

# Modeling Entities as Semantic Points for Visual Information Extraction in the Wild

Zhibo Yang<sup>1,2\*</sup> Rujiao Long<sup>2\*†</sup> Pengfei Wang<sup>2\*</sup> Sibao Song<sup>2</sup> Humen Zhong<sup>2</sup>  
 Wenqing Cheng<sup>1</sup> Xiang Bai<sup>1</sup> Cong Yao<sup>2</sup>  
<sup>1</sup>Huazhong University of Science and Technology <sup>2</sup>Alibaba Group




Figure 1. Demonstration of the proposed SIBR dataset, where real-world complexity and challenges (e.g., overlap, blur, printing shift, rotation, and complex linking) are well represented. The rectangles in red and green stand for keys and values, while the yellow and blue arrows indicate inter-entity and intra-entity linkings, respectively. (Best viewed in color.)

## Abstract

Recently, Visual Information Extraction (VIE) has been becoming increasingly important in both the academia and industry, due to the wide range of real-world applications. Previously, numerous works have been proposed to tackle this problem. However, the benchmarks used to assess these methods are relatively plain, i.e., scenarios with real-world complexity are not fully represented in these benchmarks. As the first contribution of this work, we curate and release a new dataset for VIE, in which the document images are much more challenging in that they are taken from real applications, and difficulties such as blur, partial occlusion, and printing shift are quite common. All these factors may lead to failures in information extraction. Therefore, as the second contribution, we explore an alternative approach to precisely and robustly extract key information from document images under such tough conditions. Specifically, in contrast to previous methods, which usually either incorporate visual information into a multi-modal architecture or train text spotting and information extraction in an end-to-end fashion, we explicitly model entities as semantic points, i.e., center points of entities are enriched with semantic information describing the attributes and relationships of different entities, which could largely benefit entity labeling and linking. Extensive experiments on standard benchmarks in this field as well as the proposed

dataset demonstrate that the proposed method can achieve significantly enhanced performance on entity labeling and linking, compared with previous state-of-the-art models. Dataset is available at <https://www.modelscope.cn/datasets/damo/SIBR/summary>.

## 1. Introduction

Visually Rich Documents (VRDs) are ubiquitous in daily, industrial, and commercial activities, such as receipts of shopping, reports of physical examination, product manuals, and bills of entry. Visual Information Extraction (VIE) aims to automatically extract key information from these VRDs, which can significantly facilitate subsequent processing and analysis. Due to its broad applications and grand technical challenges, VIE has recently attracted considerable attention from both the Computer Vision community [33, 34, 39] and the Natural Language Processing community [16, 35, 37]. Typical techniques for tackling this challenging problem include essential electronic conversion of image (OCR) [28, 30, 41], intermediate procedure of structure analysis [25] and high-level understanding of contents [35], among which entities play an important role as an aggregation of vision, structure, and language.

Though substantial progresses [11, 19, 35] have been made, it is still challenging to precisely and reliably extract key information from document images in unconstrained conditions. As shown in Fig. 1, in real-world scenarios documents may have various formats, be captured casually with

\*Equal Contribution.

†Correspondence Author.

a mobile phone, or exist occlusion or shift in printing, all of which would pose difficulties for VIE algorithms.

To highlight the challenges in real applications and promote the development of research in VIE, we establish a new dataset called **Structurally-rich Invoices, Bills and Receipts in the Wild (SIBR for short)**, which contains 1,000 images with 71,227 annotated entity instances and 39,004 entity links. The challenges of SIBR lie in: (1) The documents are from different real-world scenarios, so their formats and structures might be complicated and varying; (2) The image quality may be very poor, i.e., blur, noise, and uneven illumination are frequently seen; (3) The printing process is imperfect that shift and rotation might happen.

To deal with these difficulties, we explore an novel approach for information extraction from VRDs. Different from previous methods, which usually employ a sequential pipeline that first uses an off-the-shelf OCR engine to detect and read textual information (location and content) and then fuses such information with visual cues for follow-up entity labeling (*a.k.a.* entity extraction) and linking in a multi-modal architecture (mostly a Transformer) [19, 37], the proposed method adopts a unified framework that all components, including text detection, text recognition, entity extraction and linking, are jointly modeled and trained in an integrated way. This means that in our method a separate OCR engine is no longer necessary. The benefits are two-fold: (1) The accuracy of entity labeling and linking will not be limited by the capacity of the OCR engine; (2) The running speed of the whole pipeline could be boosted.

Drawing inspirations from general object detection [4, 15, 40, 42] and vision-language joint learning [11, 14, 31], we put forward to model entities as semantic points (**ESP for short**). Specifically, as shown in Fig. 3, entities are represented using their center points, which are enriched with semantics, such as geometric and linguistic information, to perform entity labeling and linking. To better learn a joint vision-language representation, we also devise three training tasks that are well integrated into the paradigm. The entity-image text matching (**EITM**) task, which is only used in the pre-training stage, learns to align entity-level vision vectors and language vectors (encoded with off-the-shell BERT) with a contrastive learning paradigm. Entity extraction (**EE**) and Entity linking (**EL**), the main tasks for VIE, are used in the pre-training, fine-tuning, and inference stages. In these two modules, region features and position embedding (from ground truth or detection branch) are encoded with transformer layers and then decoded to entity classes and relations. Owing to the joint vision-language representation, text recognition is no longer a necessary module in our framework, and we will discuss the impact of the text recognition branch in Sec. 5.4.

Extensive experiments have been conducted on standard benchmarks for VIE (such as FUNSD, XFUND, and

CORD) as well as the proposed SIBR dataset. We found that compared with previous state-of-the-art methods, the proposed ESP algorithm can achieve highly competitive performance. Especially, it shows an advantage in the task of entity linking. Our main contributions can be summarized as follows: (1) We curate and release a new dataset for VIE, in which the document images are with real-world complexity and difficulties. (2) We devise a unified framework for spotting, labeling and linking entities, where a separate OCR engine is unnecessary. (3) We adopt three vision-language joint modeling tasks for learning informative representation for VIE. (4) Extensive experiments demonstrate the effectiveness and advantage of our approach.

## 2. Related Work

The proposed approach has been inspired by ideas from object detection and document understanding. Exemplar approaches include detection with grouping [24], text spotting [23, 43] and information extraction joint training [34], and vision-language pre-training for VIE [35].

**Detection with Grouping Methods** consider an object as a collection of parts arranged in specific relations. Object as points methods [4, 15, 40] firstly detected four corner points and one center point of an object and then grouped them with the geometric-based algorithm or associative embedding features. SegLink [27] and its variant SegLink++ [32] took a small segment of a text instance as the fundamental element and linked them together to form the bounding boxes. In [8, 9], hierarchical part-whole ideas were approached at the model level and feature space. Visual entity labeling and linking can also be learned in this manner to some extent. We model visual entities as points and group them with their semantic relations.

**Vision-Language Pre-training Methods** usually extract visual document information by using a 2D positional encoding of BERT-like architecture. The innovation of these works lies in novel pre-training tasks or new multi-modal transformer layers. LayoutLM [35] and its variant LayoutLM-v2 [37] and layoutLM-v3 proposed several novel pretraining tasks, including Masked Visual-Language Model, Text-Image Alignment, Text-Image Matching, and Word-Patch Alignment. StrucText [19] introduced a multi-granularity image and text alignment for entity labeling. BROS [10] proposed an area-masking self-supervision strategy that reflected the 2D nature of text blocks. In addition to pre-training tasks, DocFormer [1] and SelfDoc [18] also introduced cross-attention layers for better fusing of visual and language features. LiLT [33] extended document understanding to multi-languages by designing language-independent attention layers. The above methods usually depended on offline OCR engines, which may lead to error amplification and redundant computation.

**Jointly Training Methods** aim at improving information




Figure 2. Comparisons of color distribution across different datasets. A distribution of background from 5 to 250 is for better readability. (Best viewed in color.)

Element	SROIE	FUNSD	XFUND	CORD	EPHOIE	SIBR
Box	45279	9529	97012	19887	15754	71227
Link	-	10420	102291	7470	-	39004
Image	1139	199	1393	1000	1494	1000
Scenario	Scan	Scan	Scan	CamC	Scan	CamC
Link Type	-	1	1	1	-	2
Language	1	1	7	1	2	2
Overlap	7.4%	9.5%	5.1%	10.1%	1.6%	42.7%

Table 1. Comparison between SIBR and other datasets. “Scan” is short for scanned receipts or documents, “CamC” is short for Camera-captured images. “Overlap” showcases the proportion of images with overlapping entity boxes.

extraction by a unified end-to-end trainable framework. TRIE [39] and VIES [34] both proposed networks for simultaneous text detection, text recognition, and information extraction. VIES introduced a vision coordination mechanism and semantics coordination mechanism to gather rich visual and semantic features from text detection and recognition respectively. TRIE designed a multi-modal context block to bridge the OCR and IE modules. Indeed, TRIE and VIES both explicitly classified multi-modal information to get final entity labels without any general visual-language pre-training, but their drawback lies in the poor generalization ability to different downstream tasks.

### 3. The SIBR Dataset

In this section, we introduce Structurally-rich Invoices, Bills and Receipts in the Wild (SIBR) dataset.

**Dataset Description** With the wide spread of mobile phones, photoed document images are ubiquitous in industrial and commercial activities, among which photoed receipts are of great value for their probative force to business transactions. There are 1000 images in the SIBR, including 600 Chinese invoices, 300 English bills of entry, and 100 bilingual receipts. SIBR is well annotated with 71227 entity-level boxes and 39004 links. The dataset is divided into a training set of 600 images and a test set of 400 images. Details are shown in Table 1.

**Dataset Challenge** Previously, several datasets [13, 26] have been constructed for VIE. However, they are relatively plain, i.e., scenarios with real-world complexity are not well

represented in these benchmarks. Compared with real scene datasets such as (SROIE [12] and EPHOIE [34]), SIBR has more diverse appearances and richer structures. The document images in SIBR are much more challenging in that they are taken from real applications and difficulties (Fig. 1) such as severe noise, uneven illumination, image deformation, printing shift and complicated links. As shown in Fig. 2, SIBR has a wider distribution of RGB value. Due to image deformation or printing shift, SIBR has 42.7% of images with overlapping boxes, as shown in the bottom line of Table 1, which brings great challenges to text spotting and information extraction.

**Annotation Details** Following [34], four vertices are used to represent the bounding box of the entity. The label of entity is categorized as “question”, “answer”, “header”, and “other”. The link between two entities is categorized as either an intra-link or an inter-link, which represents links in an entity or between two different entities. In addition, text content is also annotated for end-to-end information extraction. It is worth noting that the granularity of text boxes is different in SIBR and FUNSD [13]. FUNSD is annotated at the semantic block level which may contain multiple lines, while the entity with multiple lines in SIBR is represented by text segments and intra-links between segments. As shown in Fig. 1, the blue lines represent intra-entity links, while the yellow lines represent inter-entity links. With the help of [38], the sensitive information (names, telephone numbers, etc.) is erased and re-edited with new text.

## 4. Methodology

In this section, we introduce our ESP framework for visual information extraction tasks, including entity extraction (EE) and entity linking (EL). The proposed ESP is first pre-trained with pseudo-labels and objectives which are carefully designed to learn entity semantics as well as the entity relations. Then we fine-tune our model for downstream tasks using task-specific annotations.

### 4.1. Overall Architecture

The goal of a VIE algorithm is to learn to locate key text contents and extract entity labels and linking information from visually-rich documents. We approach VIE from a point-based perspective, i.e., we learn to represent each entity, its label and relations with other entities by a single point. Specifically, we adopt the CenterNet framework [40], owing to its simple and concise pipeline. By utilizing the unified paradigm for feature learning, the model can benefit from joint learning across different modalities and tasks.

**Visual Backbone** We choose ConvNeXt-tiny-FPN as the visual backbone for its excellent trade-off between performance and efficiency. ConvNeXt [22] is a pure CNN architecture, which compares favorably with state-of-the-art hierarchical vision Transformers across multiple vision

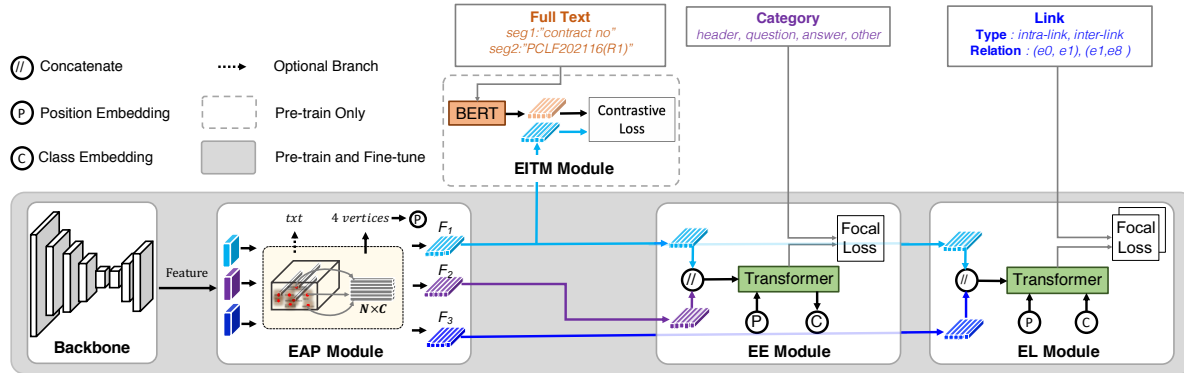


Figure 3. Schematic illustration of the proposed ESP. A ConvNeXt-FPN backbone is used to extract image features. Fed with three convolutional layers appended to the ConvNeXt, EAP module generates three kinds of outputs: (a) Detections, (b) Text contents, (c) Center-based Features: vision-language embedding  $F_1$ , entity type embedding  $F_2$ , and entity relational embedding  $F_3$ . Then, EITM aligns  $F_1$  with text embeddings given by BERT by a contrastive learning strategy. In the final decoding stage, EE module takes concatenated features of  $F_1$  and  $F_2$  as input and predicts categories of the corresponding entities. The process of entity linking (EL) is the same as EE. Note that the recognition branch is optional, as it has little influence on entity extraction, which will be discussed in ablation studies.

benchmarks while retaining the simplicity and efficiency of ConvNets. Moreover, we adopt feature pyramid network (FPN) [20] to fully exploit multi-scale visual features. Let  $I \in \mathbb{R}^{w \times h \times 3}$  be the input image of width  $w$  and height  $h$ . The visual backbone then produces a feature map  $P_2 \in \mathbb{R}^{\frac{w}{4} \times \frac{h}{4} \times 256}$  as the input to the EAP module.

**Entities-As-Points (EAP)** Given an image, CenterNet [40] extracts a set of bounding boxes for each object class with the help of a heatmap head, a size head and an offset head. We build the EAP module upon the original CenterNet framework with the following architectural modifications.

First, we devise the size head into a quadrangle head, which predicts the location of four vertices to detect arbitrary-oriented entity text instances. Therefore, the total text detection loss is computed as follows:

$$\mathcal{L}_{\text{Det}} = \mathcal{L}_{\text{Heatmap}} + \mathcal{L}_{\text{Quad}} + \mathcal{L}_{\text{Offset}} \quad (1)$$

Second, we propose three feature transformation branches to encode cross-modal features, entity type features, and entity relational features respectively. Each feature transformation branch consists of a  $3 \times 3$  convolutional layer, a ReLU layer, and a final  $1 \times 1$  convolutional layer. Moreover, the transformed feature maps are extracted based on the predicted points from the text detection branch. In this way, cross-modal features and entity-related features are produced in a unified point-based representation. Third, we adopt a simplified version of CRNN [29] as the text recognition branch. The predicted entity boxes will be fed into a grid-sample module to sample the text features  $F_t \in \mathbb{R}^{N \times C \times S_h \times S_w}$ .  $S_h$  and  $S_w$  denote the sampling size of height and width respectively while  $N$  denotes the number of text segments. Our CRNN variant is composed of 6 convolutional layers for feature extraction, 1 bi-directional LSTM layer for sequential modeling, and 1 linear projection layer for label prediction. Connectionist Temporal Classification (CTC) loss [5] is adopted as recognition

loss  $\mathcal{L}_{\text{Rec}}$  to align sequence predictions with the labels. In fact, our model does not depend on the recognition module, which is only used to output the end-to-end results, so we did not use the recognition results on FUNSD, XFUND, and CORD.

Overall, the point-based nature of entity representation possesses several advantages. First, our proposed method learns visual, textual, and layout representations simultaneously on a grid-based feature map, avoiding sophisticated serialization steps as for sequence-based methods [11, 35, 36], especially in real-world scenarios (e.g., misaligned second-time printing, receipts with complex layout) where such linearization step is nontrivial. Second, the text is usually truncated with a preset maximum number of tokens for sequence-based approaches. Nevertheless, our proposed model is not limited by a token number and can learn more holistic representations to mine inter-text relationships without any text truncation.

**Entity-Image Text Matching (EITM)** Recent progress in VIE has been mainly powered by multi-modal pre-training. Therefore, we harness an entity-image text matching module during the pre-training stage to enhance cross-modal feature learning. An image-text matching task is adopted to encourage cross-modal interactions and mutual alignment between visual and textual representations. EITM is only applied to the pre-training stage, thus enabling recognition-free prediction for entity labeling and linking. We employ the idea of contrastive learning to mutually align entity image and entity text representations into semantic space. It aims to find the closest entity image embedding from a batch given an entity text embedding. Similarly, for a given entity image embedding, the objective is to find the closest entity text embedding. This results in an InfoNCE loss over each batch of entity image-text pairs for pre-training our model,

$$L_{v2t} = - \sum_{j=1}^N \log \frac{\exp(v_i \cdot t_j / \tau)}{\sum_{k=1}^N \exp(v_i \cdot t_k / \tau)} \quad (2)$$

where  $N$  denotes the number of image-text pairs which equals the number of text segments,  $v_i$  and  $t_i$  represent  $i$ -th entity image embedding and entity text embedding respectively, and  $\tau$  denotes the hyperparameter of temperature. Similarly, for each  $t_i$ , the loss is formulated as,

$$L_{t2v} = - \sum_{j=1}^N \log \frac{\exp(t_i \cdot v_j / \tau)}{\sum_{k=1}^N \exp(t_i \cdot v_k / \tau)} \quad (3)$$

Finally, the total loss for EITM is defined as,

$$L_{\text{EITM}} = L_{v2t} + L_{t2v} \quad (4)$$

**Entity Extraction (EE)** To strengthen the entity type feature learning, we first concatenate vision-language features  $F_1$  and entity type features  $F_2$  to leverage both visual and semantic representations.

We then employ two self-attention layers to aggregate features from all entities. Last, the aggregated features are fed into the entity classifier, which is a ReLU-fc-ReLU-fc-Sigmoid block. We set the hidden size, number of self-attention heads, feed-forward/filter size, and number of transformer layers to 832, 8, 64, and 2 respectively. The dimension of classifier output is  $N \times N_c$ , where  $N$  denotes the number of text segments and  $N_c$  denotes the number of entity classes.

To address the extreme imbalance of entity types, we adopt Focal loss [21] as our loss function,

$$L_{\text{EE}} = \begin{cases} -(1 - p_t)^\gamma \log(p_t), & p_t^* = 1 \\ -p_t^\gamma \log(1 - p_t), & \text{otherwise} \end{cases} \quad (5)$$

where  $p_t$  is the predicted probability vector and  $p_t^*$  is the ground truth. We set  $\gamma$  to 4.

**Entity Linking (EL)** In many VIE applications, an entity often contains multiple text segments or text boxes. Therefore, it is important to model intra-entity links in addition to inter-entity links. Inter-entity links rely more on the spatial correlations of text segments, however, intra-entity links rely heavily on text semantics. Thus, we concatenate the cross-modal feature  $F_1$  with entity relational feature  $F_3$  to handle the lack of semantics and apply two layers' self-attention modules to capture inter-text dependencies.

The structure of the module is the same as Entity Extraction Module. The only difference is that the Entity Linking Module constructs a relationship matrix between entity text instances. We denote the output of the self-attention layers by  $E_l' \in \mathbb{R}^{N_l \times C}$ , where  $C$  is the dimension of input feature embeddings and  $N_l$  is the number of key-value pairs. First, we broadcast  $E_l'$  to obtain  $E_{l1} \in \mathbb{R}^{1 \times N_l \times d}$  and  $E_{l2} \in \mathbb{R}^{N_l \times 1 \times d}$ . Since entity links are vectors with

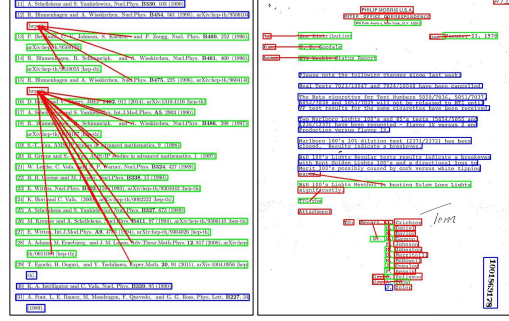


Figure 4. Example of pseudo labels. The boxes in red, green, and blue represent pseudo keys, pseudo values, and unmatched text contents, respectively. Key-value pairs are linked with red lines.

directions, we formulate the relational matrix via element-wise subtraction as given in Eq. (6) to encourage learning directed linking representations,

$$E_l^* = E_{l1} - E_{l2}, \quad E_l^* \in \mathbb{R}^{N_l \times N_l \times C} \quad (6)$$

$E_l^*$  is input into the classifier to obtain link results. The classifier is a ReLU-fc-ReLU-fc-Sigmoid block. Since the entity linking matrices are usually sparse, we use Focal loss as the same as the entity labeling task and denote as  $L_{\text{EL}}$ .

## 4.2. Training Objectives

**Pre-Training** The pre-training is conducted on DocBank and RVL-CDIP datasets with all aforementioned modules. Unlike previous methods that focus on designing mask-based or cross-alignment tasks, we generate pseudo-labels to mine semantic and relational information in these pre-training datasets. See Sec. 5.1 for more details. Thus, the overall pre-training objective is,

$$\mathcal{L}_{\text{PT}} = \mathcal{L}_{\text{Det}} + \mathcal{L}_{\text{EITM}} + \lambda_{\text{EL}} \mathcal{L}_{\text{EL}} + \lambda_{\text{EE}} \mathcal{L}_{\text{EE}} \quad (7)$$

**Fine-Tuning** Since EITM task only exists in the pre-training stage and recognition tasks only exist in fine-tuning, the loss term for fine-tuning is,

$$\mathcal{L}_{\text{FT}} = \mathcal{L}_{\text{Det}} + \mathcal{L}_{\text{Rec}}^* + \lambda_{\text{EL}} \mathcal{L}_{\text{EL}} + \lambda_{\text{EE}} \mathcal{L}_{\text{EE}} \quad (8)$$

where  $\lambda_{\text{EL}}$  is set to 1 for pre-training and 10 for fine-tuning and  $\lambda_{\text{EE}}$  is set to 10 for pre-training and 200 for fine-tuning. It is worth noting that, the recognition branch in our framework is optional, i.e., it can be removed or replaced by any off-the-shelf text recognizer given entity images.

## 5. Experiments

### 5.1. Datasets

DocBank and RVL-CDIP are used for pre-training, while the FUNSD [13], XFUND [34] and CORD [26] are adopted to compare with SOTA methods to verify the effectiveness of the ESP on EE and EL tasks.

Following [19], we use both **DocBank** and **RVL-CDIP** for pre-training. DocBank contains 500K images, of which

Model	Parameters	Pretrain Data Size	FUNSD		XFUND		CORD	
			EE	EL	EE	EL	EE	EL
BERT <sub>BASE</sub> [35]	110M + $\alpha$	-	60.92	27.65	-	-	89.68	92.83
RoBERTa <sub>BASE</sub> [2]	125M + $\alpha$	-	66.48	-	-	47.69	93.54	-
LayoutLM <sub>BASE</sub> [35]	160M + $\alpha$	11M	88.41	45.86	-	-	96.07	95.21
SelfDoc [18]	-	0.32M	83.36	-	-	-	-	-
UDoc [6]	272M + $\alpha$	11M	87.93	-	-	-	-	-
StrucText [19]	107M + $\alpha$	0.9M	83.09	44.10	-	-	-	-
LayoutLMv2 <sub>BASE</sub> [37]	200M + $\alpha$	11M	82.76	42.91	-	-	94.95	95.59
LiLT <sub>BASE</sub> [33]	-	11M	85.74	62.76	85.85	81.25	-	-
XYLayout <sub>BASE</sub> [7]	345M + $\alpha$	30M	83.35	-	82.04	67.79	-	-
BROS <sub>BASE</sub> [10]	110M + $\alpha$	11M	83.05	71.46	-	-	96.50	95.73
LayoutLMv3 <sub>BASE</sub> [11]	133M + $\alpha$	11M	90.29	-	-	-	96.56	-
BERT <sub>LARGE</sub> [35]	340M + $\alpha$	-	65.63	29.11	-	-	90.25	94.31
RoBERTa <sub>LARGE</sub> [2]	355M + $\alpha$	-	70.72	-	-	-	93.80	-
LayoutLM <sub>LARGE</sub> [35]	343M + $\alpha$	11M	77.89	42.83	-	-	-	95.41
LayoutLMv2 <sub>LARGE</sub> [37]	426M + $\alpha$	11M	84.20	70.57	-	-	96.01	97.29
BROS <sub>LARGE</sub> [10]	340M + $\alpha$	11M	84.52	77.01	-	-	97.28	97.40
DocFormer <sub>LARGE</sub> [1]	536M + $\alpha$	5M	84.55	-	-	-	96.99	-
LayoutLMv3 <sub>LARGE</sub> [11]	368M + $\alpha$	11M	<b>92.08</b>	80.35	-	-	<b>97.46</b>	98.28
ESP (ours)	50M	0.9M	91.12	<b>88.88</b>	<b>89.13</b>	<b>92.31</b>	95.65	<b>98.80</b>

Table 2. Comparisons on Entity Extraction and Entity Linking tasks.  $\alpha$ : parameters for OCR engine should be considered.

400K are for training, 50K for validation, and 50K for testing. RVL-CDIP includes 400K grayscale images in 16 categories, with 25K images per category. 320K are for training, 40K for validation, and 40K for testing.

**FUNSD** is composed of 149 training samples and 50 testing samples, where the images are scanned English documents. **XFUND** extends the FUNSD dataset to 7 other languages, with the same number of images in each language as the FUNSD. **CORD** consists of 1000 fully annotated receipts, including 800 training samples, 100 validation samples, and 100 testing samples. The FUNSD, XFUND, and CORD are annotated with both word and block locations. Combining coordinate information of text words and blocks, the segment-level annotation is generated for fine-tuning the ESP model.

**Pre-Training Pseudo Label Generation** The existing pre-training methods mainly focus on designing self-supervised tasks to boost performance, such as masked language models or word-patch alignment. We approach this by exploiting knowledge from pre-training datasets. The pipeline to generate pseudo-labels for pre-training EE and EL tasks of key-value pairs is described as follows:

First, the key library is generated by counting the occurrences of text contents. The number of occurrences more than 100 and 15 will be considered as key in DocBank and RVL-CDIP, respectively. Generally, the keys should not be too long or too short, thus the text, with more than 50 characters or less than 3, is skipped. Second, the key-value pairs are generated based on the spatial relationship. If a text appears in the key library, it will be considered as a key, and then we can find its corresponding values based on rules, such as the values will appear below or right of the keys empirically. The category of other text is assigned as “other”.

For the key-value pairs, the key/value attributes are actually the pseudo labels of the EE task, and the key-value relationships are the pseudo labels of the EL task.

The visualization of pseudo-label is shown in Fig. 4. For there are some errors, the EL task is only trained with matched key/value links and the others are masked in the pre-training.

## 5.2. Implementation Details

**Pre-Training** We pre-train the ESP models for 10 epochs with an initial learning rate of  $1.25e-4$ , which is decayed by a factor of 0.1 at the 5th and 8th epoch, respectively. The longer side of images is resized to 1024, with the aspect ratio kept. The OCR information, obtained from PaddleOCR [17], is adopted to generate pseudo labels, including key-value labeling and linking information with the algorithm mentioned in Sec. 5.1.

**Fine-Tuning** Keeping the input size the same as the pre-training stage, the batch size is set to 8 and the initial learning rate is  $1.25e-4$ . Since there are only 149 training samples in FUNSD, we train our model for 300 epochs, and the learning rate drops to  $1.25e-5$  and  $1.25e-6$  at the 200th and 250th epochs, respectively. CORD and XFUND are trained for 150 epochs, and the learning rate drops to  $1.25e-5$  and  $1.25e-6$  at the 100th and 130th epochs.

**Evaluation** Following [19], we aggregate segment features of the text into one entity if the entity has multiple segments on FUNSD, XFUND, and CORD. But when testing end-to-end on SIBR, model inference does not rely on any annotation information. After decoding, the boxes, categories, and links are directly matched with GT. We use F1-score as our evaluation metrics for both EE and EL tasks.

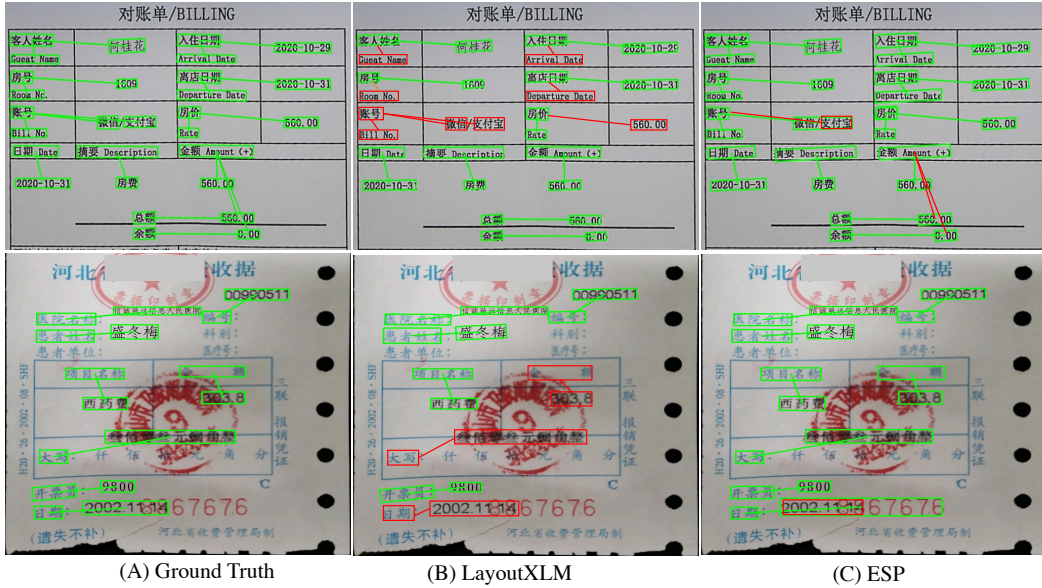


Figure 5. Qualitative results of ESP on the proposed dataset. The columns from left to right represent the visualization of ground truth, LayoutXML, and ESP. The rectangles in green and red stand for the results of correct and incorrect information extraction, respectively. The LayoutXML is more error-prone, due to the lower quality of OCR results. Only key-value pairs are shown for better visualization.

Task	Model	FUNSD	XFUND							
			ZH	JA	ES	FR	IT	DE	PT	Avg.
EE	RoBERTa [2]	66.7	87.7	77.6	61.1	67.4	66.9	68.1	68.2	71.0
	LayoutXML [36]	79.4	89.2	79.2	75.5	70.2	80.8	82.2	79.0	79.5
	LiLT [33]	84.2	89.4	79.6	79.1	79.5	83.7	82.3	82.2	82.3
	ESP	<b>91.1</b>	<b>90.3</b>	<b>81.1</b>	<b>85.4</b>	<b>90.5</b>	<b>88.9</b>	<b>87.2</b>	<b>87.5</b>	<b>87.3</b>
EL	RoBERTa [2]	26.6	51.1	58.0	52.9	49.7	53.1	50.4	39.8	50.7
	LayoutXML [36]	54.8	70.7	69.6	68.9	63.5	64.2	65.5	57.2	65.6
	LiLT [33]	62.76	72.9	70.4	71.9	69.7	70.4	65.6	58.7	68.5
	ESP	<b>88.6</b>	<b>90.8</b>	<b>88.3</b>	<b>85.2</b>	<b>90.9</b>	<b>90.0</b>	<b>85.2</b>	<b>86.2</b>	<b>88.1</b>

Table 3. Evaluation on FUNSD and XFUND (fine-tuning on X, testing on X). Note that, LayoutXML adopts language-specific fine-tuning and LiLT applies a plain text model (InfoXML). However, ESP does not involve any multilingual language models as EITM is frozen during the fine-tuning stage.

### 5.3. Comparison with the State-of-the-Arts

In this section, we compare ESP with LayoutLM-based models [11, 35–37], BERT [3], RoBERTa [2], XLayout [7], StrucText [19], and BROS [10] on three public benchmarks, and the results are summarized in Tab. 2.

We observe that ESP outperforms all existing methods on the EL task, and it boosts the upper bound of FUNSD and XFNSD to new record (88.88% and 92.31%). This proves that visual and geometric features can largely benefit entity linking prediction. On the EE task, ESP also achieves competitive performance, but its performance is slightly lower than that of LayoutLMv3 on FUNSD and CORD. Note that during evaluation LayoutLMv3 additionally used the ground truth of text contents, while ESP did not, so the comparison is not totally fair. Moreover, the model size (50M parameters) and the image set for pre-training (0.9M samples) of ESP are much smaller, compared with LayoutLMv3 (over 368M parameters and 11M samples).

We also conduct experiments on XFUND and FUNSD to verify the importance of visual features, and the results

Tag	Task				FUNSD		
	EAP	EITM	EE	EL	EE	EL	Box
(a)	✓	✓	✓	✓	<b>91.12</b>	<b>88.89</b>	<b>78.00</b>
(b)		✓	✓	✓	89.86	88.49	70.07
(c)	✓	✓		✓	88.60	87.97	77.85
(d)	✓		✓	✓	88.67	88.44	76.87
(e)	✓	✓	✓		89.13	61.20	77.05
(f)	✓				85.10	-	73.65
(g)					85.04	60.45	66.09

Table 4. Ablation study for pre-training tasks. The segment-level boxes are generated according to the locations of the words. However, the results are not high due to inaccurate word locations.

are shown in Tab. 3. Previous LayoutLM and its variants extract image feature by additional un-trainable visual encoders, leading to inevitably losses of OCR related details. We argue that fine-grained visual cues, such as fonts, colors, positions and typographic information, are crucial for VIE. We compare ESP with three multi-language approaches, LayoutXML, RoBERTa, and LiLT. We pre-train ESP on English datasets of DocBank and RVL-CDIP, but fine-tune and test it on language-specifically FUNSD (English) and XFUND (7 other languages). Note that the EITM task is not trained in the fine-tuning stage, meaning that the model does not perceive the other 7 kinds of languages. We find that ESP maintains at a plateau of 87.3% on XUFNSD, much better than language-independent LiLT, which demonstrates the effectiveness of visual information extractors.

### 5.4. Ablation Study

**Effect of Pre-training Tasks** This group of experiments is conducted on FUNSD to illustrate the impact of various

Tag	Task	Model	EE	EL
(a)	EE/EL	TRIE [39]	85.62	-
		LayoutXLM [36]	94.72	83.99
		ESP <sup>†</sup>	94.93	85.87
		ESP	<b>95.27</b>	<b>85.96</b>
(b)	Det + EE/EL	TRIE [39]	46.40	-
		LayoutXLM [36]	75.07	53.63
		ESP <sup>†</sup>	91.69	73.22
		ESP	<b>92.29</b>	<b>74.44</b>
(c)	Det + Rec + EE/EL	TRIE [39]	-	-
		LayoutXLM [36]	68.55	46.71
		ESP <sup>†</sup>	67.76	48.17
		ESP	<b>70.45</b>	<b>51.47</b>

Table 5. Evaluations and ablation studies on SIBR. “Det” and “Rec” denote “detection” and “recognition”, respectively. The recognition of text is considered correct when the normalized edit distance is less than 0.3. The † denotes a co-training process with text recognition branch.

combinations of the pre-training tasks, and the results are concluded in Tab. 4. By comparing (a) with each of (b) to (e), we find that every pre-training task benefits the proposed ESP. Specifically, (c) and (e) indicate that EE and EL pre-training tasks benefit corresponding downstream tasks most, resulting in an increase of 2.52% and 27.69%, respectively. (d) shows that the EITM module indeed learns effective vision-language representation and improves EE and EL by 2.45% and 0.45%. By comparing (f) with (g) and (a), we find that the benefits of pre-training tasks far exceed pre-trained data (6.02% vs 0.06%). We also experimentally prove that the last three semantic pre-training tasks can boost entity detection. By comparing (f) and (a), we find that the addition of EITM, EE, and EL tasks can improve the accuracy of detection by 4.35%. It demonstrates that the detection and semantic tasks can be mutually reinforced.

**Evaluations under Different Restrictions** We also evaluate the results under different restrictions, which are compliant with various input combinations in testing: (a) ground truth of text boxes and contents, (b) predicted text boxes and ground truth of text contents, (c) predicted text boxes and contents. Here, we re-implement LayoutXLM [36] and TRIE [39] and compare them with our ESP on SIBR. LayoutXLM is more suitable for SIBR as it is the multi-lingual variant of LayoutLMv2. TRIE and ESP are both trained in an end-to-end manner. As for LayoutXLM, it is trained with OCR-matched truth and evaluated with ground truth. In order to conduct experiments on the third condition (c), we use Duguang OCR engine<sup>1</sup> to get character-level locations and contents which are not provided by PaddleOCR. We closely follow the experimental setup of TRIE<sup>2</sup> and LayoutXLM<sup>3</sup>. The ablation studies under different restrictions

<sup>1</sup><https://duguang.aliyun.com/>

<sup>2</sup><https://github.com/hikopendata/DAVAR-Lab-OCR>

<sup>3</sup><https://github.com/microsoft/unilm>

are shown in Tab. 5.

(a) *With GT Input.* ESP achieves SOTA results on both EE and EL tasks, outperforming the other two methods by 0.55% and 1.97%, respectively. This proves that ESP can effectively handle the challenges in SIBR.

(b) *With Predicted Boxes.* In the second comparison, the performance of ESP exceeds the other two methods by a large margin. Specifically, taking the detection errors of PaddleOCR into consideration, the performance of LayoutXLM decreases by 19.65% in EE and 30.36% in EL. The performance of TRIE decreases by a larger margin. However, ESP only decreases by 2.98% and 11.52%, respectively, demonstrating its robustness to detection error.

(c) *With Predicted Boxes and Text.* Taking the recognition results into consideration, two of the three methods’ F1-scores decrease by a significant margin (>20%). Since TRIE gives an egregious result, we leave the corresponding result blank. We have several conclusions in this experiment: (1) The advantages of ESP become smaller with predicted text contents, (2) The text recognition branch is not easily trained with small data, (3) The connection between OCR and VIE should be studied more. The ultimate goal of VIE is to get not only the correct entity labels but also the correct contents. A qualitative comparison is shown in Fig. 5. Compared with LayoutXLM, ESP can better handle challenges such as overlap and printing shift.

**Effect of Text Recognition** We conduct the following experiments to discuss the effectiveness of the optional text recognition branch, and the results are presented in Tab. 5. ESP<sup>†</sup> refers to a co-training process of ESP with a text recognition branch. It can be observed from the first comparison that the result of entities slightly decreases by 0.34% with a recognition branch. The reason for the drop is the insufficient fine-tuning of text recognition on small data. Besides, we presume that the pre-training tasks have learned a better vision-language representation, which means explicit text information is not essential to the proposed ESP.

## 6. Conclusion and Future Work

We have released an in-the-wild document dataset for VIE, in which the images are challenging in their diverse appearances and complex structures. To deal with these difficulties, we further propose a unified framework to model entities as semantic points, in which entity spotting, labeling, and linking can be learned in an end-to-end manner. We also introduce three pre-training tasks to enforce the network to learn a better vision-language representation. Extensive experiments demonstrate that the proposed method can achieve significantly enhanced performance. The proposed ESP can also be extended to more general VIE problems, which we will leave for future research.

**Acknowledgement.** This work was supported by the National Natural Science Foundation of China 62225603.



## References

- [1] Srikar Appalaraju, Bhavan Jasani, Bhargava Urala Kota, Yusheng Xie, and R. Manmatha. Docformer: End-to-end transformer for document understanding. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 973–983, 2021. [2](#), [6](#)
- [2] Yiming Cui, Wanxiang Che, Ting Liu, Bing Qin, Shijin Wang, and Guoping Hu. Revisiting pre-trained models for Chinese natural language processing. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 657–668, Online, 2020. Association for Computational Linguistics. [6](#), [7](#)
- [3] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, 2019. Association for Computational Linguistics. [7](#)
- [4] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 6568–6577. IEEE, 2019. [2](#)
- [5] Alex Graves, Santiago Fernández, Faustino J. Gomez, and Jürgen Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In William W. Cohen and Andrew W. Moore, editors, *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006)*, Pittsburgh, Pennsylvania, USA, June 25-29, 2006, volume 148 of *ACM International Conference Proceeding Series*, pages 369–376. ACM, 2006. [4](#)
- [6] Jiuxiang Gu, Jason Kuen, Vlad I Morariu, Handong Zhao, Rajiv Jain, Nikolaos Barmpalios, Ani Nenkova, and Tong Sun. Unidoc: Unified pretraining framework for document understanding. *Advances in Neural Information Processing Systems*, 34:39–50, 2021. [6](#)
- [7] Zhangxuan Gu, Changhua Meng, Ke Wang, Jun Lan, Weiqiang Wang, Ming Gu, and Liqing Zhang. 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, 2022. [6](#), [7](#)
- [8] Geoffrey Hinton. How to represent part-whole hierarchies in a neural network. *ArXiv preprint*, abs/2102.12627, 2021. [2](#)
- [9] Geoffrey E Hinton. Mapping part-whole hierarchies into connectionist networks. *Artificial Intelligence*, 46(1-2):47–75, 1990. [2](#)
- [10] Teakgyu Hong, Donghyun Kim, Mingi Ji, Wonseok Hwang, Daehyun Nam, and Sungrae Park. 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, 2022. [2](#), [6](#), [7](#)
- [11] Yupan Huang, Tengchao Lv, Lei Cui, Yutong Lu, and Furu Wei. Layoutlmv3: Pre-training for document ai with unified text and image masking. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2022. [1](#), [2](#), [4](#), [6](#), [7](#)
- [12] Zheng Huang, Kai Chen, Jianhua He, Xiang Bai, Dimosthenis Karatzas, Shijian Lu, and CV Jawahar. Icdar2019 competition on scanned receipt ocr and information extraction. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 1516–1520. IEEE, 2019. [3](#)
- [13] Guillaume Jaume, Hazim Kemal Ekenel, and Jean-Philippe Thiran. 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, 2019. [3](#), [5](#)
- [14] Wonjae Kim, Bokyoung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 5583–5594. PMLR, 2021. [2](#)
- [15] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European conference on computer vision (ECCV)*, pages 734–750, 2018. [2](#)
- [16] Chenliang Li, Bin Bi, Ming Yan, Wei Wang, Songfang Huang, Fei Huang, and Luo Si. 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, Online, 2021. Association for Computational Linguistics. [1](#)
- [17] Chenxia Li, Weiwei Liu, Ruoyu Guo, Xiaoyue Yin, Kaitao Jiang, Yongkun Du, Yuning Du, Lingfeng Zhu, Baohua Lai, Xiaoguang Hu, Dianhai Yu, and Yanjun Ma. Pp-ocrv3: More attempts for the improvement of ultra lightweight ocr system. *ArXiv preprint*, abs/2206.03001, 2022. [6](#)
- [18] Peizhao Li, Jiuxiang Gu, Jason Kuen, Vlad I Morariu, Handong Zhao, Rajiv Jain, Varun Manjunatha, and Hongfu Liu. Selfdoc: Self-supervised document representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5652–5660, 2021. [2](#), [6](#)
- [19] Yulin Li, Yuxi Qian, Yuechen Yu, Xiameng Qin, Chengquan Zhang, Yan Liu, Kun Yao, Junyu Han, Jingtuo Liu, and Errui Ding. Structext: Structured text understanding with multi-modal transformers. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 1912–1920, 2021. [1](#), [2](#), [5](#), [6](#), [7](#)
- [20] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 936–944. IEEE Computer Society, 2017. [4](#)
- [21] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *IEEE International Conference on Computer Vision, ICCV*

- 2017, Venice, Italy, October 22-29, 2017, pages 2999–3007. IEEE Computer Society, 2017. 5
- [22] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022. 3
- [23] Shangbang Long, Xin He, and Cong Yao. Scene text detection and recognition: The deep learning era. *International Journal of Computer Vision*, 129:161–184, 2018. 2
- [24] Alejandro Newell, Zhiao Huang, and Jia Deng. Associative embedding: End-to-end learning for joint detection and grouping. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 2277–2287, 2017. 2
- [25] Lawrence O’Gorman. The document spectrum for page layout analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 15(11):1162–1173, 1993. 1
- [26] Seunghyun Park, Seung Shin, Bado Lee, Junyeop Lee, Jaeheung Surh, Minjoon Seo, and Hwalsuk Lee. Cord: a consolidated receipt dataset for post-ocr parsing. In *Workshop on Document Intelligence at NeurIPS 2019*, 2019. 3, 5
- [27] Baoguang Shi, Xiang Bai, and Serge J. Belongie. Detecting oriented text in natural images by linking segments. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 3482–3490. IEEE Computer Society, 2017. 2
- [28] Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39:2298–2304, 2015. 1
- [29] Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *IEEE transactions on pattern analysis and machine intelligence*, 39(11):2298–2304, 2016. 4
- [30] Ray Smith. An overview of the tesseract ocr engine. In *Ninth international conference on document analysis and recognition (ICDAR 2007)*, volume 2, pages 629–633. IEEE, 2007. 1
- [31] Sibong Song, Jianqiang Wan, Zhibo Yang, Jun Tang, Wenqing Cheng, Xiang Bai, and Cong Yao. Vision-language pre-training for boosting scene text detectors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15681–15691, 2022. 2
- [32] Jun Tang, Zhibo Yang, Yongpan Wang, Qi Zheng, Yongchao Xu, and Xiang Bai. Seglink++: Detecting dense and arbitrary-shaped scene text by instance-aware component grouping. *Pattern Recognition*, 96:106954, 2019. 2
- [33] Jiapeng Wang, Lianwen Jin, and Kai Ding. Lilt: A simple yet effective language-independent layout transformer for structured document understanding. *ArXiv preprint*, abs/2202.13669, 2022. 1, 2, 6, 7
- [34] Jiapeng Wang, Chongyu Liu, Lianwen Jin, Guozhi Tang, Jiaxin Zhang, Shuaitao Zhang, Qianying Wang, Yaqiang Wu, and Mingxiang Cai. Towards robust visual information extraction in real world: new dataset and novel solution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 2738–2745, 2021. 1, 2, 3, 5
- [35] Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. Layoutlm: Pre-training of text and layout for document image understanding. In Rajesh Gupta, Yan Liu, Jiliang Tang, and B. Aditya Prakash, editors, *KDD ’20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*, pages 1192–1200. ACM, 2020. 1, 2, 4, 6, 7
- [36] Yiheng Xu, Tengchao Lv, Lei Cui, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, and Furu Wei. Layoutxlm: Multimodal pre-training for multilingual visually-rich document understanding. *ArXiv preprint*, abs/2104.08836, 2021. 4, 7, 8
- [37] Yang Xu, Yiheng Xu, Tengchao Lv, Lei Cui, Furu Wei, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, Wanxiang Che, Min Zhang, and Lidong Zhou. 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, Online, 2021. Association for Computational Linguistics. 1, 2, 6, 7
- [38] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S. Huang. Free-form image inpainting with gated convolution. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 4470–4479. IEEE, 2019. 3
- [39] Peng Zhang, Yunlu Xu, Zhazhan Cheng, Shiliang Pu, Jing Lu, Liang Qiao, Yi Niu, and Fei Wu. TRIE: end-to-end text reading and information extraction for document understanding. In *MM ’20: The 28th ACM International Conference on Multimedia, Virtual Event / Seattle, WA, USA, October 12-16, 2020*, pages 1413–1422, 2020. 1, 3, 8
- [40] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *ArXiv preprint*, abs/1904.07850, 2019. 2, 3, 4
- [41] Xinyu Zhou, Cong Yao, He Wen, Yuzhi Wang, Shuchang Zhou, Weiran He, and Jiajun Liang. East: An efficient and accurate scene text detector. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2642–2651, 2017. 1
- [42] Xingyi Zhou, Jiacheng Zhuo, and Philipp Krahenbuhl. Bottom-up object detection by grouping extreme and center points. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 850–859, 2019. 2
- [43] Yingying Zhu, Cong Yao, and Xiang Bai. Scene text detection and recognition: recent advances and future trends. *Frontiers of Computer Science*, 10:19–36, 2016. 2