

TRIE: End-to-End Text Reading and Information Extraction for Document Understanding

Peng Zhang^{1*}, Yunlu Xu^{1*}, Zhanzhan Cheng^{21†}, Shiliang Pu¹⁺, Jing Lu¹, Liang Qiao¹, Yi Niu¹, Fei Wu²

¹Hikvision Research Institute, Hangzhou, China

(zhangpeng23,xuyunlu,pushiliang,hri,lujing6,qiaoliang6,niuyl)@hikvision.com

²Zhejiang University, Hangzhou, China

11821104@zju.edu.cn, wufei@cs.zju.edu.cn

ABSTRACT

Since real-world ubiquitous documents (*e.g.*, invoices, tickets, resumes and leaflets) contain rich information, automatic document image understanding has become a hot topic. Most existing works decouple the problem into two separate tasks, (1) *text reading* for detecting and recognizing texts in images and (2) *information extraction* for analyzing and extracting key elements from previously extracted plain text. However, they mainly focus on improving *information extraction* task, while neglecting the fact that *text reading* and *information extraction* are mutually correlated. In this paper, we propose a unified *end-to-end text reading and information extraction network*, where the two tasks can reinforce each other. Specifically, the multimodal visual and textual features of *text reading* are fused for *information extraction* and in turn, the semantics in *information extraction* contribute to the optimization of *text reading*. On three real-world datasets with diverse document images (from fixed layout to variable layout, from structured text to semi-structured text), our proposed method significantly outperforms the state-of-the-art methods in both efficiency and accuracy.

CCS CONCEPTS

• Applied computing → Document analysis; Optical character recognition; • Computing methodologies → Information extraction.

KEYWORDS

End-to-End; Text Reading; Information Extraction; Visually Rich Documents

ACM Reference Format:

Peng Zhang^{1*}, Yunlu Xu^{1*}, Zhanzhan Cheng^{21†}, Shiliang Pu¹⁺, Jing Lu¹, Liang Qiao¹, Yi Niu¹, Fei Wu². 2020. TRIE: End-to-End Text Reading and Information Extraction for Document Understanding. In *Proceedings of the*

*Both authors contributed equally to this research.

†This work is completed under the supervision of Zhanzhan Cheng (contact email: 11821104@zju.edu.cn).

+Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '20, October 12–16, 2020, Seattle, WA, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7988-5/20/10... \$15.00

<https://doi.org/10.1145/3394171.3413900>

28th ACM International Conference on Multimedia (MM '20), October 12–16, 2020, Seattle, WA, USA. ACM, New York, NY, USA, 10 pages.
<https://doi.org/10.1145/3394171.3413900>

Table 1: Example of VRDs and their categories.

Layout	Text type	Structured	Semi-structured
	Fixed	invoice, passport, ID card	business email, sales contract
Variable		purchase receipt, business card	resume, financial report

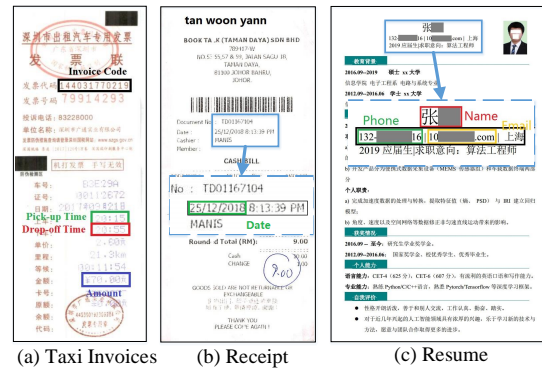


Figure 1: Examples of VRDs and example entities to extract. Sensitive information like name and contact information in resumes are pixelated to preserve privacy. Best viewed in color.

1 INTRODUCTION

Document understanding is a relatively traditional research topic that refers to the techniques to automatically handle with the text content. Among different types of documents, Visually Rich Documents (VRDs) have attracted more and more attention. It is named after its composed modality of both text and vision that offers abundant information, including but not limited to layout, tabular structure, font size and even the font color in addition to plain text. Here, we divide them into *four categories* from the dimensions of layout and text

type. *Layout* here is defined as the relative positions of texts and *text type* (i.e., structured, and semi-structured) follows the conventions of [22, 51]. In detail, *Structured* data connotes data that is organized in a predetermined schema (e.g., invoices and receipt), while *semi-structured* data denotes data that has one or more identifiers, but each portion of data is not necessarily organized in predefined fields, e.g., resumes. Table 1 summarizes common examples of VRDs and some examples are shown in Fig. 1. Automatically recognizing texts and extracting valuable contents from VRDs can facilitate information entry, retrieval and compliance check. It is of great benefit to office automation in areas like accounting, financial and much broader real-world applications.

The complex problem of understanding VRDs is always decoupled into two different stages, text reading and information extraction. Specifically, **text reading** includes text detection and recognition in images, which belongs to the optical character recognition (OCR) research field and has already been widely used in many Computer Vision (CV) applications [12, 42, 52]. In the **information extraction** (IE) stage, specific contents (e.g. entity, relation) are mined and processed from previously recognized plain text for diverse Natural Language Processing (NLP) tasks, such as Named Entity Recognition (NER) [24, 36, 48] and Question-Answer (QA) [1, 13, 57]. Note that in existing routines, the two stages are separately executed. It means that the former recognizes text from images without semantic supervision (e.g., entity name annotation) of IE stage, and the latter extracts information from only serialized plain text, not using rich visual information.

Under the traditional routines, most existing methods [9, 23, 30, 39, 40, 45, 54, 61] design frameworks in multiple stages of text reading (usually including detection and recognition) and information extraction independently in order. In earlier explorations, the task is intrinsically downgraded into a traditional OCR procedure and the downstream IE from serialized plain text [40, 45], which completely discards the visual features and layout information from images, as shown in Fig. 2(a). Noticing the rich visual information contained in document images, recent works incorporate them into network design. [9, 23, 30, 39, 61] work on recognized texts and their positions (see Fig. 2(b)), while [54, 58] further includes image embeddings (see Fig. 2(c)). They all focus on the design of information extraction task in VRD understanding.

In essence, all the above works inevitably have the following three limitations. (1) VRD understanding requires both visual and textual features, but the visual features they exploited are limited. (2) *Text reading* and *information extraction* are highly correlated, but the relations between them have rarely been explored. (3) The stagewise training strategy of *text reading* and *information extraction* brings redundant computation and time cost.

To address those limitations, in this paper, we propose a novel end-to-end *Text Reading and Information Extraction* (TRIE) network. The workflow is as shown in Fig. 2(d). Instead of just focusing on information extraction, we bridge *text reading* and *information extraction* tasks with shared features, including position features, visual features and textual features. Thus, the two tasks can reinforce each other amidst a unified framework. Specifically, in the forward pass, multimodal visual and textual features produced by *text reading* are fully fused for *information extraction*, while in the backward pass, the semantics in *information extraction* also contribute to the

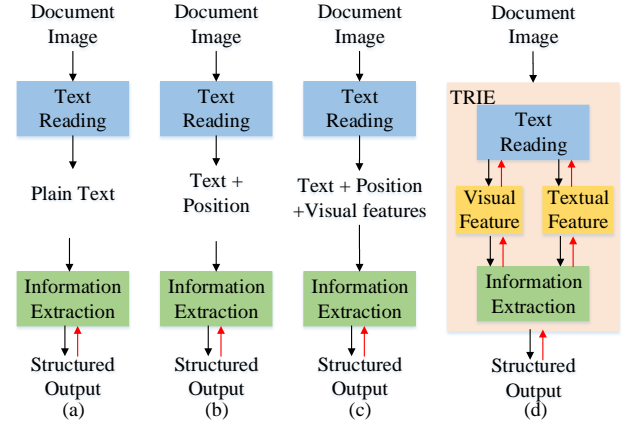


Figure 2: Comparison of VRD understanding architectures: (a) IE with plain text, (b) IE with text and position, (c) IE with position and visual features in addition to plain text, (d) our proposed TRIE. The black and red arrows mean the forward and backward processing, respectively. Best viewed in color.

optimization of *text reading*. Since all the modules in the network are differentiable, the whole network can be trained end-to-end. To the best of our knowledge, this is the first end-to-end trainable VRD framework, which can handle various types of VRDs, from fixed to variable layouts and from structured to semi-structured text type.

Major contributions are summarized as follows:

- (1) We propose an end-to-end trainable framework for simultaneous text reading and information extraction in VRD understanding. The whole framework can be trained end-to-end from scratch, with no need of stagewise training strategies.
- (2) We design a multimodal context block to bridge the OCR and IE modules. To the best of our knowledge, it is the first work to mine the mutual influence of text reading and information extraction.
- (3) We perform extensive evaluations on our framework and show superior performance compared with the state-of-the-art counterparts both in efficiency and accuracy on three real-world benchmarks. Note that those three benchmarks cover diverse types of document images, from fixed to variable layouts, from structured to semi-structured text types.

2 RELATED WORKS

Here, we briefly review the recent advances in text reading and information extraction.

2.1 Text Reading

Text reading are formally divided into two sub-tasks: text detection and text recognition. In text detection, methods are usually divided into two categories: *anchor-based methods* and *segmentation-based methods*. *Anchor-based methods* [2, 17, 28, 32, 49] predict the existence of texts and regress their location offsets at pre-defined grid points of the input image, while *segmentation-based methods* [33, 53, 62] learn the pixel-level classification tasks to separate

text regions apart from the background. In text recognition, the mainstreaming CRNN framework was introduced by [50], using recurrent neural networks (RNNs) [6, 19] combined with CNN-based methods for better sequential recognition of text lines. Then, the attention mechanism replaced existing CTC decoder [50] and was applied to a stacked RNN on top of the recursive CNN [25], whose performance surpassed the state-of-the-art among diverse variations.

To sufficiently exploit the complementary between detection and recognition, [3, 12, 18, 26, 31, 35, 42, 52] were proposed to jointly detect and recognize text instances in an end-to-end manner. They all achieved impressive results compared to traditional pipeline approaches due to capturing *relations* within detection and recognition sub-tasks through joint training. This inspires us to dive into the broader field and for the first time pay attention to the relations between text reading and information extraction.

2.2 Information Extraction

Information extraction has been studied for decades. Before the advent of learning based models, rule-based methods [8, 11, 21, 38, 44, 47] were proposed. Pattern matching were widely used in [21, 44] to extract one or multiple target values. Afterwards, Intellix [47] required predefined template with relevant fields annotated and SmartFix [8] employed specifically designed configuration rules for each template. Though rule-based methods work in some cases, they rely heavily on the predefined rules, whose design and maintenance usually require deep expertise and large time cost. Besides, they can not generalize across document templates.

Learning-based networks [7, 10, 24, 36, 55, 56] were firstly proposed to work on plain text documents. [40, 45] adopted the idea from natural language processing and used recurrent neural networks to extract entities of interest from VRDs. However, they discard the layout information during the text serialization, which is crucial for document understanding. Recently, observing the rich visual information contained in document images, works tended to incorporate more details from VRDs. Some works [9, 23, 39, 61] took the layout into consideration, and worked on the reconstructed character or word segmentation of the document. Position information had also been utilized. [30] combined texts and their positions through a Graph Convolutional Network (GCN) and [54] further integrated position and image embeddings for pre-training inspired by BERT [10]. [37] presented a representation learning approach to extract structured information from templatic documents, which worked in pipeline of candidate generation, scoring and assignment. However, these works all focus on the design of information extraction task. They miss lots of informative details because multi-modality of inputs are not fully explored.

Two related concurrent works were presented in [4, 14]. [14] proposed an entity-aware attention text extraction network to extract entities from VRDs. However, it could only process documents of relatively fixed layout and structured text, like train tickets, passports and bussiness cards. [4] localized, recognized and classified each word in the document. Since it worked in the word granularity, it not only required much more labeling efforts (positions, content and category of each word) but also had difficulties in extracting those entities which were embedded in word texts (e.g. extracting ‘51xxxx@xxx.com’ from ‘153-xxx97|51xxxx@xxx.com’). Besides,

in its entity recognition branch, it still worked on the serialized word features, which were sorted and packed in the left to right and top to bottom order. Thus, it can only handle documents of simple layout and non-structured text, like paragraph pages. Both of the two existing works are strictly limited to documents of relative fixed layout and one type of text (structured or semi-structured). While our proposed framework has the ability to handle documents of both fixed and variable layouts, structured and semi-structured text types.

3 METHODOLOGY

We first introduce the overall architecture of TRIE in Sec 3.1. In the following three sections, text reading module, multimodal context block and information extraction module are illustrated in details respectively. Optimization functions are as shown in Sec 3.5.

3.1 Overall Architecture

An overview of the architecture is as shown in Fig. 3. It mainly consists of three parts: text reading module, multimodal context block and information extraction module. *Text reading* module is responsible for localizing and recognizing all texts in document images and *information extraction* module is to extract entities of interest from them. The *multimodal context block* is novelly designed to bridge the text reading and information extraction modules.

Specifically, in text reading, the network takes the original image as input and outputs text region coordinate positions. Once the positions obtained, we apply RoIAlign [15] on the shared convolutional features to get text region features. Then an attention-based sequence recognizer is adopted to get textual features for each text. Since context information of a text is necessary to tell it apart from other entities, we design a special *multimodal context block* to provide both visual and textual context. In the block, the visual context features summarize the local patterns of a text while textural context features model implicit relations among the whole document image. These two context features are complementary to each other and fused as final context features. Finally, entities are extracted through the combined context and textual features. The whole framework learns to read text and extract entities jointly and can be trained in an end-to-end way from the scratch, which allows bi-directional information flows and brings reinforcing effects among text reading and information extraction modules.

3.2 Text Reading Module

Text reading module commonly includes a shared backbone, a text detection branch as well as a sequential recognition branch. We adopt ResNet [16] and Feature Pyramid Network (FPN) [29] as our backbone to extract shared convolutional features. For an input image x , we denote $\mathcal{I} \in \mathbb{R}^{h \times w \times d}$ as the shared feature maps, where h , w and d are the height, width and channel number of \mathcal{I} . The text detection branch takes \mathcal{I} as input and predicts the locations of all possible text regions.

$$\mathcal{B} = \text{Detector}(\mathcal{I}) \quad (1)$$

where the *Detector* can be any anchor-based [17, 28, 32, 49] or segmentation-based [33, 53, 62] text detection methods. Here, $\mathcal{B} = (b_1, b_2, \dots, b_m)$ is a set of m text bounding boxes and $b_i = (x_0, y_0, x_1, y_1)$ denotes the top-left and bottom-right positions of the i -th

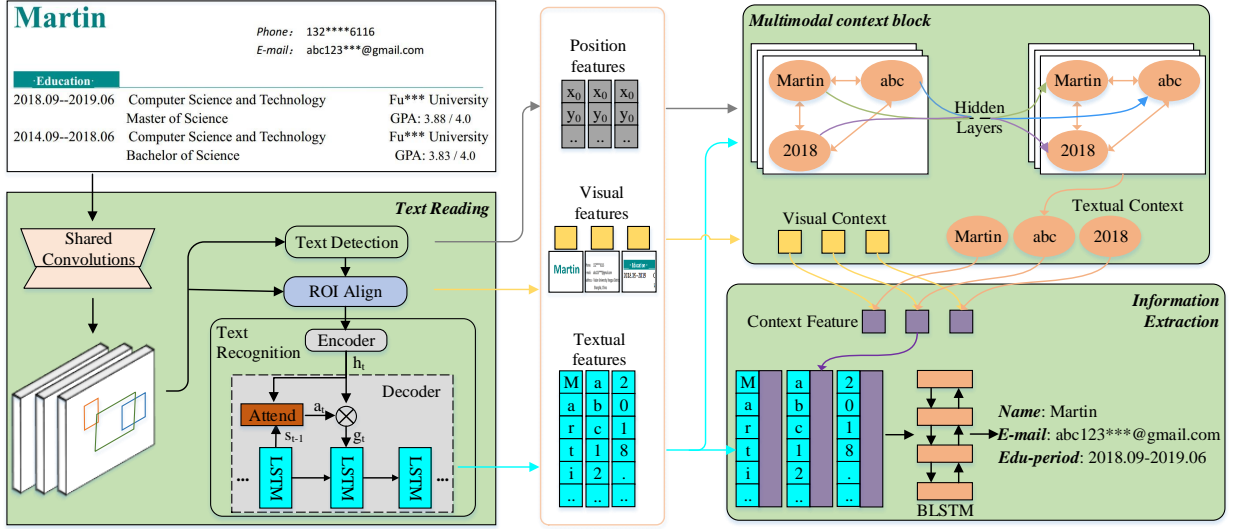


Figure 3: Overall architecture. The network predicts text regions, text content and extract entities of interest in a single forward pass.

text. Given text positions \mathcal{B} , RoIAlign [15] is applied on the shared convolutional features \mathcal{I} to get their text region features, denoted as $C = (c_1, c_2, \dots, c_m)$, where $c_i \in \mathbb{R}^{h' \times w' \times d}$, h' and w' are the spatial dimensions, and d is the vector dimension the same as in \mathcal{I} .

Afterwards, the text recognition branch predicts a character sequence from each text region features c_i . Firstly, c_i is fed into an encoder to extract a higher-level feature sequence $\mathcal{F} \in \mathbb{R}^{l \times d_r}$, where l is the length of the sequence and d_r is vector dimension. Then, attention-based decoder is adopted to recurrently generate the sequence of characters $y = (y_1, y_2, \dots, y_T)$, where T is the length of label sequence. Specifically, at step t , the attention weights α_t and glimpse vector g_t are computed as follows,

$$g_t = \sum_{k=1}^l \alpha_{t,k} \mathcal{F}_k \quad (2)$$

$$\alpha_{t,k} = \exp(e_{t,k}) / \sum_{j=1}^l \exp(e_{t,j}) \quad (3)$$

where,

$$e_{t,j} = w^\top \tanh(Ws_{t-1} + M\mathcal{F}_j + b) \quad (4)$$

and w , W , M and b are trainable weights.

The state s_{t-1} is updated via,

$$s_t = RNN(s_{t-1}, g_t, y_{t-1}) \quad (5)$$

where y_{t-1} is the $(t-1)$ -th ground-truth label in training, while in testing, it is the label predicted in the previous step. In our experiment, LSTM [19] is adopted as RNN unit. Finally, the probability distribution over the vocabulary label space is estimated by,

$$p_t^{rcg} = \text{softmax}(W^{rcg}s_t + b^{rcg}) \quad (6)$$

where W^{rcg} and b^{rcg} are learnable weights. We denote that $z_i = (s_1, s_2, \dots, s_T)$ as the textual features of i -th text, as they contain semantic features for each character in it.

In a nutshell, the text reading module outputs visual features $C = (c_1, c_2, \dots, c_m)$ and textual features $\mathcal{Z} = (z_1, z_2, \dots, z_m)$ of m texts in the image in addition to their positions $\mathcal{B} = (b_1, b_2, \dots, b_m)$.

3.3 Multimodal Context Block

The context of a text provides necessary information to tell it apart from other entities, which is crucial for information extraction. Unlike the most existing works which rely only on textual features and/or position features, we design a multimodal context block to consider position features, visual features and textual features all together. This block provides both visual context and textual context of a text, which are complementary to each other and further fused in the information extraction module.

3.3.1 Visual Context. As mentioned, visual details such as the obvious color, font, layout and other informative features are equally important as textual details (text content) for document understanding. A natural way of capturing the local visual context of a text is resort to the convolutional neural network. Different from [54] which extracts these features from scratch, we directly reuse $C = (c_1, c_2, \dots, c_m)$ produced by the text reading module. Thanks to the deep backbone and lateral connections introduced by FPN, each c_i summarizes the rich local visual patterns of the i -th text.

3.3.2 Textual Context. Unlike visual context which focuses on local visual patterns, textual context models the fine-grained long distance dependencies and relationships between texts, providing complementary context information. Inspired by [10, 27, 34], we apply the self-attention mechanism to extract textual context features, supporting variable number of texts.

Self-attention Recap. The input of popular scaled dot-product attention consists of queries Q and keys K of dimension d_k , and values V of dimension d_v . The output is obtained by weighted summation over all values V and the attention weights are learned from Q and K . Please refer to [10] for details.

Textual Context Modeling. To retain the document layout and content as much as possible, we make the use of positions $\mathcal{B} = (b_1, b_2, \dots, b_m)$ and textual features $\mathcal{Z} = (z_1, z_2, \dots, z_m)$ at the same time. We first introduce *position embeddings* to preserve layout information. Then, a *ConvNet* of multiple kernels similar to [60] is used to aggregate semantic character features in z_i and outputs \widehat{z}_i ,

$$\widehat{z}_i = \text{ConvNet}(z_i). \quad (7)$$

Finally, the i -th text’s embedding is fused of \widehat{z}_i and position embedding, followed by the *LayerNorm* normalization, defined as

$$\text{embedding}_i = \text{LayerNorm}(\text{Linear}(\widehat{z}_i) + \text{PE}(b_i)) \quad (8)$$

where, *Linear* is a fully-connected layer, projecting \widehat{z}_i to the same dimension with $\text{PE}(b_i)$. b_i is the bounding box coordinates of i -th text (computed by Equa. 1) and PE is the position embedding layer. Afterwards, we pack all the texts’ embedding vector together, which serves as the K , Q and V in the scaled dot-product attention.

The textual context features $\widetilde{C} = (\widetilde{c}_1, \widetilde{c}_2, \dots, \widetilde{c}_m)$ is obtained by,

$$\begin{aligned} \widetilde{C} &= \text{Attention}(Q, K, V) \\ &= \text{softmax}\left(\frac{QK^T}{\sqrt{d_{info}}}\right)V \end{aligned} \quad (9)$$

where d_{info} is the dimension of text embeddings, and $\sqrt{d_{info}}$ is the scaling factor. To further improve the representation capacity of the attended feature, multi-head attention is introduced. Each head corresponds to an independent scaled dot-product attention function and the text context features \widetilde{C} is given by:

$$\begin{aligned} \widetilde{C} &= \text{MultiHead}(Q, K, V) \\ &= [\text{head}_1, \text{head}_2, \dots, \text{head}_n]W^{info} \end{aligned} \quad (10)$$

$$\text{head}_j = \text{Attention}(QW_j^Q, KW_j^K, VW_j^V) \quad (11)$$

where W_j^Q , W_j^K and $W_j^V \in \mathbb{R}^{(d_{info} \times d_n)}$ are the learned projections matrices for the j -th head, n is the number of heads, and $W^{info} \in \mathbb{R}^{(d_{info} \times d_{info})}$. To prevent the multi-head attention model from becoming too large, we usually have $d_n = \frac{d_{info}}{n}$.

3.4 Information Extraction Module

Both the context and textual features matter in entity extraction. The context features (including both visual context features C and textual context features \widetilde{C}) provide necessary information to tell entities apart while the textual features \mathcal{Z} enable entity extraction in the character granularity, as they contain semantic features for each character in the text. So we first perform multimodal fusion of visual context features C and textual context features \widetilde{C} , which are further combined with textual features \mathcal{Z} to extract entities.

3.4.1 Adaptive Multi-modality Context Fusion. Given the visual context features $C \in \mathbb{R}^{m \times h' \times w' \times d}$ and textual context features $\widetilde{C} \in \mathbb{R}^{m \times d_{info}}$, we fuse them adaptively to get the final context of each text. Specifically, for the i -th text, $c_i \in \mathbb{R}^{h' \times w' \times d}$ is spatially averaged to output the visual context vector. The final context vector \overline{c}_i is a weighted sum of visual context vector and textual context vector. That is, \overline{c}_i

$$\overline{c}_i = \alpha \cdot \text{Linear}\left(\frac{1}{h'w'} \sum_j c_{ij}\right) + \beta \cdot \text{Linear}(\widetilde{c}_i), \quad (12)$$

where α and β are automatically learnable weights. *Linear* is a fully-connected layer, projecting the visual context vector and text context vector to the same dimension.

3.4.2 Entity Extraction. We further combine the context vector \overline{c}_i and textual features z_i by concatenating \overline{c}_i to each s_j in z_i ,

$$U_i = (u_{i,1}, u_{i,2}, \dots, u_{i,T}), \text{ where } u_j = s_{i,j} \parallel \overline{c}_i \quad (13)$$

Then, a Bidirectional-LSTM is applied to further model the dependencies within the characters,

$$H_i = (h_{i,1}, h_{i,2}, \dots, h_{i,T}) = \text{BiLSTM}(U_i), \quad (14)$$

which is followed by a fully connected network, projecting the output to the dimension of IOB [46] label space.

$$p_{i,j}^{info} = \text{softmax}(\text{Linear}(h_{i,j})) \quad (15)$$

3.5 Optimization

The proposed network can be trained in an end-to-end manner and the losses are generated from three parts,

$$\mathcal{L} = \mathcal{L}_{det} + \lambda_{recog} \mathcal{L}_{recog} + \lambda_{info} \mathcal{L}_{info} \quad (16)$$

where hyper-parameters λ_{recog} and λ_{info} control the trade-off between losses.

\mathcal{L}_{det} is the loss of text detection branch, which consists of a classification loss and a regression loss, as defined in [43]. The recognition loss and information extraction loss are formulated as,

$$\mathcal{L}_{recg} = -\frac{1}{T} \sum_{i=1}^m \sum_{t=1}^T \log p_{i,t}^{rcg}(y_{i,t}^{rcg}) \quad (17)$$

$$\mathcal{L}_{info} = -\frac{1}{T} \sum_{i=1}^m \sum_{t=1}^T \log p_{i,t}^{info}(y_{i,t}^{info}) \quad (18)$$

where $y_{i,t}^{rcg}$ is the ground-truth label of t -th character in i -th text from recognition branch and $y_{i,t}^{info}$ is its corresponding label of information extraction.

Note that since *text reading* and *information extraction* modules are bridged with shared visual and textual features, they can reinforce each other. Specifically, the visual and textual features of text reading are fully fused and essential for information extraction, while the semantic feedback of information extraction also contributes to the optimization of the shared convolutions and text reading module.

4 EXPERIMENTS

In this section, we perform experiments to verify the effectiveness of the proposed method.

Table 2: Statistics of benchmark datasets used in this paper.

Dataset	Training	Testing	Entities	Layout	Text Type
Taxi Invoices	4000	1000	9	Fixed	Struct
SROIE	626	347	4	Variable	Struct
Resumes	1978	497	6	Variable	Semi-struct

4.1 Datasets

We validate our model on three real-world datasets. One is the public SROIE [20] benchmark, and the other two are self-built datasets, Taxi Invoices and Resumes, respectively. Note that, the three benchmarks differ largely in layout and text type, from fixed to variable layout and from structured to semi-structured text. Statistics of the datasets are listed in Table 2 and some examples are shown in Fig. 1.

- *Taxi Invoices* consists of 5000 images and has 9 entities to extract (Invoice Code, Invoice Number, Date, Pick-up time, Drop-off time, Price, Distance, Waiting, Amount). The invoices are in Chinese and can be grouped into roughly 13 templates. So it is kind of document of fixed layout and structured text type.
- *SROIE* [20] is a public dataset for receipt information extraction in ICDAR 2019 Challenge. It contains 626 receipts for training and 347 receipts for testing. Each receipt is labeled with four types of entities, which are Company, Date, Address and Total. It has variable layouts and structured text.
- *Resumes* is a dataset of 2475 Chinese scanned resumes, which has 6 entities to extract (Name, Phone Number, Email Address, Education period, Universities and Majors). As an owner can design his own resume template, this dataset has variable layouts and semi-structured text.

4.2 Implementation Details

4.2.1 Network Details: The backbone of our model is ResNet-50 [16], followed by the FPN [29] to further enhance features. The text detection branch in *text reading module* adopts the Faster R-CNN [43] network and outputs the predicted bounding boxes of possible texts for later sequential recognition. For each text region, its features of shape 32×256 are extracted from the shared convolutional features by RoIAlign [15] and decoded by LSTM-based attention [5], where the number of hidden units is set to 256. In the *information extraction module*, we use convolutions of three kernel size [3, 5, 7] followed by max pooling to extract representations of text. The dimension of text’s embedding vector is 256. The textual context modeling module consists of multiple layers of multi-head scaled dot-product attention modules. The number of hidden units of BiLSTM used in entity extraction is set to 128. Hyper-parameters λ_{recog} and λ_{info} in Equa. 16 are all empirically set to 1 in our experiments.

4.2.2 Training: Our model and all the reimplemented counterparts are implemented under the PyTorch framework [41]. For our model, with the ADADELTA [59] optimization method, we set the learning rate to 1.0 at the beginning and decreased it to a tenth every 40 epochs. The batch size is set to 2 images per GPU and we train our model for 150 epochs. For compared counterparts, we train the separate text reading and information extraction tasks independently

until they fully converge. All the experiments are carried out on a workstation with two NVIDIA Tesla V100 GPUs.

4.3 State-of-the-Art Comparisons

We compare our proposed method with pipeline of SOTA text reading and information extraction.

4.3.1 Network Settings of SOTA counterparts and Evaluation Metric. Most existing works only focus on information extraction task for document understanding. For fair comparisons, we train the *text reading* network which is identical to our model individually and take the results as the input of information extraction counterparts. We re-implement three top information extraction methods [23, 30, 36]. Their network settings are as follows:

- (1) **Chargrid(TR):** Pipeline of text reading and Chargrid. For Chargrid[23], the input character embedding is set to 128 and the rest of network is identical to the paper.
- (2) **NER(TR):** Pipeline of text reading and LSTM-CRF. The input embedding of LSTM-CRF[36] is set to 256, followed by a BiLSTM with 128 hidden units and CRF.
- (3) **GCN(TR):** Pipeline of text Reading and GCN. In GCN, we build the network as the paper [30] specified and set the input embedding to 256, which performs slightly better than 128.

In consistence with [30] and the SROIE Challenge [20], we use F1-score to evaluate the performance of all experiments.

4.3.2 Results. Evaluation on Taxi Invoices: In this dataset, the noise of low-quality and taint may lead to failures of detection and recognition of entities. Besides, the contents may be misplaced, e.g. the content of Pick-up time may appear after the ‘Date’. Table 3 shows the results. We see that our model outperforms counterparts by significant margins except for the Pick-up time (illustrated in the tail of the paragraph). Concretely, *NER(TR)* discards the layout information and serializes all texts into one-dimensional text sequences, reporting inferior performance than other methods. Benefiting from the layout information, *Chargrid(TR)* and *GCN(TR)* work much better. However, *Chargrid(TR)* conducts pixel segmentation task and is prone to omit characters or include extra characters. For *GCN(TR)*, it only exploits the positions of text segments. Obviously, our *TRIE* has the ability to boost performances by using more useful visual features in VRDs. In addition, we attribute the only slight lower score of Pick-up entity compared with *GCN(TR)* to the annotations. For example in Fig. 4, when an entity such as the Pick-up time ‘18:47’ is too blurred to read, it is tagged as NULL. However, our model can still correctly read and extract this entity out, which leads to lower statistics.

Evaluation on SROIE: We perform two sets of experiments and the results are as shown in Table 4.

Setting 1: We train text reading module all by ourselves and report comparisons. Notice that, we do *not* employ tricks of data synthesis and model ensemble in the training of text reading. Since entities of ‘Company’ and ‘Total’ often have distinguishing visual features (e.g., bold type or large font), as shown in Fig. 1(b), benefiting from fusion of visual and textual features, our model outperforms three counterparts by a large margin.

Setting 2: For fair comparisons on the leaderboard, we use the groundtruth of text bounding-box and transcript provided officially.

Table 3: Performance comparisons (F1-score) on Taxi Invoices.

Entities	Chargrid(TR)	NER(TR)	GCN(TR)	Our Model
Code	89.4	94.5	97.0	98.2
Number	85.3	92.4	93.7	95.4
Date	89.8	82.5	93.0	94.9
Pick-up time	82.9	60.0	86.3	84.6
Drop-off time	87.4	81.1	91.0	93.6
Price	93.0	94.5	93.6	94.9
Distance	92.7	93.6	91.4	94.4
Waiting	89.2	85.4	91.0	92.4
Amount	80.2	86.3	88.7	90.9
Avg	87.77	85.59	91.74	93.26



Figure 4: The Pick-up time which is too blur to read is annotated as NULL, but our model correctly read and extract it out, which results in lower accuracy.

Table 4: Performance comparisons (F1-score) on SROIE.

Setting	Model	F1-Score
Setting 1: Prediction of bboxes and transcript of texts	Chargrid(TR)	78.24
	NER(TR)	69.09
	GCN(TR)	76.51
	Our model	82.06
Setting 2: Groundtruth of bboxes and transcript of texts	Character-Word LSTM [24]	90.85
	LayoutLM[54]	95.24
	PICK[58]	96.12
	Our model	96.18

Character-Word LSTM is similar to NER [24], which applies LSTM on character and word level sequentially. LayoutLM [54] makes use of large pre-training data and fine-tunes on SROIE. Similar to LayoutLM, PICK [58] extracts rich semantic representation containing the textual and visual features as well as global layout. Compared with these methods, our model shows competitive performance.

Evaluation on Resumes: On the variable layout and semi-structured dataset, our method still shows impressive results. Table 5 shows

Table 5: Performance comparisons (F1-score) on Resumes.

Entities	Chargrid(TR)	NER(TR)	GCN(TR)	Our Model
Name	43.4	42.7	42.9	45.7
Phone	87.0	86.6	83.3	88.0
E-mail	70.9	69.6	68.0	74.9
Edu-period	77.1	68.7	62.2	81.4
University	74.7	86.0	82.3	87.4
Major	72.1	80.4	78.7	80.8
Avg	70.87	72.33	69.57	76.3
Speed(fps)	1.13	1.69	1.62	1.76

Table 6: Effects of multimodal features on information extraction.

	✓	✓	✓	✓
Text feat only	✓	✓	✓	✓
Textual Context feat		✓		✓
Visual Context feat			✓	✓
Accuracy	74.33	92.30	92.70	93.26

the results. We find that *GCN(TR)* has difficulty in adapting to such flexible layout. *Chargrid(TR)* obtains impressive performances on isolated entities such as Name, Phone and Education period. Since University and Major entities are often blended with other texts, *Chargrid(TR)* may be failed to extract these entities. As expected, *NER(TR)* performs better on this dataset, thanks to the inherent serializable property. While it is inferior to *Chargrid(TR)* on isolated entities due to the missing of layout information. Our model inherits the advantages of both *Chargrid(TR)* and *NER(TR)*, providing context features for identifying entities and performing entity extraction in the character-level. In short, our model gets comprehensive gain.

4.4 Discussion

In this section, we offer more detailed discussion and deeper understanding on different part of the proposed *TRIE* framework.

4.4.1 Effect of multi-modality features on information extraction: To further examine the contributions of visual and textual context features to information extraction, we perform the following ablation study on the Taxi Invoices and the result is shown in Table 6. *Text feat only* means that entities are extracted using features from text reading module only. Since the layout information is completely lost, this method presents worst performance. Introducing either the *visual context features* or *textual context features* brings significant performance gains. Notice that, introducing *visual context features* outperforms *textual context features* slightly, revealing the effects of visual context features. Further fusion of the above multi-modality features gives the best performance, which verifies the effect of multi-modality features. Visual examples in Fig. 5 illustrates the effects of each piece of context.

4.4.2 Effect of the end-to-end framework on text reading: To verify the effect of end-to-end framework on text reading, we evaluate GCN [30] model on two sets of text reading results, and the



(a) Effect of visual context (b) Effect of textual context

Figure 5: Illustration of effects of visual and textual context (highlighted with box). Taxi Code in (a) has distinguishing visual features, such as bold type and large font, which can boost information extraction performance. Invoices in (b) have unusual layouts where the visual context often fails. The textual context models dependencies and relationships between texts.

Table 7: Effect of end-to-end framework on text reading.

Text Reading Results		IE Model	Accuracy
TR only	End-to-End (TRIE)		
✓		GCN [30]	91.70
	✓	GCN [30]	92.60
	✓	TRIE	93.26

comparison results are shown in Table 7. *TR only* means training text reading module individually, and *End-to-End (TRIE)* is our proposed model. GCN on results of *End-to-End (TRIE)* shows 0.9 F1-score gain over that on *TR only*, demonstrating the effectiveness of end-to-end training. Some visualization examples are shown in Fig. 6. The Distance ‘23.3km’ and Date ‘2017-12-21’ are incorrectly recognized as ‘23.38m’ and ‘2010-02-21’ in pipeline training, leading to wrong information extraction results. While we can obtain their correct results when evaluating in an end-to-end training way. These quantitative and qualitative results uncover the effect of end-to-end framework on text reading. Considering effects of multimodal features (See Section 4.4.1), we can conclude that joint modeling of text reading and information extraction can benefit both tasks.

4.4.3 Speed: We evaluate the running time of our model and its counterparts in *fps* (frames per second). Results are shown in Table 5. Since *text reading* needs to detect and recognize all the texts in images, it takes up most of the time and runs at 1.83 *fps*. Compared to *text reading*, the *information extraction* is much more lightweight. That is, *Chargrid*, *NER* and *GCN* run at 2.97, 22.40 and 14.58 *fps*, respectively. In total, their full pipeline *Chargrid(TR)*, *NER(TR)* and *GCN(TR)* report speed of 1.13, 1.69 and 1.62 respectively. Benefiting from feature sharing between *text reading* and *information extraction*, *TRIE* achieves the highest speed of 1.76 *fps*. Note that, all methods are tested on Resumes. *Text reading* uses 2666×1600 size



Pipeline End-to-end Pipeline End-to-end
(a) Distance (b) Date

Figure 6: Visualization of text reading results of pipeline training and our end-to-end training. Best viewed in color.

Table 8: Effects of Layers and Heads in Textural Context Block.

Datasets	Layers	Heads			
		2	4	8	16
Taxi Invoices	1	92.97	92.98	92.86	92.72
	2	93.00	92.98	93.26	92.71
	3	92.55	92.81	93.06	92.83
Resumes	1	75.20	75.21	75.47	74.53
	2	75.62	76.25	76.28	75.86
	3	75.55	75.74	76.35	76.35

image as input and 32×512 image patches for recognition. *Chargrid* resolution is set to 1344×950 to reduce computations.

4.4.4 Ablation on layers and heads in Textural Context Block: To analyze the effects of different numbers of layers and heads in the *Textural Context Block*, we perform the following ablation studies, and the results are shown in Table 8. Taxi Invoices is relatively simple and has fixed layout, thus model with 1 or 2 layers and small number of heads gives the best result. As the layers and heads grow deeper, the model is prone to overfit. The layout of Resumes is much more variable than Taxi Invoices, thus it requires deeper layers and heads. In practice, one can adjust these settings according to the complexity of a specific task.

5 CONCLUSION

In this paper, we presented an end-to-end network to bridge the text reading and information extraction for document understanding. These two tasks can mutually reinforce each other through joint training. The visual and textual features of text reading can boost the performances of information extraction while the loss of information extraction can also supervise the optimization of text reading. On a variety of benchmarks, from structured to semi-structured text type and fixed to variable layout, our proposed method significantly outperforms three state-of-the-art methods in both aspects of efficiency and accuracy.

REFERENCES

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-Up and Top-Down Attention for Image Captioning and Visual Question Answering. In *CVPR*. 6077–6086.
- [2] Fedor Borisov, Albert Gordo, and Viswanath Sivakumar. 2018. Rosetta: Large Scale System for Text Detection and Recognition in Images. In *KDD*. 71–79.
- [3] Michal Busta, Lukas Neumann, and Jiri Matas. 2017. Deep TextSpotter: An End-to-End Trainable Scene Text Localization and Recognition Framework. In *ICCV*. 2223–2231.
- [4] Manuel Carbonell, Alicia Fornés, Mauricio Villegas, and Josep Lladós. 2019. TreyNet: A Neural Model for Text Localization, Transcription and Named Entity Recognition in Full Pages. *arXiv preprint arXiv:1912.10016* (2019).
- [5] Zhanzhan Cheng, Fan Bai, Yunlu Xu, Gang Zheng, Shiliang Pu, and Shuigeng Zhou. 2017. Focusing Attention: Towards Accurate Text Recognition in Natural Images. In *ICCV*. 5086–5094.
- [6] Junyoung Chung, Çağlar Gülçehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. *arXiv preprint arXiv:1412.3555* (2014).
- [7] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G. Carbonell, Quoc Viet Le, and Ruslan Salakhutdinov. 2019. Transformer-XL: Attentive Language Models beyond a Fixed-Length Context. In *ACL*. 2978–2988.
- [8] Andreas Dengel and Bertin Klein. 2002. smartFXT: A Requirements-Driven System for Document Analysis and Understanding. In *DAS (Lecture Notes in Computer Science)*, Vol. 2423. 433–444.
- [9] Timo I. Denk and Christian Reisswig. 2019. BERTgrid: Contextualized Embedding for 2D Document Representation and Understanding. *arXiv preprint arXiv:1909.04948* (2019).
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*. 4171–4186.
- [11] Daniel Esser, Daniel Schuster, Klemens Muthmann, Michael Berger, and Alexander Schill. [n.d.]. Automatic indexing of scanned documents: a layout-based approach. In *Document Recognition and Retrieval XIX, part of the IS&T-SPIE Electronic Imaging Symposium (SPIE Proceedings)*, Vol. 8297. 82970H.
- [12] Wei Feng, Wenhao He, Fei Yin, Xu-Yao Zhang, and Cheng-Lin Liu. 2019. TextDragon: An End-to-End Framework for Arbitrary Shaped Text Spotting. In *ICCV*. 9075–9084.
- [13] Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. 2016. Multimodal Compact Bilinear Pooling for Visual Question Answering and Regional Grounding. In *EMNLP*. 457–468.
- [14] He Guo, Xiameng Qin, Jiaming Liu, Junyu Han, Jingtuo Liu, and Errui Ding. 2019. EATEN: Entity-Aware Attention for Single Shot Visual Text Extraction. In *ICDAR*. 254–259.
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. 2017. Mask R-CNN. In *ICCV*. 2980–2988.
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *CVPR*. 770–778.
- [17] Pan He, Weilin Huang, Tong He, Qile Zhu, Yu Qiao, and Xiaolin Li. 2017. Single Shot Text Detector with Regional Attention. In *ICCV*. 3066–3074.
- [18] Tong He, Zhi Tian, Weilin Huang, Chunhua Shen, Yu Qiao, and Changming Sun. 2018. An End-to-End TextSpotter With Explicit Alignment and Attention. In *CVPR*. 5020–5029.
- [19] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation* 9, 8 (1997), 1735–1780.
- [20] Zheng Huang, Kai Chen, Jianhua He, Xiang Bai, Dimosthenis Karatzas, Shijian Lu, and C. V. Jawahar. 2019. ICDAR2019 Competition on Scanned Receipt OCR and Information Extraction. In *ICDAR*. 1516–1520.
- [21] Scott B. Huffman. 1995. Learning information extraction patterns from examples. In *Connectionist, Statistical, and Symbolic Approaches to Learning for Natural Language Processing (Lecture Notes in Computer Science)*, Vol. 1040. 246–260.
- [22] Douglass Russell Judd, Bruce Karsh, Ram Subbaroyan, Troy Toman, Rahul Lahiri, and Patrick Lok. 2004. Apparatus and method for searching and retrieving structured, semi-structured and unstructured content. US Patent App. 10/439,338.
- [23] 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*. 4459–4469.
- [24] Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. Neural Architectures for Named Entity Recognition. In *NAACL-HLT*. 260–270.
- [25] Chen-Yu Lee and Simon Osindero. 2016. Recursive Recurrent Nets with Attention Modeling for OCR in the Wild. In *CVPR*. 2231–2239.
- [26] Hui Li, Peng Wang, and Chunhua Shen. 2017. Towards End-to-End Text Spotting with Convolutional Recurrent Neural Networks. In *ICCV*. 5248–5256.
- [27] Liumin Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. VisualBERT: A Simple and Performant Baseline for Vision and Language. *arXiv preprint arXiv:1908.03557* (2019).
- [28] Minghui Liao, Baoguang Shi, Xiang Bai, Xinggong Wang, and Wenyu Liu. 2017. TextBoxes: A Fast Text Detector with a Single Deep Neural Network. In *AAAI*. 4161–4167.
- [29] Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. 2017. Feature Pyramid Networks for Object Detection. In *CVPR*. 936–944.
- [30] Xiaojing Liu, Feiyu Gao, Qiong Zhang, and Huasha Zhao. 2019. Graph Convolution for Multimodal Information Extraction from Visually Rich Documents. In *NAACL-HLT*. 32–39.
- [31] Xuebo Liu, Ding Liang, Shi Yan, Dagui Chen, Yu Qiao, and Junjie Yan. 2018. FOTS: Fast Oriented Text Spotting With a Unified Network. In *CVPR*. 5676–5685.
- [32] Yuliang Liu and Lianwen Jin. 2017. Deep Matching Prior Network: Toward Tighter Multi-oriented Text Detection. In *CVPR*. 3454–3461.
- [33] Shangbang Long, Jiaqiang Ruan, Wenjie Zhang, Xin He, Wenhao Wu, and Cong Yao. 2018. Textsnake: A Flexible Representation for Detecting Text of Arbitrary Shapes. In *ECCV*. 19–35.
- [34] Jiaseen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks. In *NeurIPS*. 13–23.
- [35] Pengyuan Lyu, Minghui Liao, Cong Yao, Wenhao Wu, and Xiang Bai. 2018. Mask TextSpotter: An End-to-End Trainable Neural Network for Spotting Text with Arbitrary Shapes. In *ECCV*. 71–88.
- [36] Xuezhe Ma and Eduard H. Hovy. 2016. End-to-end Sequence Labeling via Bi-directional LSTM-CNNs-CRF. In *ACL*.
- [37] Bodhisattwa Prasad Majumder, Navneet Potti, Sandeep Tata, James Bradley Wendt, Qi Zhao, and Marc Najork. 2020. Representation Learning for Information Extraction from Form-like Documents. In *ACL*. 6495–6504.
- [38] Ion Muslea et al. 1999. Extraction patterns for information extraction tasks: A survey. In *AAAI*, Vol. 2.
- [39] Rasmus Berg Palm, Florian Laws, and Ole Winther. 2019. Attend, Copy, Parse End-to-end Information Extraction from Documents. In *ICDAR*. 329–336.
- [40] Rasmus Berg Palm, Ole Winther, and Florian Laws. 2017. CloudScan - A Configuration-Free Invoice Analysis System Using Recurrent Neural Networks. In *ICDAR*. 406–413.
- [41] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *NeurIPS*. 8024–8035.
- [42] Liang Qiao, Sanli Tang, Zhanzhan Cheng, Yunlu Xu, Yi Niu, Shiliang Pu, and Fei Wu. 2020. Text Perceptron: Towards End-to-End Arbitrary-Shaped Text Spotting. In *AAAI*.
- [43] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *NeurIPS*. 91–99.
- [44] Ellen Riloff. 1993. Automatically Constructing a Dictionary for Information Extraction Tasks. In *Proceedings of the 11th National Conference on Artificial Intelligence*. 811–816.
- [45] Clément Sage, Alexandre Aussem, Haytham Elghazel, Véronique Eglin, and Jérémy Espinas. 2019. Recurrent Neural Network Approach for Table Field Extraction in Business Documents. In *ICDAR*. 1308–1313.
- [46] Erik F. Tjong Kim Sang and Jorn Veenstra. 1999. Representing Text Chunks. In *EACL*. 173–179.
- [47] Daniel Schuster, Klemens Muthmann, Daniel Esser, Alexander Schill, Michael Berger, Christoph Weidling, Kamil Aliyev, and Andreas Hofmeier. 2013. Intellix - End-User Trained Information Extraction for Document Archiving. In *ICDAR*. 101–105.
- [48] Khaled Shaalan. 2014. A Survey of Arabic Named Entity Recognition and Classification. *Comput. Linguistics* 40, 2 (2014), 469–510.
- [49] Baoguang Shi, Xiang Bai, and Serge J. Belongie. [n.d.]. Detecting Oriented Text in Natural Images by Linking Segments. In *CVPR*.
- [50] Baoguang Shi, Xiang Bai, and Cong Yao. 2017. An End-to-End Trainable Neural Network for Image-Based Sequence Recognition and Its Application to Scene Text Recognition. *IEEE TPAMI*. 39, 11 (2017), 2298–2304.
- [51] Stephen Soderland. 1999. Learning Information Extraction Rules for Semi-Structured and Free Text. *Mach. Learn.* 34, 1-3 (1999), 233–272.
- [52] Hao Wang, Pu Lu, Hui Zhang, Mingkun Yang, Xiang Bai, Yongchao Xu, Mengchao He, Yongpan Wang, and Wenyu Liu. 2020. All You Need Is Boundary: Toward Arbitrary-Shaped Text Spotting. In *AAAI*.
- [53] Wenhao Wang, Enze Xie, Xiang Li, Wenbo Hou, Tong Lu, Gang Yu, and Shuai Shao. 2019. Shape Robust Text Detection With Progressive Scale Expansion Network. In *CVPR*.
- [54] 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. *KDD* (2020).
- [55] Vikas Yadav and Steven Bethard. 2018. A Survey on Recent Advances in Named Entity Recognition from Deep Learning models. In *COLING*. 2145–2158.
- [56] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. XLNet: Generalized Autoregressive Pretraining for

- Language Understanding. In *NeurIPS*. 5754–5764.
- [57] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alexander J. Smola. 2016. Stacked Attention Networks for Image Question Answering. In *CVPR*. 21–29.
- [58] Wenwen Yu, Ning Lu, Xianbiao Qi, Ping Gong, and Rong Xiao. 2020. PICK: Processing Key Information Extraction from Documents using Improved Graph Learning-Convolutional Networks. *ICPR (2020)*.
- [59] Matthew D. Zeiler. 2012. ADADELTA: An Adaptive Learning Rate Method. *arXiv preprint arXiv:1212.5701* (2012).
- [60] Xiang Zhang, Junbo Jake Zhao, and Yann LeCun. 2015. Character-level Convolutional Networks for Text Classification. In *NeurIPS*. 649–657.
- [61] Xiaohui Zhao, Zhuo Wu, and Xiaoguang Wang. 2019. CUTIE: Learning to Understand Documents with Convolutional Universal Text Information Extractor. *arXiv preprint arXiv:1903.12363* (2019).
- [62] Xinyu Zhou, Cong Yao, He Wen, Yuzhi Wang, Shuchang Zhou, Weiran He, and Jiajun Liang. 2017. EAST: An Efficient and Accurate Scene Text Detector. In *CVPR*. 2642–2651.