

LayoutMask: Enhance Text-Layout Interaction in Multi-modal Pre-training for Document Understanding

Yi Tu, Ya Guo, Huan Chen, Jinyang Tang

Ant Group, China

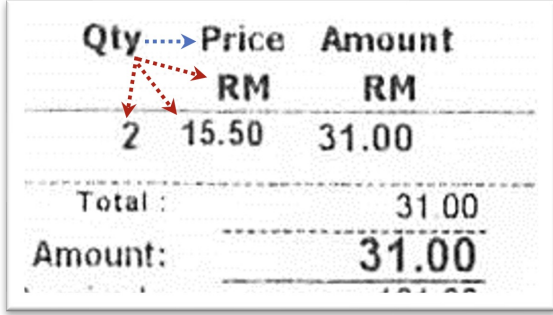
{qianyi.ty, guoya.gy, chenhuan.chen, jinyang.tjy}@antgroup.com

Abstract

Visually-rich Document Understanding (VrDU) has attracted much research attention over the past years. Pre-trained models on a large number of document images with transformer-based backbones have led to significant performance gains in this field. The major challenge is how to fusion the different modalities (text, layout, and image) of the documents in a unified model with different pre-training tasks. This paper focuses on improving text-layout interactions and proposes a novel multi-modal pre-training model, LayoutMask. LayoutMask uses local 1D position, instead of global 1D position, as layout input and has two pre-training objectives: (1) Masked Language Modeling: predicting masked tokens with two novel masking strategies; (2) Masked Position Modeling: predicting masked 2D positions to improve layout representation learning. LayoutMask can enhance the interactions between text and layout modalities in a unified model and produce adaptive and robust multi-modal representations for downstream tasks. Experimental results show that our proposed method can achieve state-of-the-art results on a wide variety of VrDU problems, including form understanding, receipt understanding, and document image classification.

1 Introduction

Visually-rich Document Understanding (VrDU) is an important research area that aims to understand various types of documents (*e.g.*, forms, receipts, and posters), and it has attracted much attention from both academia and industry. In recent years, pre-training techniques (Devlin et al., 2019; Zhang et al., 2019) have been introduced into this area and self-supervised pre-training multi-modal models have demonstrated great successes in various VrDU tasks (Xu et al., 2020, 2021; Hong et al., 2022; Li et al., 2021a).



Segments	Qty	Price	Amount	RM	...		
Tokens	_Q	ty	_Price	_A	mount	_RM	...
Global 1D	1	2	3	4	5	6	...
Local 1D	1	2	1	2	3	1	...

Figure 1: A receipt image from SROIE dataset and the global/local 1D positions of tokens based on global/in-segment reading orders. Local 1D positions restart with “1” for each individual segment. **Blue Arrow:** When using global 1D position, the reading order is explicitly implied by the ascending numbers, so the word after “Qty” is “Price”. **Red Arrows:** When using local 1D position, the successor of “Qty” is not directly given and can have more possible choices, so their semantic relations and 2D positions will be considered during pre-training.

However, existing document pre-training models suffer from reading order issues. Following the idea of BERT (Devlin et al., 2019), these methods (Xu et al., 2020, 2021; Hong et al., 2022) usually adopt ascending numbers (*e.g.*, 0, 1, 2, ..., 511) to represent the global reading order of tokens in the document. Then, these numbers are encoded into 1D position embeddings to provide explicit reading order supervision during pre-training, which are called “global 1D position”. While such global 1D positions are widely used in NLP models for textual data, it is not a good choice for document data. Firstly, plain texts always have definite and linear reading orders, but the reading order of a document may not be unique or even linear, which

cannot be simply encoded with monotonically increasing numbers. Secondly, the global reading order of a document is usually obtained by ordering detected text segments from OCR tools with empirical rules, so it heavily relies on stable and consistent OCR results, affecting the generalization ability in real-world applications. Moreover, the empirical rules to obtain reading orders (*e.g.*, “top-down and left-right”) may not be able to handle documents with complex layouts, thus providing inaccurate supervision.

Some previous studies have attempted to solve the above reading order issues. LayoutReader (Wang et al., 2021) proposes a sequence-to-sequence framework for reading order detection with supervised reading order annotations. XYLayoutLM (Gu et al., 2022) utilizes an augmented XY Cut algorithm to generate different proper reading orders during pre-training to increase generalization ability. ERNIE-Layout (Peng et al., 2022) rearranges the order of input tokens in serialization modules and adopts a reading order prediction task in pre-training. While these studies propose data-based or rule-based solutions to provide explicit reading order supervision, we believe that the self-supervised pre-training process on a large number of documents without using extra supervision is sufficient to help the model to learn reading order knowledge, and such knowledge can be implicitly encoded into the pre-trained model with better adaptiveness and robustness to various document layouts.

We proposed a novel multi-modal pre-training model, **LayoutMask**, to achieve this goal. LayoutMask only uses text and layout information as model input and aims to enhance text-layout interactions and layout representation learning during pre-training. It differs from previous studies in three aspects: choice of 1D position, masking strategy, and pre-training objective.

Instead of global 1D position, LayoutMask proposes to use the in-segment token orders as 1D position, which is referred to as “**local 1D position**” (See illustration in Figure 1). As local 1D position does not provide cross-segment orders, LayoutMask is supposed to infer global reading order by jointly using 1D position, 2D position, and semantic information, thus bringing in-depth text-layout interactions. To further promote such interactions, we equip the commonly used pre-training objective, Masked Language Modeling (MLM), with

two novel masking strategies, **Whole Word Masking** and **Layout-Aware Masking**, and design an auxiliary pre-training objective, **Masked Position Modeling**, to predict masked 2D positions during pre-training. With the above designs, we increase the difficulty of pre-training objectives and force the model to focus more on layout information to obtain reading order clues in various document layouts in self-supervised learning, thus producing more adaptive and robust text-layout representations for document understanding tasks.

Experimental results show that our proposed method can bring significant improvements to VrDU tasks and achieve SOTA performance with only text and layout modalities, indicating that previous studies have not fully explored the potential power of layout information and text-layout interactions. The contributions of this paper are summarized as follows:

1. We propose LayoutMask, a novel multi-modal pre-training model focusing on text-layout modality, to generate adaptive and robust multi-modal representations for VrDU tasks.
2. In LayoutMask, we use local 1D position instead of global 1D position to promote reading order learning. We leverage Whole Word Masking and Layout-Aware Masking in the MLM task and design a new pre-training objective, Masked Position Modeling, to enhance text-layout interactions.
3. Our method can produce useful multi-modal representations for documents and significantly outperforms many SOTA methods in multiple VrDU tasks.

2 Related Work

The early studies in VrDU area usually use unimodal models or multi-modal models with shallow fusion (Yang et al., 2016, 2017; Katti et al., 2018; Sarkhel and Nandi, 2019). In recent years, pre-training techniques in NLP (Devlin et al., 2019; Zhang et al., 2019; Bao et al., 2020) and CV (Bao et al., 2021; Li et al., 2022) have become more and more popular, and they have been introduced into this area. Inspired by BERT (Devlin et al., 2019), LayoutLM (Xu et al., 2020) first improved the masked language modeling task by using the 2D coordinates of each token as layout embeddings, which can jointly model interactions be-

tween text and layout information and benefits document understanding tasks. Following this idea, LayoutLMv2 (Xu et al., 2021) propose to concatenate image patches with textual tokens to enhance text-image interactions, and LayoutLMv3 (Huang et al., 2022) proposed to learn cross-modal alignment with unified text and image masking.

While the above methods focus on text-image interactions, some other studies have realized the importance of layout information. StructuralLM (Li et al., 2021a) utilizes segment-level layout features to provide word-segment relations. DocFormer (Appalaraju et al., 2021) combines text, vision, and spatial features with a novel multi-modal self-attention layer and shares learned spatial embeddings across modalities. LiLT (Wang et al., 2022) proposes a language-independent layout transformer where the text and layout information are separately embedded. ERNIE-Layout (Peng et al., 2022) adopts a reading order prediction task in pre-training and rearranges the token sequence with the layout knowledge.

3 Methodology

LayoutMask is a multi-modal transformer that can encode text and layout information of documents and produce multi-modal representations. The pipeline of LayoutMask can be seen in Figure 2. LayoutMask uses the transformer model with a spatial-aware self-attention mechanism proposed in LayoutLmv2 (Xu et al., 2021) as the backbone and follows its preprocessing settings for text and layout embeddings. In Section 3.1, we will discuss the different choices of layout information in LayoutMask. In Section 3.2, we will introduce the pre-training tasks and masking strategies used in LayoutMask.

3.1 Selection of Layout Information

For VrDU tasks, there are two types of commonly used layout information: 1D position and 2D position. We list the 1D and 2D positions used in previous studies in Table 1.

1D Position: As we discussed in Section 1, using global 1D position will bring read order issues and could damage the adaptiveness and robustness of pre-trained models. Different from some previous models that leverage global 1D position as model input, we propose to use local 1D position in LayoutMask. Local 1D position only encodes the token orders within each segment and always

Method	Position	
	1D	2D
LayoutLM (Xu et al., 2020)	Global	Word
StructuralLM (Li et al., 2021a)	Global	Segment
LayoutLMv2 (Xu et al., 2021)	Global	Word
BROS (Hong et al., 2022)	Global	Segment [†]
LiLT (Wang et al., 2022)	Global	Segment
LayoutLMv3 (Huang et al., 2022)	Global	Segment
LayoutMask(Ours)	Local	Segment

Table 1: The 1D position and 2D position choices in previous studies. Our method uses local 1D position and segment-level 2D position. †: BROS leverages relative 2D positions instead of absolute positions.

restarts with 1 for each individual segment. Illustrations of the global and local 1D positions can be seen in Figure 1 and Figure 2. Compared with global 1D position, the major difference of using local 1D position is the lack of cross-segment orders, so the global reading order has to be inferred with other layout and semantic clues. Besides, the in-segment orders implied by local 1D position are more reliable and trustworthy than cross-segment orders when meeting complex document layouts.

2D Position: The 2D position is represented as a 4-digit vector like $[x_1, y_1, x_2, y_2]$, where $[x_1, y_1]$ and $[x_2, y_2]$ are the normalized coordinates of the top-left and bottom-right corners of a text box. There are two commonly used types of 2D positions: word-level 2D position (Word-2D) and segment-level 2D position (Segment-2D). For Word-2D, tokens of the same word will have the same word-level boxes as their 2D position. While for Segment-2D, the segment coordinates are shared by tokens within each segment.

In our model, we choose local 1D position and segment-level 2D position as our model input, where local 1D position can provide in-segment orders, and segment-level 2D position can provide cross-segment reading order clues, so the pre-trained model can learn the correct global reading order by jointly using 1D and 2D positions. We will compare the experimental results using different 1D & 2D position combinations in Section 4.3.1 and provide detailed discussions.

3.2 Pre-training Objectives

3.2.1 Masked Language Modeling

The Masked Language Modeling task is the most essential and commonly used pre-training task in multi-modal pre-training. In this task, we randomly mask some tokens with a given probability P_{mlm} (e.g., 15%) and recover these tokens during pre-

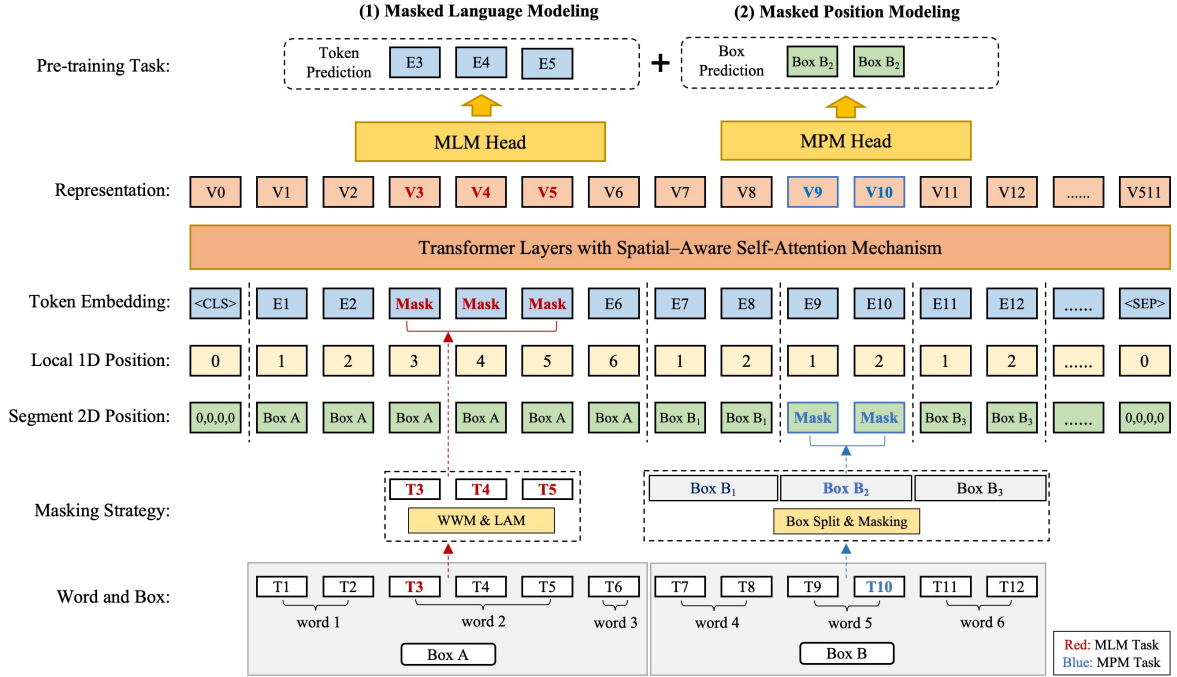


Figure 2: The model pipeline of LayoutMask. **Red Text:** Illustration of the Masked Language Modeling task. **Blue Text:** Illustration of the Masked Position Modeling task.

training.

For each document, we use M to denote the number of masked tokens. y_i and \bar{y}_i represent the ground truth and prediction of the i -th masked token. Then the loss of this task is the average cross entropy loss of all masked tokens:

$$\mathcal{L}_{\text{mlm}} = -\frac{1}{M} \sum_{i=1}^M \text{CE}(y_i, \bar{y}_i). \quad (1)$$

In preliminary experiments, we find that the naive MLM task is not optimal for multi-modal pre-training. Thus we propose to adopt two novel strategies, Whole Word Masking (WWM) and Layout-Aware Masking (LAM), to enhance this task.

Whole Word Masking: The WWM strategy was first proposed for Chinese language models to increase the task difficulty (Cui et al., 2021). Following this strategy, we set masks at word-level instead of token-level, which is much more challenging. When using WWM, the semantic relations between masked and unmasked tokens of the same words are eliminated, so the model has to find more context to predict masked words, which can promote text-layout interactions.

Layout-Aware Masking: As we use Local-1D and Segment-2D as model input, the global reading order should be obtained by jointly using 1D and 2D positions, where Local-1D provides in-segment orders and segment-2D provides cross-segment clues.

We find that the cross-segment orders are harder to be learned, so we propose Layout-Aware Masking (LAM) strategy to address this issue. Unlike naive masking strategy where each token has an equal masking probability P_{mlm} , in LAM strategy, the first and last word of each segment has a higher probability (*i.e.*, $3 \times P_{\text{mlm}}$) to be masked. In order to predict such masked words, the model has to pay more attention to finding their contexts in the preceding or succeeding segment, thus promoting learning cross-segment orders.

3.2.2 Masked Position Modeling

To further promote the representation learning of layout information in the MLM task, we design an auxiliary task, Masked Position Modeling (MPM), which has a symmetric pre-training objective: recovering randomly masked 2D positions during pre-training (See illustration in Figure 2). Inspired by WWM, we also apply the MPM task at word-level instead of token-level. For each pre-training document, we randomly choose some unduplicated words with a given probability P_{mpm} . Then, for each selected word, we mask their 2D positions with the following two steps:

Box Split: We first split the selected word out of its segment so the original segment box becomes 2 or 3 segment pieces (depending on if the word is at the start/end or in the middle). The selected word

Method	#Parameters	Modality	FUNSD(F1↑)	CORD(F1↑)	SROIE(F1↑)
BERT _{Base} (Devlin et al., 2019)	110M	T	60.26	89.68	90.99
RoBERTA _{Base} (Liu et al., 2019)	125M	T	66.48	93.54	-
UniLMv2 _{Base} (Bao et al., 2020)	125M	T	68.90	90.92	94.59
BROS _{Base} (Hong et al., 2022)	110M	T+L	83.05	95.73	95.48
LiLT _{Base} (Wang et al., 2022)	-	T+L	88.41	96.07	-
LayoutLM _{Base} (Xu et al., 2020)	160M	T+L+I	79.27	-	94.38
LayoutLMv2 _{Base} (Xu et al., 2021)	200M	T+L+I	82.76	94.95	96.25
TILT _{Base} (Powalski et al., 2021)	230M	T+L+I	-	95.11	97.65 [†]
DocFormer _{Base} (Appalaraju et al., 2021)	183M	T+L+I	83.34	96.33	-
LayoutLMv3 _{Base} (Huang et al., 2022)	133M	T+L+I	90.29	96.56	-
LayoutMask_{Base} (Ours)	182M	T+L	92.91±0.34	96.99±0.30	96.87±0.19
BERT _{Large} (Devlin et al., 2019)	340M	T	65.63	90.25	92.00
RoBERTA _{Large} (Liu et al., 2019)	355M	T	70.72	93.80	-
UniLMv2 _{Large} (Bao et al., 2020)	355M	T	72.57	92.05	94.88
LayoutLM _{Large} (Xu et al., 2020)	343M	T+L	77.89	-	95.24
BROS _{Large} (Hong et al., 2022)	340M	T+L	84.52	97.40	-
LayoutLMv2 _{Large} (Xu et al., 2021)	426M	T+L+I	84.2	96.01	97.81
TILT _{Large} (Powalski et al., 2021)	780M	T+L+I	-	96.33	98.10 [†]
DocFormer _{Large} (Appalaraju et al., 2021)	536M	T+L+I	84.55	96.99	-
LayoutLMv3 _{Large} (Huang et al., 2022)	368M	T+L+I	92.08	97.46	-
ERNIE-Layout _{Large} (Peng et al., 2022)	-	T+L+I	93.12	97.21	97.55
LayoutMask_{Large} (Ours)	404M	T+L	93.20±0.29	97.19±0.20	97.27±0.32

Table 2: F1 scores (%) of different methods on FUNSD, CORD, and SROIE. The best results are denoted in boldface. †: TILT utilized supervised datasets during pre-training, so the scores are not directly comparable.

becomes a one-word segment piece with just itself. Then we update the local 1D positions (restarting with 1) and segment 2D positions for each new segment piece. With the above operations, we can eliminate the local reading order clues implied by original 1D and 2D positions, so the model has to focus on semantical clues and new 2D positions.

Box Masking: For each selected word, we mask its 2D position with pseudo boxes: $[0, 0, 0, n]$ where $n \in [0, 1, 2, \dots]$ is a random number. Notice that segment 2D position is shared among tokens in the same segment, so the pseudo boxes will act as identifiers to distinguish identical tokens from different masked boxes, thus avoiding ambiguity.

During pre-training, our model is supposed to predict the masked 2D positions with GIoU loss (Rezatofighi et al., 2019):

$$\mathcal{L}_{\text{mpm}} = -\frac{1}{N} \sum_{i=1}^N \left(\frac{|B_i \cap \bar{B}_i|}{|B_i \cup \bar{B}_i|} - \frac{|C_i \setminus (B_i \cup \bar{B}_i)|}{|C_i|} \right). \quad (2)$$

Here, $i \in [1, 2, \dots, N]$ is the index of N masked 2D positions. B_i is the ground truth box normalized to $[0, 1]$, and \bar{B}_i denotes the predicted 2D position. C_i is the smallest convex shapes that covers B_i and \bar{B}_i . \mathcal{L}_{mpm} is the average GIoU loss of N masked 2D positions.

The MPM task is very similar to the cloze test, where a group of randomly selected words is supposed to be refilled at the right positions in the

original document. To predict the masked 2D positions of selected words, the model has to find the context for each word based on semantic relations and then infer with 2D position clues from a spatial perspective. The joint learning process with both semantic and spatial inference can promote text-layout interactions and help the model to learn better layout representations.

With the above two pre-training objectives, the model is pre-trained with the following loss:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{mlm}} + \lambda \mathcal{L}_{\text{mpm}}, \quad (3)$$

where λ is a hyper-parameter that controls the balance of the two pre-training objectives.

4 Experiments

4.1 Pre-training Settings

LayoutMask is pre-trained with IIT-CDIP Test Collection (Lewis et al., 2006). It contains about 42 million scanned document pages, and we only use 10 million pages. We use a public OCR engine, PaddleOCR¹ to obtain the OCR results.

We train LayoutMask with two parameter sizes. LayoutMask_{Base} has 12 layers with 16 heads, and the hidden size is 768. LayoutMask_{Large} has 24 layers with 16 heads where the hidden size is 1024. LayoutMask_{Base} and LayoutMask_{Large} are

¹<https://github.com/PaddlePaddle/PaddleOCR>

Method	Modality	Accuracy \uparrow	
		Base	Large
VGG-16 (Afzal et al., 2017)	I	90.97	
Ensemble (Das et al., 2018)	I	93.07	
LadderNet (Sarkhel and Nandi, 2019)	I	92.77	
BERT (Devlin et al., 2019)	T	89.81	89.92
RoBERTa (Liu et al., 2019)	T	90.06	90.11
UniLMv2 (Bao et al., 2020)	T+L	90.06	90.20
LayoutLM (Xu et al., 2020)	T+L	91.78	91.90
StructuralLM (Li et al., 2021a)	T+L	-	96.08
SelfDoc (Li et al., 2021b)	T+L+I	92.81	-
TITL (Powalski et al., 2021)	T+L+I	95.25	95.52
LayoutLMv2 (Xu et al., 2021)	T+L+I	95.25	95.64
DocFormer (Appalaraju et al., 2021)	T+L+I	96.17	95.50
LiLT (Wang et al., 2022)	T+L+I	95.68	-
LayoutLMv3 (Huang et al., 2022)	T+L+I	95.44	95.93
ERNIE-Layout (Peng et al., 2022)	T+L+I	-	96.27
LayoutMask (Ours)	T+L	93.26	93.80

Table 3: The accuracies (%) of different methods on RVL-CDIP dataset. For transformer-based models, we provide results for both base and large versions.

initialized with pre-trained XLM-RoBERTa models (Conneau et al., 2020).

For hyper-parameters, we have $P_{\text{mlm}}=25\%$ and $P_{\text{mpm}}=15\%$ (See ablation study in Section A of the Appendix). The weight of MPM loss λ is set to be 1.

4.2 Comparison with the State-of-the-Art

In this section, we compare LayoutMask with SOTA models on two VrDU tasks: form & receipt understanding and document image classification.

4.2.1 Form and Receipt Understanding

In this task, we conduct entity extraction task on three document understanding datasets: FUNSD (Jaume et al., 2019), CORD (Park et al., 2019), and SROIE (Huang et al., 2019). The FUNSD dataset is a form understanding dataset, which contains 199 documents (149 for training and 50 for test) and 9707 semantic entities. The CORD dataset is a receipt understanding dataset, and it contains 1000 receipts (800 for training, 100 for validation, and 100 for test) with 30 semantic labels in 4 categories. The SROIE dataset is another receipt understanding dataset with four types of entities, containing 626 receipts for training and 347 receipts for test.

For evaluation, we adopt the word-level F1 score as the evaluation metric for FUNSD and CORD and use the entity-level F1 score for SROIE. Since these datasets are quite small, in order to provide stable and reliable results, we repeat our experiments ten times for each test and report the average F1 scores and standard errors as the final results.

The results of previous methods and LayoutMask on these datasets are listed in Table 2. We have categorized them by the modalities used in

pre-training: “T” for text, “L” for layout, and “I” for image. Notice that LayoutMask is a “T+L” model that does not use image modality.

For the base version, LayoutMask_{Base} outperforms other methods, including “T+L+I” models, on all three datasets (FUNSD+2.62%, CORD+0.43%, SROIE+0.62%). For the large version, LayoutMask_{Large} ranks first on FUNSD and has comparable results on CORD and SROIE.

These results show that LayoutMask has competitive performance with SOTA methods, demonstrating the effectiveness of our proposed modules. Since LayoutMask only uses text and layout information, we believe that the potential power of layout information has not been fully explored in previous studies.

4.2.2 Document Image Classification

In the document image classification task, we aim to classify document images in RVL-CDIP dataset (Harley et al., 2015). This dataset is a subset of the IIT-CDIP collection with 400,000 labeled document images (320,000 for train, 40,000 for validation, and 40,000 for test) in 16 categories. We use PaddleOCR to extract text and layout information as model input. We compare different methods with the overall classification accuracies on RVL-CDIP, and the results are in Table 3.

It is observed that LayoutMask has beaten all unimodality models (“I” and “T”). For “T+L” models, LayoutMask_{Base} outperforms other base models with a margin of 1.48%, while LayoutMask_{Large} takes the second place in large models. Compared with “T+L+I” models where image modality is utilized, LayoutMask falls behind due to the lack of visual features from image modality. We have found that the image modality plays an important role in this task because RVL-CDIP images contain many elements that cannot be recognized by OCR engines (*e.g.*, figures, table lines, and handwritten texts) and have orientation issues (See examples in Figure 5 of the Appendix). So the lack of image modality will bring difficulties that cannot be solved with only text and layout information.

4.3 Ablation Study on LayoutMask

4.3.1 Comparison of Layout Information

We first compare the performance of LayoutMask using different layout information. To make a fair comparison, we use LayoutMask with only the MLM task and the WWM strategy during pre-training. For each test, LayoutMask is pre-trained

Position Settings		Datasets		
1D	2D	FUNSD (F1↑)	CORD (F1↑)	SROIE (F1↑)
Global	Word	82.17±0.45	95.95±0.43	96.02±0.34
Global	Segment	91.61±0.42	96.69±0.24	96.20±0.26
Local	Word	91.65±0.36	95.86±0.22	96.54±0.23
Local	Segment	92.30±0.24	96.68±0.12	96.56±0.21

Table 4: The average F1 scores (%) with different 1D position and 2D position combinations. The best results are denoted in boldface.

#	Position Settings 1D & 2D	Swap Probability	SROIE (F1↑)				Overall
			Address	Company	Date	Total	
1	Local+Segment	-	96.69±0.37	95.88±0.28	99.66±0.13	94.02±0.49	96.56±0.21
2	Global+Segment	-	96.54±0.51	95.84±0.59	99.69±0.26	92.73±0.57	96.20±0.26
3		10	91.73±2.00	95.22±0.61	99.65±0.34	91.87±1.33	94.62±0.69
4		20	90.03±3.77	94.93±0.53	99.60±0.32	91.67±1.35	94.06±1.02
5		30	88.12±4.59	94.88±0.82	99.55±0.28	91.19±1.38	93.44±1.14

Table 5: F1 scores (%) on SROIE dataset with difference 1D positions and increasing segment swap probabilities (%). We report both entity-level scores (“Address”, “Company”, “Date”, and “Total”) and overall results (“Overall”) for detailed comparison.

and fine-tuned with a specific 1D and 2D position combination. The results are listed in Table 4.

Performance of 1D Position: For 1D position, Local-1D outperforms Global-1D on both FUNSD (+9.48%/+0.69% with Word-2D/Segment-2D) and SROIE (+0.52%/+0.36%) and falls a little behind on CORD (-0.07%/-0.01%).

To understand the benefits of using Local-1D, we provide entity-level F1 score on SROIE dataset in Table 5 (#1 for Local+Segment and #2 for Global+Segment). It is obvious that the performance gap between Local+Segment and Global+Segment mainly comes from entity “Total” (from 94.02% to 92.72%), while other entities have similar F1 scores. We illustrate two example images of SROIE and their entities annotations in Figure 3. The right image, which contains entity “Total”, has both vertical layout (first two lines) and horizontal layout and has multiple misleading numbers with the same content as ground truth (*i.e.*, “193.00”). So it is hard to recognize the entity “Total” by using the ordinary reading order implied by Global-1D. Therefore, using Local-1D can perform better since it is more adaptive to such cases.

Performance of 2D Position: For 2D position, using segment-level 2D position brings better results on all three datasets, regardless of the 1D position types. An important reason is that the segment information is highly indicative of recognizing entities. For example, every entity in FUNSD and CORD exactly shares the same segment. Therefore, although Word-2D contains more layout details, it will break the alignments between 2D positions

and entities, thus bringing performance drops. A typical result of such phenomenon² can be seen on FUNSD, where replacing Global+Segment to Global+Word will result in a significant decrease of 9.44%.

Robustness Comparison: Besides performance superiority, another important reason to choose the local 1D position is its robustness to layout disturbance. In real-world cases, a typical layout disturbance is “Segment Swap”, where segments in the same line are indexed with wrong orders due to document rotation or OCR issues. In such scenarios, the incorrect cross-segment order will lead to incorrect global 1D positions and can be harmful to model inference. Fortunately, the local 1D position is naturally immune to such disturbance since it does not rely on cross-segment orders, making it more robust than global 1D position.

To quantify such differences in robustness, we demonstrate how the segment swap will influence the performance of using global 1D position by simulating it on test datasets. For each test document, we randomly choose some lines with a given probability P_{swap} and then swap the segments in it. We conduct experiments on LayoutMask_{Base} (MLM+WWM) in Global+Segment setting with different P_{swap} (*i.e.*, 10%, 20%, and 30%) and the results are reported in Table 5 (#3-5).

During our experiments, we have found that the segment swap does not bring significant perfor-

²Similar phenomenon can also be observed in the LayoutLM series models, where using Segment-2D increase F1 scores for about 8% on FUNSD dataset.

#	Pre-training Setting				Datasets			
	MLM	WWM	LAM	MPM	FUNSD (F1 \uparrow)	CORD (F1 \uparrow)	SROIE (F1 \uparrow)	RVL-CDIP (ACC \uparrow)
1	✓				89.73 \pm 0.50	96.32 \pm 0.15	95.76 \pm 0.34	92.17
2	✓	✓			92.30 \pm 0.24	96.68 \pm 0.12	96.56 \pm 0.21	92.89
3	✓	✓	✓		92.66 \pm 0.26	96.89 \pm 0.24	96.64 \pm 0.22	93.03
4	✓	✓		✓	92.77 \pm 0.30	96.84 \pm 0.17	96.66 \pm 0.32	93.11
5	✓	✓	✓	✓	92.91 \pm 0.34	96.99 \pm 0.30	96.87 \pm 0.19	93.26

Table 6: Performance analysis with different pre-training objectives and masking strategies.

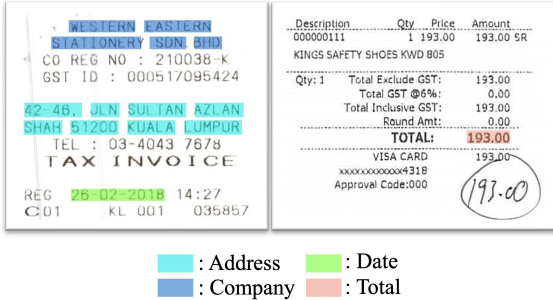


Figure 3: Two images from SROIE dataset. Colored boxes denote the ground truth of entities. The left image contains two cross-line and cross-segment entities (“Address” and “Company”). The right image, with a mixture of vertical and horizontal layouts, contains the “Total” entity.

performance changes on FUNSD and CORD datasets (so these results are not listed due to the limited space). A possible reason is that FUNSD and CORD do not contain cross-segment entities, so the segment swap can not break the order of words in each entity. Evidence for this explanation is that the SROIE dataset is significantly affected by segment swap, and its cross-segment entities (“Address” and “Company”) have obvious performance drops. In SROIE, the majority of “Address” entities and a few “Company” entities are printed in multiple lines (See examples in Figure 3), so the segment swap can change the in-entity orders of entity words. The results show that the “Address” entity has the largest drop among all entities (-4.81%, -6.51%, and -8.42% for $P_{\text{swap}}=10\%$, 20%, 30%). Besides, the “Total” entity has the second largest decrease (-0.86%, -1.06%, and -1.54%). As aforementioned, the “Total” entities are usually surrounded by complex layouts and misleading numbers, so the segment swap will bring extra difficulties in recognizing the correct entities.

The above performance decreases of using global 1D position prove the superiority of using local 1D position since it is not affected by such layout disturbance and can have more robust performance in real-world scenarios.

4.3.2 Effectiveness of Proposed Methods

In Table 6, we provide results using different pre-training tasks and masking strategies to demonstrate the effectiveness of our proposed modules.

Comparing #1 and #2 in Table 6, we observe that WWM brings significant performance improvements on all datasets. The reason is that it increases the difficulty of the MLM task, so we can obtain a stronger language model. We also find that LAM can also bring consistent improvements on all dataset because LAM can force the model to learn better representations for layout information, which is beneficial to downstream tasks.

Comparing #2 to #4 and #3 to #5, it is observed that the MPM task also brings considerable improvements on all datasets. MPM works as an auxiliary task to help the MLM task and can increase the pre-training difficulty, contributing to learning better and more robust layout representations.

Moreover, the full-version LayoutMask (#5) outperforms the naive version (#1) by a large margin (FUNSD+3.18%, CORD+0.67%, SROIE+1.11%, and RVL-CDIP+1.09%), demonstrating the effectiveness of our proposed modules when working together. To better illustrate the effectiveness of our model design, we list category-level accuracy improvements on RVL-CDIP dataset and provide detailed discussions in Section B of the Appendix.

5 Conclusion

In this paper, we propose LayoutMask, a novel multi-modal pre-training model, to solve the reading order issues in VrDU tasks. LayoutMask adopts local 1D position as layout input and can generate adaptive and robust multi-modal representations. In LayoutMask, we equip the MLM task with two masking strategies and design a novel pre-training objective, Masked Position Modeling, to enhance the text-layout interactions and layout representation learning. With only using text and layout modalities, our method can achieve excellent results and significantly outperforms many SOTA methods in VrDU tasks.

Limitations

Our method has the following limitations:

Datasets: In multi-modal pre-training, we rely on downstream datasets to evaluate the performance of pre-trained models. The commonly used entity extraction datasets are relatively small and lack diversity, so the proposed method may not generalize well to real word scenarios.

Lack of Image Modality: In LayoutMask, we focus on text-layout interactions, leaving the image modality unexplored. However, documents in the real world contain many elements that can not be described by text and layout modalities, like figures and lines, so incorporating image modality is important in building a universal multi-modal pre-training model for document understanding.

References

- Muhammad Zeshan Afzal, Andreas Kölsch, Sheraz Ahmed, and Marcus Liwicki. 2017. Cutting the error by half: Investigation of very deep cnn and advanced training strategies for document image classification. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, volume 1, pages 883–888. IEEE.
- 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.
- Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. 2021. Beit: Bert pre-training of image transformers. In *International Conference on Learning Representations*.
- Hangbo Bao, Li Dong, Furu Wei, Wenhui Wang, Nan Yang, Xiaodong Liu, Yu Wang, Jianfeng Gao, Songhao Piao, Ming Zhou, et al. 2020. Unilmv2: Pseudo-masked language models for unified language model pre-training. In *International Conference on Machine Learning*, pages 642–652. PMLR.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Édouard 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.
- Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, and Ziqing Yang. 2021. Pre-training with whole word masking for chinese bert. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3504–3514.
- Arindam Das, Saikat Roy, Ujjwal Bhattacharya, and Swapan K Parui. 2018. Document image classification with intra-domain transfer learning and stacked generalization of deep convolutional neural networks. In *2018 24th international conference on pattern recognition (ICPR)*, pages 3180–3185. IEEE.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- 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.
- 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 for 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.
- 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.
- Anoop R Katti, Christian Reisswig, Cordula Guder, Sebastian Brarda, Steffen Bickel, Johannes Höhne, and Jean Baptiste Faddoul. 2018. Chargrid: Towards understanding 2d documents. In *EMNLP*.
- David Lewis, Gady Agam, Shlomo Argamon, Ophir Frieder, David Grossman, and Jefferson Heard. 2006. Building a test collection for complex document information processing. In *Proceedings of the 29th*

- annual international ACM SIGIR conference on Research and development in information retrieval*, pages 665–666.
- Chenliang Li, Bin Bi, Ming Yan, Wei Wang, Songfang Huang, Fei Huang, and Luo Si. 2021a. Structurallm: Structural pre-training for form 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 6309–6318.
- Junlong Li, Yiheng Xu, Tengchao Lv, Lei Cui, Cha Zhang, and Furu Wei. 2022. Dit: Self-supervised pre-training for document image transformer. In *MM '22: The 30th ACM International Conference on Multimedia, Lisboa, Portugal, October 10 - 14, 2022*, pages 3530–3539. ACM.
- Peizhao Li, Jiuxiang Gu, Jason Kuen, Vlad I Morariu, Handong Zhao, Rajiv Jain, Varun Manjunatha, and Hongfu Liu. 2021b. Selfdoc: Self-supervised document representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5652–5660.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- 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 *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*.
- 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.
- Hamid Rezatofghi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. 2019. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 658–666.
- Ritesh Sarkhel and Arnab Nandi. 2019. Deterministic routing between layout abstractions for multi-scale classification of visually rich documents. In *28th International Joint Conference on Artificial Intelligence (IJCAI), 2019*.
- Jiapeng Wang, Lianwen Jin, and Kai Ding. 2022. 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.
- Zilong Wang, Yiheng Xu, Lei Cui, Jingbo Shang, and Furu Wei. 2021. Layoutreader: Pre-training of text and layout for reading order detection. In *EMNLP (1)*.
- Yang Xu, Yiheng Xu, Tengchao Lv, Lei Cui, Furu Wei, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, Wanxiang Che, et al. 2021. 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, 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*, pages 1192–1200.
- Xiao Yang, Ersin Yumer, Paul Asente, Mike Kraley, Daniel Kifer, and C Lee Giles. 2017. Learning to extract semantic structure from documents using multimodal fully convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5315–5324.
- Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: human language technologies*, pages 1480–1489.
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. Ernie: Enhanced language representation with informative entities. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1441–1451.

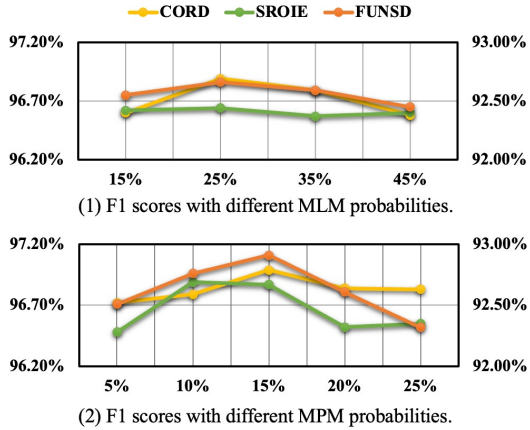


Figure 4: The F1 scores on FUNSD, CORD, and SROIE with different masking probabilities. FUNSD dataset uses the x-axis on the right side.

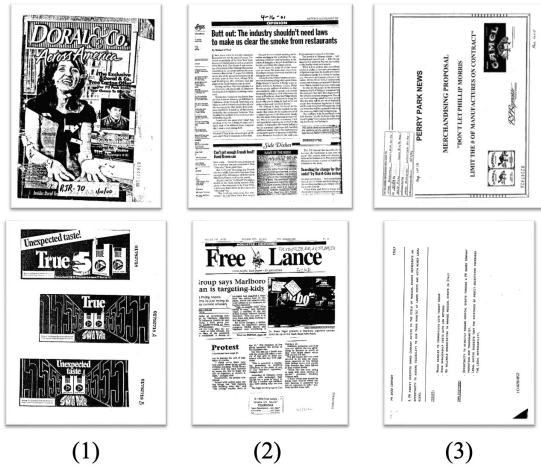


Figure 5: RVL-CDIP images from different categories. (1) Advertisement; (2) News article; (3) Presentation (with incorrect orientations).

A Ablation Study of Masking Probabilities

We compare LayoutMask using different P_{mlm} and P_{mpm} , and the results are in Figure 4. We first find the best P_{mlm} without using the MPM task, and the optimal value is 25%. Then we fix such optimal P_{mlm} to find the best P_{mpm} , which is 15% as the results show.

B Ablation Study on RVL-CDIP

To further understand the effectiveness of our model design, we list the detailed classification results on RVL-CDIP dataset with the naive version and the full version in Table 7. It is observed that the major performance improvements come from three categories: presentation (+3.36%), ad-

Category	Model Settings		Diff. (%)
	Naive	Full	
letter	90.30	90.86	0.56
form	85.71	86.77	1.07
email	98.17	98.33	0.15
handwritten	93.96	94.26	0.30
advertisement	88.47	91.40	2.93
sci-report	87.87	89.38	1.51
sci-publication	93.08	93.73	0.65
specification	95.91	96.56	0.64
file folder	91.29	92.71	1.42
news article	90.09	92.44	2.35
budget	94.01	94.96	0.95
invoice	94.02	94.54	0.52
presentation	86.14	89.50	3.36
questionnaire	92.44	92.88	0.44
resume	98.31	98.70	0.39
memo	94.93	95.12	0.19
Overall	92.17	93.26	1.09

Table 7: The category-level accuracies (%) on RVL-CDIP dataset of LayoutMask on the naive version and the full version. Categories with top-3 accuracy improvements are denoted in boldface.

vertisement (+2.93%), and news article (+2.35%). We find these categories have more diverse layouts (See examples in Figure 5), so classifying these documents requires a better understanding of the document structure, which also indicates the effectiveness of our methods in helping layout understanding.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
in Section Limitations
- A2. Did you discuss any potential risks of your work?
in Section 4
- A3. Do the abstract and introduction summarize the paper’s main claims?
in Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

in Section 4

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

C Did you run computational experiments?

In Section 4

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
In Section 4

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

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

In Section 4

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

In Section 4

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

In Section 4

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

Left blank.

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

No response.

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

No response.

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

No response.

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

No response.

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

No response.