

Definition 4 (Total Order, Strict Total Order).

A partial order TO over A is a total order iff. $\forall s_1, s_2 \in A : (s_1, s_2) \in TO \vee (s_2, s_1) \in TO$, i.e. every two elements of A are comparable.

A strict partial order STO over A is a strict total order iff. $\forall s_1, s_2 \in A \wedge s_1 \neq s_2 : (s_1, s_2) \in STO \vee (s_2, s_1) \in STO$, i.e. every two different elements of A are comparable.

Comment 2. Each TO/STO over $A = \{s_i\}_{1 \leq i \leq |A|}$ can be represented as a permutation of A as $(s_{P(1)}, \dots, s_{P(|A|)})$ where P is a bijection from $\{i\}_{1 \leq i \leq |A|}$ to itself.

B A Brief Overview of Layout Reading Order Prediction

Generally, the term layout reading order refers to the order in which people will read layout contents.

Modeling the order itself plays an important role to visually-rich document understanding, while it is also beneficial to incorporate the reading order information in addressing various VrD tasks. Also, the cognition of this problem is advancing along with the development of this field. This section introduces the progressive development of modeling of layout reading order, together with its application in document intelligence.

Reading Order as a Permutation The assumption that layout reading order could be represented as a permutation on layout elements can be traced back to (NAGY, 1984), which proposed XY-cut tree to represent the hierarchical structure of layout elements. A permutation of layout elements can be derived from the XY-cut tree as the layout reading order. Rule-based methods (Ha et al., 1995; Ishitani, 2003; Meunier, 2005) and statistical-ML-based methods (Han et al., 2007; Kovanen and Aizawa, 2015) have been proposed to predict the optimal XY-cut tree of layouts. However, the layout reading order is not always from-top-to-bottom and from-left-to-right, which constrains the maximum effectiveness of XY-cut tree in representing the reading order. Therefore, Malerba and Ceci (2007); Malerba et al. (2008); Ferilli et al. (2014) have proposed to directly predict a permutation of layout elements, rather than a XY-cut tree.

Coming to the deep learning era, Li et al. (2020) proposed a pointer network with a GCN layout encoder for reading order prediction. Following this paradigm, ORO (Quirós and Vidal, 2021, 2022), DocReL (Li et al., 2022), TPP (Zhang et al., 2023a) and MLARP (Qiao et al., 2024) tackled the task via devising relation extraction models between layout elements and decoding strategies from predicted relation pairs. Besides, LayoutReader (Wang et al., 2021) predicted a word sequence as the reading order sequence without guaranteeing its property as a permutation of input words. SGS (Wang et al., 2023d) introduced two major patterns of reading order as column-wise and row-wise order, and tackled the task by segmenting and sorting ordered blocks.

In leveraging reading order information in downstream application, ERNIE-Layout (Peng et al., 2022) includes a pre-training objective of reading order prediction; TPP (Zhang et al., 2023a) verifies the effectiveness of a reading order prediction model to correct the input token sequence of layout-

aware models. Relevant data resources that are publicly available include ReadingBank (Wang et al., 2021) and DocTrack (Wang et al., 2023c) on scanned VrDs, and OHG, FCR and ABP (Quirós and Vidal, 2022) on handwritten VrDs. However, all the above works failed to capture or leverage the complete reading order information due to the inappropriate modeling of reading order.

Reading Order beyond a Permutation Several previous works have noted that layout reading order information cannot be fully conveyed by one single permutation of layout elements, i.e. a globally total-ordered element sequence. All these works treated the incomplete reading order information issue as a task-irrelevant inductive bias injecting noise into the inputs, and attempted to mitigate its negative effects. In transformer-based layout-aware models, the noise brought by noisy (*incomplete*) reading order information is usually introduced by absolute position embeddings. LSPE (Wang et al., 2022a) introduced a learnable fully-connected feed-forward sinusoidal positional embedding network to handle the noise. XYLayoutLM (Gu et al., 2022) reduced the random noise by utilizing multiple globally total-ordered element sequences generated by an augmented XY-cut algorithm. UTel (Tao et al., 2022) introduced 1D clipped relative position embeddings to mitigate the use of long-term reading order information which is much more noisy. LayoutMask (Tu et al., 2023) utilized locally total-ordered word sequences of segments with local 1D position embeddings. These works increased the robustness of models on the noisy reading order information, but did not make use of the non-linear reading order information which is essential in complex layouts. From theoretical perspective, Ceci et al. (2007) modeled the layout reading order as multiple disjoint chains (i.e. total-ordered subsets) of the set of layout elements. However, this assumption did not consider the numerous intertwined reading order relationships in irregular layouts, such as design layouts (Fig. 3a) and tables (Fig. 3c).

Reading Order as Ordering Relations In this work, we model the reading order as ordering relations over layout elements. The presented modeling draws inspiration from previous works which modeled other non-linear structures of VrD layouts as binary relations or ordered structures. Aiello and Smeulders (2003, 2004) illustrated to represent the reasoning dependencies about

temporal intervals between layout elements as a bi-dimensional Allen relation (Allen, 1983). Clausner et al. (2013); Shrivastava et al. (2021) introduced tree structures of ordered and unordered groups to represent the logical structure of layout elements using a tree of ordered and unordered groups.

The presented modeling is also motivated by order-theoretic modeling of linguistic objects. Vendrov et al. (2015) proposed to model the visual-semantic hierarchy as a partial order over images and language, as it can be displayed as a DAG. Athiwaratkun and Wilson (2018) introduced density order embeddings to embed the hierarchy relation into vector space. Liu et al. (2023) presented an order-theoretic framework to model the relation between token pairs, which universally tackled a series of structured prediction tasks. These works highlights the motivation of predicting and leveraging the inherent structures within textual data by modeling them as ordering structures.

C Baseline Selection for Semantic Entity Recognition and Entity Linking

Semantic entity recognition and entity linking are the most typical and important tasks in VrD-IE. On tackling these tasks, task-specific methods are developed together with universal methods suited to multiple tasks. Also, the performance on these tasks are adopted as an evaluation of the performance of pre-trained text-and-layout models. In order to conduct a comprehensive comparison between our methods and other diverse methods, we roughly categorize the existing methods, selecting the most representative methods from each category for comparison. In principle, we only consider layout-based fully-supervised methods, and mainly focus on the most popular methods from the collection of top-tier venues (conferences and journals) in AI applications:

- **Attention:** These methods attempt to better capture task-relevant features by modulating the original attention mechanism. Especially, entity linking methods propose relation-aware attention modules to model the linking relationship between layout elements. We select representative text-and-layout representation learning methods (Wang et al., 2022b; Hong et al., 2022; Huang et al., 2022; Zhu et al., 2023) and entity linking methods (Hwang et al., 2021; Zhang et al., 2021; Liao et al., 2023) for comparison.
- **Graph:** These methods tackle the task by

modeling graph structures, which is especially useful in modeling the linking relationships between entities in EL task. In representative methods, GNN-based structures are adopted for either boosting relation-aware pre-training (Appalaraju et al., 2021; Lee et al., 2022; Wang et al., 2022c; Zhang et al., 2022; Wang et al., 2023a; Li et al., 2023) or entity linking modeling (Tang et al., 2021; Gemelli et al., 2022; Ramu et al., 2024).

- **Pre-trained:** These methods take advantage of unsupervised pre-training on large-scale data resources. They achieve the goal only by designing pre-training objectives, rather than introducing specific architectures. The selected methods are all pre-trained text-and-layout models aiming to serve as a representation model for document layouts (Li et al., 2021, 2022; Peng et al., 2022; Luo et al., 2022a; Tu et al., 2023; Bai et al., 2023; Luo et al., 2023).
- **Supervised pre-trained:** These methods leverage extra supervised data of the downstream tasks, therefore their performance are not fairly comparable to our methods (Yang et al., 2023; Tu et al., 2024). Nevertheless, considering the similar research aim and model paradigm, the results of these methods are also reported in the main table.
- **LLM-based:** These methods are based on large-scale generative pre-training or supervised fine-tuning on domain-specific data resources. Considering their text-generation paradigm and disproportionate parameter scale, the application scenario of these models should be totally different from the extraction-based methods introduced above. Therefore, we separately listed their performance in Tab. 5 to compare with our methods.
- **Span-extraction-based:** These methods are also worth mentioned as they have also achieved good performance on VrD-IE tasks (Zhang et al., 2023a; Li et al., 2024). However, due to the difference on the evaluation metrics, these methods are not involved in the comparison.

D Implementation Details

Reading Order Prediction Methods For training the baseline sequence-based ROP methods (LR and TPP) on ROOR, these models are supervised by the topology-sorted sequence of the original label for each sample. The training process

Method	#Params	FUNSD	CORD	SROIE
UDOP (2023)	794M	91.62	97.58	-
DocLLM (2023b)	7B	51.8	67.4	91.9
LayoutLLM (2024)	7B	79.98	63.10	72.12
RORE-GeoLayoutLM	~399M	91.84	98.52	96.97

Table 5: Performance of LLM-based generative methods and our methods on three VrD-SER benchmarks.

generally follows the settings reported in the original papers, with the only exception that we fine-tune each model by 500 epoches. During inference, these models produce word sequences as outputs, and each two adjacent words are extracted as a predicted reading order relation pair. In segment-level evaluation of these models, we order the segments according to the first appearance of their words in the original outputs, constructing segment sequences as outputs.

The proposed relation-extraction-based ROP model adopts LayoutLMv3-base (Huang et al., 2022) and LayoutMask (Tu et al., 2023) as the backbones. The maximum sequence length of textual tokens for both of them is 2048. The maximum layout elements is set to 512 for word-level task and 256 for segment-level task, ensuring the number of layout elements within each ROOR samples should not exceeds the maximum. The head size of the global pointer decoding network (i.e. the dimension of query and key vectors in Eqn. 3) is set to 128. We use an AdamW optimizer with 5% linear warming-up steps, a 0.1 dropout rate, and a 1e-5 weight decay, together with a cosine scheduler. The learning rate is searched from {1e-5, 2e-5, 5e-5, 1e-4, 1.5e-4, 2e-4} for LayoutLMv3-base and {1e-5, 2e-5, 5e-5, 1e-4} for LayoutLMv3-large, and the best learning rate is 1e-4/2e-5, respectively. We fine-tune the model by 1,000 epoches (10,000 steps) on 1 A800 GPUs with batch size being 16 and patience being 50 for early stopping.

Reading Order Relation Enhanced Methods for Downstream VrD Tasks For information extraction tasks, the enhanced semantic entity recognition (SER) and entity linking (EL) models strictly keep the same settings of the original implementation (Zhang et al., 2024), with all experiments using the learning rate of 1e-5, batch size of 16, number of fine-tuning epoches of 500 and early stopping patience of 50. The initial value for λ^l is all set to 10 in all layers, and the learning rate of the λ s are searched between the original learning rate and 1e-2, where the

Task Level	Method	Bounding Box Level	
		Word	Segment
Word	LayoutLMv3-base	90.26	87.93
	LayoutLMv3-large	93.01	92.80
Segment	LayoutLMv3-base	63.00	68.60
	LayoutLMv3-large	76.04	82.38

Table 6: Ablation study on the level of bounding box information used in the baseline ROP model.

latter is for faster convergence. Specially, in these three groups of experiment: LayoutLMv3-base-FUNSD-SER, GeoLayoutLM-FUNSD-SER and GeoLayoutLM-FUNSD-EL, since we used the fine-tuned weights provided by the official repositories of LayoutLMv3² and GeoLayoutLM³ which are fairly close to the optimal state, we only use the reading order information in the first 4 layers, and the value for $\lambda^l, 1 \leq l \leq 4$ are set to 0.1 initially. In Tab. 2 and 3, for the baseline models, we report the bias-removed running results from (Zhang et al., 2024) together with our running results using its implementation. Among other baseline methods, the performance of GraphLayoutLM (Li et al., 2023) can only be reproduced on CORD according to its official release⁴. We are unable to reproduce its performance on other datasets since the pre-trained weights are not released. Therefore, we only report its performance on CORD from its official release. These reproduced results are marked with *.

For generative VrD-QA tasks, the baseline model is comprised of a LayoutLMv3-base encoder and a BART-base decoder (Lewis et al., 2019), and is pre-trained on open-sourced VrD-QA-form data filtered from InstructDoc (Tanaka et al., 2024) for better initialization. The model is then fine-tuned on each dataset using learning rate of 1e-5, batch size of 48, with an AdamW optimizer with 2% linear warming-up steps and a cosine scheduler for 6 epochs. For the RORE settings, the reading-order-relation-aware attention module is only applied to its decoder, and the other settings keep the same of RORE-VrD-IE experiments.

E Ablation Studies

Impact of Layout-level on ROP Models For better understanding the usage of text-and-layout

²<https://github.com/microsoft/unilm/tree/master/layoutlmv3>

³<https://github.com/AlibabaResearch/AdvancedLiterateMachinery/tree/main/DocumentUnderstanding/GeoLayoutLM>

⁴<https://github.com/Line-Kite/GraphLayoutLM>

Task	RORE		
	None	Pseudo	Ground Truth
SER	83.62	83.96 ($\uparrow 0.34$)	84.34 ($\uparrow 0.72$)
EL	86.18	86.56 ($\uparrow 0.38$)	87.42 ($\uparrow 1.24$)

Table 7: Ablation study on the impact of the quality of reading order relation supervision. We report the performance of GeoLayoutLM and its RORE variant on EC-FUNSD with pseudo or ground truth reading order relation inputs are reported. The pseudo labels are generated by the same ROP model used in Tab. 3 and 4.

information in the proposed ROP models, we conduct an ablation study by configuring the level of bounding box inputs. The results in Tab. 6 demonstrate that the proposed ROP models generally perform better when the bounding box level coincides with the task level.

Impact of the Quality of Pseudo Reading Order Relation Labels for Downstream Tasks

In the introduced RORE pipeline, pseudo reading order relation labels generated by off-the-shelf ROP models are utilized for any VrD task lacking manual reading order annotations. It is of interest to determine the extent to which performance is affected when using pseudo labels instead of ground truth labels in downstream tasks.

To quantify the adverse effects of employing pseudo labels, we conduct an ablation study based on the EC-FUNSD dataset, which has available ground truth reading order labels (from ROOR). We conduct experiments on GeoLayoutLM for VrD-IE, choosing either pseudo or ground truth labels during training and inference. According to Tab. 7, there is a noticeable performance gap when using pseudo labels. These results underscore the importance of creating corresponding reading order relation annotations when implementing the RORE pipeline in real-world VrD applications. Additionally, these results reflect the effectiveness of RORE methods in Tab. 3 and 4, as they have achieved satisfying performance even when the ground truth labels are unavailable and the target task differs from the domain of the ROP model.

Impact of Different Reading Order Relations on Application Performance

In this work, two types of reading order relationships are proposed, of which ISDR is believed to represent local adjacent relationship between layout element, and GSDR is known as a global order representation indicating the possible order relationship between each two layout elements. To determine the role

Dataset	Task	None	ISDR	GSDR
FUNSD	SER	91.10	91.84 ($\uparrow 0.74$)	91.75 ($\uparrow 0.65$)
	EL	88.06	88.46 ($\uparrow 0.40$)	88.59 ($\uparrow 0.53$)
EC-FUNSD	SER	83.62	84.34 ($\uparrow 0.72$)	84.13 ($\uparrow 0.51$)
	EL	86.18	87.42 ($\uparrow 1.24$)	86.53 ($\uparrow 0.35$)

Table 8: Ablation study on the impact of different types of reading order relation supervision. We report the performance of GeoLayoutLM on EC-FUNSD incorporating with different types of reading order information.

of each of them in enhancing downstream VrD tasks, we conduct experiments on GeoLayoutLM for VrD-IE by separately incorporating the ISDR and GSDR information into the model.

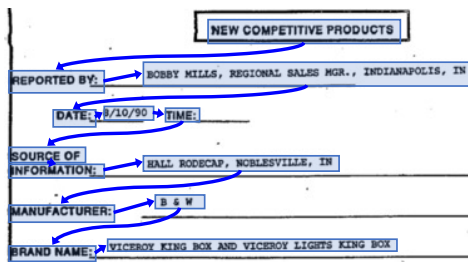
Tab. 8 illustrates that both types of reading order information substantially benefits the downstream tasks. Since the method with ISDR behaves better in the majority of task settings, we choose to leverage the ISDR information in our main experiments.

F Case Studies

Strengths and Weaknesses of the Proposed ROP Models For better understanding the strengths and weaknesses of the proposed ROP models, we conduct a case study for analyzing the prediction results by the LayoutLMv3-large model, both on in-domain datasets (e.g. FUNSD as scanned documents) and out-of-domain datasets (e.g. CORD as receipts, InfoVQA as posters). It is expected that the proposed model is able to recognize the reading pattern characteristic of in-domain samples, while also exhibiting some degree of generalization ability to out-of-domain samples. Fig. 5 visualizes the predict result of interested cases, revealing that: (1) **The proposed model has been acquainted with the major reading patterns of in-domain samples.** According to Fig. 5 (a-b), the proposed model links spatially-separated text regions which constitute a single content, while it also distinguishes irrelevant elements from the main content, such as page header and footer, seals, and watermarks. Besides, the proposed model performs well in identifying common non-linear reading order occasions in layouts, including table structures, as shown in Fig. 5 (d, f-g). It is also revealed that the proposed model does not exhibit a high false positive rate on table structures, unlike the other relation extraction models. We attribute the good phenomenon to the high annotation quality of the training data. (2)

The proposed model tends to over-rely on layout features, yet paying less attention to textual features. Fig. 5 (c, e) depict that the model shows the tendency to predict a linking from an element to the nearest element on its right, disregarding their semantic correlation. In Fig. 5 (h), confronted with a bar chart which is unseen during training, the model applies the pattern of processing tables, introducing false linking between horizontally-aligned elements. In Fig. 5 (g), a single table is falsely recognized as two tables. The lower part of the table is recognized as a separate table since it has larger fonts and its bounding boxes are misaligned with the upper part. These sub-optimal outcomes suggest potential improvements of the proposed model to leverage the textual and vision features in a more efficient way.

Positive Impact of RORE for Downstream VrD Tasks To determine the specific positive impact of RORE for downstream tasks, we conduct a case study by visualizing and comparing the prediction results of baseline models and their corresponding RORE variants for each task. According to Fig. 6, it reveals that the reading order relation inputs benefit the task by serving as a strong and reliable cue in the task. In EL and QA, the relationship between key/question entity and value/answer entity is highly overlapped with the successive relationship during reading, thus the latter information would be fully leveraged during training and serves as a key determinant in prediction. In SER, the reading order inputs hint the successive relationship between the words before and after the line break, which benefits the recognition of cross-row entities since RORE models would rely more on the reading order information than baseline models. These results also explain the success of utilizing pseudo reading order information in training RORE models, as moderate noise during training makes RORE models to focus on the more general reading order patterns that are also more robust in prediction, thus lead to steady performance improvement.



(a) In-domain (Scanned document)




(b) Out-of-domain (Receipt)



(c) Out-of-domain (Poster)



(d) Out-of-domain (Poster with non-linear order)



(e) Out-of-domain (Poster with WordArt)


REQUISITION NO.	ISSUED BY
ACCT. NO.	DEPT. NO.
1111	B700
	APPROVED BY
	H. J. Momeneppe

(f) In-domain (Table in scanned document)

V.B. BAL	15000
V. BASU PRON	21500
TOTAL	31000
CASH	31000

(g) Out-of-domain (Table in receipt, without grid-lines)

Underlying cash balance



Source: Budget 2019-20; MYEFO 2018-19
 The Conversation

(h) Out-of-domain (Chart)




Figure 5: Case study of the proposed reading order prediction model. Each arrow represents a predicted relation linking between segments.

Case	Baseline Models	Reading-order-relation-enhanced Models
Cross-row Entity Recognition	Acceptance/Response: What is the retailers response to Lorillard's Excel Merchandising plan?	Acceptance/Response: What is the retailers response to Lorillard's Excel Merchandising plan?
	Hardware Evaluation/Effectiveness: Comment on the assembly of displays and application of shields:	Hardware Evaluation/Effectiveness: Comment on the assembly of displays and application of shields: TP FP
Cross Linking (False Positive)		
Complex Linking (False Negative)		
Visual Document Informative QA		<p>Question: What is the budget offered for 3 Posters (including production)?</p> <p>Baseline Model: \$15,000 X</p> <p>RORE Model: \$27,000 ✓</p>

Figure 6: Case study of baseline models and their corresponding reading-order-relation-enhanced variants for VrD-SER, VrD-EL and VrD-QA. Entities are marked with shades and are distinguished by the shade of color. Entity linking are marked with arrows.