

Levenshtein OCR

Cheng Da ^{*}, Peng Wang ^{*}, and Cong Yao [†]

Alibaba DAMO Academy, Beijing, China
{dc.dacheng08,wdp0072012,yaocong2010}@gmail.com

Abstract. A novel scene text recognizer based on Vision-Language Transformer (VLT) is presented. Inspired by Levenshtein Transformer in the area of NLP, the proposed method (named Levenshtein OCR, and LevOCR for short) explores an alternative way for automatically transcribing textual content from cropped natural images. Specifically, we cast the problem of scene text recognition as an iterative sequence refinement process. The initial prediction sequence produced by a pure vision model is encoded and fed into a cross-modal transformer to interact and fuse with the visual features, to progressively approximate the ground truth. The refinement process is accomplished via two basic character-level operations: *deletion* and *insertion*, which are learned with imitation learning and allow for parallel decoding, dynamic length change and good interpretability. The quantitative experiments clearly demonstrate that LevOCR achieves state-of-the-art performances on standard benchmarks and the qualitative analyses verify the effectiveness and advantage of the proposed LevOCR algorithm.

Keywords: Scene Text Recognition, Transformer, Interpretability

1 Introduction

Scene text recognition is a long-standing and challenging problem [28, 52, 4] that has attracted much attention from the computer vision community. It aims at decoding textual information from natural scene images, which could be very beneficial to down-stream applications, such as traffic sign recognition and content-based image retrieval. However, reading text from natural images is faced with numerous difficulties: variation in text style and shape, non-uniform illumination, partial occlusion, perspective distortion, to name just a few. Recently, various text recognition methods [28] have been proposed to tackle this tough problem and substantial progresses have been observed [44, 48, 47, 8, 32, 51, 14, 26].

It has become a trend in the computer vision community to draw inspirations from methods initially proposed for NLP tasks to solve vision problems, for instance, ViT [7], DETR [3] and Swin-Transformer [27]. Also, in the field of scene text recognition, multiple recent works [47, 8] start to incorporate linguistic knowledge into the text recognition process, fusing information from both the vision and language modalities for higher text recognition accuracy.

^{*} Equal contribution. [†] Corresponding author.

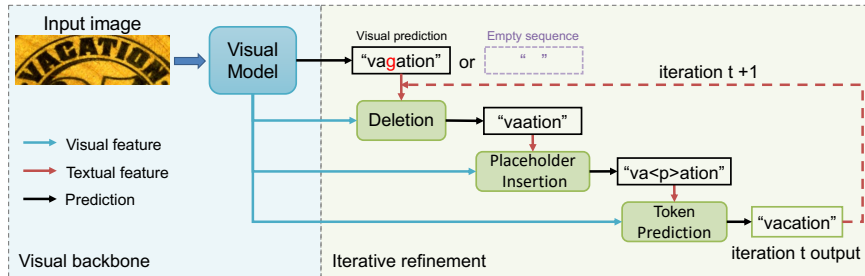


Fig. 1. Schematic overview of LevOCR. LevOCR accomplishes text recognition in an iterative way through two basic operations: *deletion* and *insertion*. Note that in LevOCR the operation of *insertion* is further decomposed into two sub-operations: Placeholder Insertion and Token Prediction.

Inspired by the wisdom from these pioneering works, we propose an alternative algorithm for scene text recognition. The backbone of the proposed model is a Vision-Language Transformer (VLT) [38, 5, 20], which is employed to perform cross-modal information fusion for more informative representations and better recognition performance. To further facilitate more flexible and interpretable text recognition, we introduce the strategy from Levenshtein Transformer (LevT) [10], which was originally designed for sequence generation and refinement tasks in NLP, into our framework. The core idea is to learn one refinement policy (deletion or insertion) for current iteration from its adversary in the previous iteration, in order to manipulate the basic units in the sequence (corresponding to characters in text recognition) to approach the target sequence. In such a way, the proposed text recognizer Levenshtein OCR (LevOCR for short) can realize text recognition in a progressive fashion, where the final prediction is obtained by iteratively refining an initial or intermediate recognition result until convergence (*i.e.*, post-editing for error correction). An intuitive illustration is depicted in Fig. 1. Note that due to the diversity of data augmentation in the training phase, the proposed model also supports generating the final recognition from an *empty* sequence, which falls back to a text generation task.

Similar to ABINet [8], we also fuse the information of both the vision modality and the language modality in an iterative procedure to predict the final recognition result. However, there are two key differences: (1) The main architecture of LevOCR is a Vision-Language Transformer (VLT), which allows for more sufficient utilization of the interactions between vision and language; (2) More importantly, while ABINet produces whole sequences at each iteration, LevOCR performs *fine-granularity predictions* through character-level operations (deletion or insertion of individual characters), endowing the system with higher flexibility and better *interpretability*, *i.e.*, when a specific decision (deletion or insertion) is made, one can trace back to the input space (image or text) to examine the supporting cues for that decision. This constitutes a unique characteristic of our LevOCR algorithm.

We have conducted both qualitative and quantitative experiments on widely-used benchmark datasets in the field of scene text recognition to verify the effectiveness of the LevOCR algorithm. LevOCR not only achieves state-of-the-art recognition performances on various benchmarks (see the tables in Sec. 4), but also provides clear and intuitive interpretation for each action prediction (see the example in Sec. 4.8 for more details).

In summary, the contributions of the work are as follows: (1) We propose a novel, cross-modal transformer based scene text recognizer, which fully explores the interactions between vision and language modalities and accomplishes text recognition via an iterative process. (2) The proposed LevOCR allows for parallel decoding and dynamic length change, and exhibits good transparency and interpretability in the inference phase, which could be very crucial for diagnosing and improving text recognition models in the future. (3) LevOCR achieves state-of-the-art results on standard benchmarks, and extensive experiments verify the effectiveness and advantage of LevOCR.

2 Related Work

2.1 Scene Text Recognition Methods

Traditional methods directly cast scene text recognition as a sequence classification task, which is purely based on visual features without any explicitly linguistic knowledge. CTC-based methods [36, 41, 13, 15] provide differentiable Connectionist Temporal Classification (CTC) loss for access to end-to-end trainable sequence recognition, among which RNN model is often employed for context modeling of feature maps extracted by CNNs[36]. Segmentation-based methods [25, 40] utilize FCN to directly predict the character labels in pixel-level and further group characters into words, in which character-level annotation is required. Opposite to parallel prediction of CTC and segmentation methods, attention-based methods [21, 6, 37] with encoder-decoder mechanism sequentially generates characters in order via RNN-attention model, where language information between characters can be implicitly captured. Due to the promising results, attention-based methods have previously dominated this field.

2.2 Enhanced Attention-based Methods

Considering irregular text, previous methods [50, 37] integrate spatial transformer module into attention-based framework, which rectifies the input with perspective and curvature distortion into a more canonical form. [6] observes the attention drift problem [42], in which the alignments between feature areas and text targets are not accurate for complicated images, and proposes a focus network to suppress the attention adrift. RobustScanner [48] utilizes positional clues to decode random character sequences effectively, by introducing a position enhancement branch into attention-based framework. Furthermore, SE-ASTER [33] employs a pre-trained language model to predict additional semantic information, which can guide the decoding process.

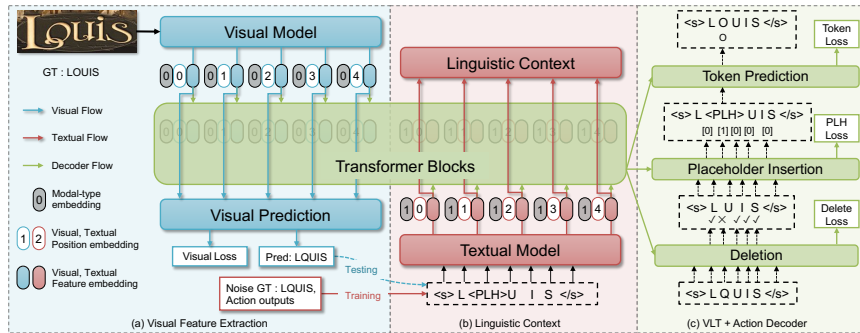


Fig. 2. Overview of the architecture of LevOCR.

2.3 Transformer-based Methods

Recently, transformer units [39] are employed in scene text recognition model to replace the complex LSTM blocks of RNN. Some prior methods [22, 35, 29] provide conventional transformer-based encoder-decoder framework for text recognition, which is based on customized CNNs feature extraction block. Furthermore, TrOCR [23] proposes an encoder-decoder structure with pre-trained transformer-based CV and NLP models, to deal with original image patches directly. ViTSTR [1] then employs vision transformers (ViT) for scene text recognition with image patches, where only transformer encoder is required and so that the characters can be predicted in parallel. To handle the linguistic context well, VisionLAN [45] presents a visual reasoning module to capture the visual and linguistic information simultaneously by masking the input image in the feature level. Additionally, SRN [47] introduces a global semantic reasoning module with transformer units to enhance semantic context. ABInet [8] goes a further step and proposes an explicit language model via transformer.

3 Methodology

LevOCR consists of three essential modules: Visual Feature Extraction (VFE), Linguistic Context (LC) and Vision-Language Transformer (VLT). Given an input image, the CNN model first extracts visual information and makes an initial visual prediction. And the initial visual prediction is further fed into transformer-based textual module to generate linguistic context. Finally, the intermediate visual features and linguistic features are directly concatenated and fed into VLT for sufficient information interaction, without any explicit cross-modal alignments. Additionally, different action decoders are built upon VLT for subsequent task. The schematic pipeline is illustrated in Fig. 2.

3.1 Visual Feature Extraction

Given a image $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$ and the corresponding text label $\mathbf{y} = (y_1, y_2, \dots, y_N)$ with N maximum text length, a modified ResNet [37, 12] backbone is utilized for visual information extraction, and then transformer units [39] are employed to generate enhanced 2D visual features $\mathbf{F}_v \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times D}$, where D is the feature dimension. We directly decrease the height of feature \mathbf{F}_v to 3 by 2 convolution layers, generating $\mathbf{V} \in \mathbb{R}^{3 \times \frac{W}{4} \times D}$ as the visual feature for subsequent refinement task. In addition, we construct another position attention [8] branch on feature \mathbf{F}_v to generate the initial visual prediction $\hat{\mathbf{y}} = (\hat{y}_1, \hat{y}_2, \dots, \hat{y}_N)$. Then, the visual loss L_v can be realized by a cross-entropy loss between \mathbf{y} and $\hat{\mathbf{y}}$. Notably, the pure visual feature \mathbf{V} is not used for visual prediction, in order to preserve more feature information for subsequent refinement task. The pipeline of visual information extraction is illustrated in Fig. 2 (a) with blue color.

3.2 Linguistic Context

NLP community has achieved substantial performance improvement. to model linguistic knowledge. The textual module is constructed with Transformer blocks [39] to model linguistic knowledge. Concretely, the input of textual module is a text sequence that need to be corrected, represented by $\tilde{\mathbf{y}} = (\tilde{y}_1, \tilde{y}_2, \dots, \tilde{y}_N)$. First, word embedding at character level is used to encode $\tilde{\mathbf{y}}$ into feature $\mathbf{F}_t \in \mathbb{R}^{N \times D}$. Then, multiple transformer blocks transcribe \mathbf{F}_t into refined text feature $\mathbf{T} \in \mathbb{R}^{N \times D}$, where N is the maximum text length and D is the feature dimension. The pipeline of linguistic context is illustrated in Fig. 2 (b) with red color.

3.3 Vision-Language Transformer

Text instances in natural scenes do not always conform with linguistic rules. For example, digits and random characters appear commonly. Therefore, LevOCR employs VLT [20] to integrate visual and linguistic features. In this way, no enforced alignments as in ABInet [8] of two modalities are required and adaptive weights of two modalities are directly driven by the objective function. Then, action decoder heads are built upon VLT for *deletion* and *insertion* action learning, which can make a complementary judgement on both two modalities and be explicitly examined for good interpretability. Specifically, the visual feature $\mathbf{V} = [\mathbf{v}_1; \mathbf{v}_2; \dots; \mathbf{v}_{N_v}] \in \mathbb{R}^{N_v \times D}$ and textual feature $\mathbf{T} = [\mathbf{t}_1; \mathbf{t}_2; \dots; \mathbf{t}_N] \in \mathbb{R}^{N \times D}$ are produced by the corresponding modules, respectively. We directly concatenate these features of two modalities as the input of VLT blocks. In order to discriminate the features with different modalities and positions, position embeddings and modal-type ones are introduced. The detailed process of VLT is formulated as follows:

$$\begin{aligned}
 \bar{\mathbf{V}} &= [\mathbf{v}_1 + \mathbf{p}_1^v; \mathbf{v}_2 + \mathbf{p}_2^v; \dots; \mathbf{v}_{N_v} + \mathbf{p}_{N_v}^v] + \mathbf{E}_v \\
 \bar{\mathbf{T}} &= [\mathbf{t}_1 + \mathbf{p}_1^t; \mathbf{t}_2 + \mathbf{p}_2^t; \dots; \mathbf{t}_N + \mathbf{p}_N^t] + \mathbf{E}_t \\
 \mathbf{H}^{(0)} &= [\mathbf{h}_1^{(0)}; \mathbf{h}_2^{(0)}; \dots; \mathbf{h}_{N+N_v}^{(0)}] = [\bar{\mathbf{T}}; \bar{\mathbf{V}}] \\
 \mathbf{H}^{(l+1)} &= \text{BERT}_l(\mathbf{H}^{(l)}).
 \end{aligned} \tag{1}$$

Here, N_v is $3 \times \frac{W}{4}$. The visual and textual position embeddings are represented as $[\mathbf{p}_1^v; \mathbf{p}_2^v; \dots; \mathbf{p}_{N_v}^v] \in \mathbb{R}^{N_v \times D}$ and $[\mathbf{p}_1^t; \mathbf{p}_2^t; \dots; \mathbf{p}_N^t] \in \mathbb{R}^{N \times D}$. Then, visual and textual modal-type embeddings are denoted as $\mathbf{E}_v \in \mathbb{R}^{1 \times D}$ and $\mathbf{E}_t \in \mathbb{R}^{1 \times D}$. And the l -th transformer block is denoted as BERT_l . Thus, the final aggregated feature $\mathbf{H} \in \mathbb{R}^{(N+N_v) \times D}$ are generated by L -th transformer block, in which even the unaligned features of two modalities can be adaptively interacted and fused.

3.4 Imitation Learning

In order to mimic how humans edit text, we cast this text sequence refinement task into a Markov Decision Process (MDP) denoted as a tuple $(\mathcal{Y}, \mathcal{A}, \mathcal{E}, \mathcal{R}, \mathbf{y}^0)$ as in [10]. We define a text as a sequence that consists of digits and characters, and thus \mathcal{Y} is a set of word vocabulary with the dictionary of symbols \mathcal{V} . Typically, $\mathbf{y}^0 \in \mathcal{Y}$ represents the initial sequence. For text refinement task, two basic actions *deletion* and *insertion* are defined as the set of actions \mathcal{A} . The reward function $\mathcal{R} = -\mathcal{D}(\mathbf{y}, \mathbf{y}^*)$ directly measures the Levenshtein distance between the prediction and the sequence of ground-truth text. Given k -th step text sequence \mathbf{y}^k , the agent interacts with the environment \mathcal{E} , executes editing actions and returns the modified sequence \mathbf{y}^{k+1} , which is denoted as $\mathbf{y}^{k+1} = \mathcal{E}(\mathbf{y}^k, \mathbf{a}^{k+1})$. Our main purpose is to learn a favourable policy π that model the probability distribution over actions \mathcal{A} for maximum reward.

Deletion Action The input text sequence for imitation learning is also denoted as $\mathbf{y} = (y_1, y_2, \dots, y_N)$ for simplicity. Deletion policy $\pi^{del}(d|i, \mathbf{y})$ aims to make a binary decision for every character $y_i \in \mathbf{y}$, in which $d = 1$ indicates that this token should be deleted or $d = 0$ for keeping it. Typically, y_1 and y_N are special symbols $\langle s \rangle$ and $\langle /s \rangle$ for sequence boundary, respectively. Thus, they can not be deleted, which is denoted as $\pi^{del}(0|1, \mathbf{y}) = \pi^{del}(0|N, \mathbf{y}) = 1$. Moreover, based on the aggregated feature \mathbf{H} , the deletion classifier can be formulated as follows:

$$\pi_{\theta}^{del}(d|i, \mathbf{y}) = \text{softmax}(\mathbf{h}_i \mathbf{W}_{del}^T), i = 2, \dots, N - 1, \quad (2)$$

where $\mathbf{W}_{del} \in \mathbb{R}^{2 \times D}$ is the weight of deletion classifier. Note that only the first N sequences of \mathbf{H} are used for prediction, and N is the maximum text length.

Insertion Action Insertion action is more complicated than deletion one, since the position of insertion need to be predefined. Technically, *insertion* is decomposed into two sub-operations: placeholder insertion and token prediction. Concretely, for each consecutive pairs (y_i, y_{i+1}) in \mathbf{y} , placeholder insertion policy $\pi^{plh}(p|i, \mathbf{y})$ predicts the number p of placeholder should be inserted at position i . Thus, the classifier of placeholder insertion is defined as follows:

$$\pi_{\theta}^{plh}(p|i, \mathbf{y}) = \text{softmax}([\mathbf{h}_i, \mathbf{h}_{i+1}] \mathbf{W}_{plh}^T), \quad (3)$$

$$i = 1, \dots, N - 1,$$

where $\mathbf{W}_{plh} \in \mathbb{R}^{M \times 2D}$ is the weight of the placeholder classifier, and M is the max number of placeholders can be inserted. $[\mathbf{h}_i, \mathbf{h}_{i+1}]$ is the concatenation of \mathbf{h}_i and \mathbf{h}_{i+1} .

Referring to the predicted number of placeholder insertion, we can insert a corresponding number of placeholders at the relevant positions. Then, token prediction policy $\pi^{tok}(t|i, \mathbf{y})$ is required to replace placeholder y_i with symbol t in the dictionary \mathcal{V} , which is formulated as follows:

$$\pi_{\theta}^{tok}(t|i, \mathbf{y}) = \text{softmax}(\mathbf{h}_i \mathbf{W}_{tok}^T), \forall y_i = \langle p \rangle. \quad (4)$$

Here, $\mathbf{W}_{tok} \in \mathbb{R}^{|\mathcal{V}| \times D}$ is the weight of token predictor, $\langle p \rangle$ is the placeholder and $|\mathcal{V}|$ is the size of dictionary.

Training Phase Notably, *deletion* and *insertion* are alternatively executed. For instance, given a text sequence, deletion policy is first called to delete wrong symbols. Then, placeholder insertion policy inserts some possible placeholders. Finally, token prediction policy replaces all placeholders with right symbols. Typically, these actions are performed in parallel for each position. Moreover, the imitation learning strategy is utilized for LevOCR training, aiming to approximating the expert policy π^* that can be directly and simply derived from the ground-truth text sequence as follows:

$$\mathbf{a}^* = \arg \min_{\mathbf{a}} \mathcal{D}(\mathbf{y}^*, \mathcal{E}(\mathbf{y}, \mathbf{a})). \quad (5)$$

Here, Levenshtein distance \mathcal{D} is used for distance measure. The optimal actions \mathbf{a}^* can represent \mathbf{d}^* , \mathbf{p}^* , and \mathbf{t}^* , which can be produced by dynamic programming efficiently. The loss function of *deletion* is formulated as follows:

$$L_{del} = \mathbb{E}_{\mathbf{y}_{del} \sim d_{\tilde{\pi}_{del}}} \sum_{d_i^* \in \mathbf{d}^*} -\log \pi_{\theta}^{del}(d_i^*|i, \mathbf{y}_{del}), \quad (6)$$

where $\mathbf{d}^* = (d_1^*, d_2^*, \dots, d_N^*) \sim \pi^*$ denotes the optimal deletion action for each position of \mathbf{y}_{del} , generated by Eq. (5). And $d_{\tilde{\pi}_{del}}$ is a text distribution induced by policy $\tilde{\pi}_{del}$ for the sequence generation with additive noise:

$$d_{\tilde{\pi}_{del}} = \begin{cases} \{\mathcal{E}(\mathcal{E}(\mathbf{y}^0, \tilde{\mathbf{p}}), \tilde{\mathbf{t}}), \tilde{\mathbf{p}} \sim \pi^R, \tilde{\mathbf{t}} \sim \pi^R\}, \alpha < \mu \\ \{\mathcal{E}(\mathcal{E}(\mathbf{y}', \mathbf{p}^*), \tilde{\mathbf{t}}), \mathbf{p}^* \sim \pi_{\theta_*}^{plh}, \tilde{\mathbf{t}} \sim \pi_{\theta}^{tok}\}, \alpha \geq \mu, \end{cases} \quad (7)$$

where π^R represents a random policy, $\alpha \sim \text{Uniform}[0, 1]$, $\mu \in [0, 1]$ is a mixture factor, \mathbf{y}^0 is the initial sequence, and \mathbf{y}' is any sequence ready to insert. For $\mu < \alpha$, we randomly add some symbols on \mathbf{y}^0 to generate \mathbf{y}_{del} . For $\mu \geq \alpha$, we use expert placeholder policy and the learned token prediction to generate \mathbf{y}_{del} based on \mathbf{y}' . This procedure can be regarded as adversarial learning in GAN [9]. Similarly, the loss function of *insertion* is as follows:

$$L_{ins} = \mathbb{E}_{\mathbf{y}_{ins} \sim d_{\tilde{\pi}_{ins}}} \left[\sum_{p_i^* \in \mathbf{p}^*} -\log \pi_{\theta}^{plh}(p_i^*|i, \mathbf{y}_{ins}) + \sum_{t_i^* \in \mathbf{t}^*} -\log \pi_{\theta}^{tok}(t_i^*|i, \mathbf{y}'_{ins}) \right], \quad (8)$$

where $\mathbf{p}^* = (p_1^*, p_2^*, \dots, p_{N-1}^*) \sim \pi^*$ represents the optimal number of placeholders for each consecutive position pair in \mathbf{y}_{ins} , generated by Eq. (5). And $\mathbf{t}^* = (t_1^*, t_2^*, \dots, t_{N-1}^*) \sim \pi^*$ denotes the optimal symbol for each placeholder in \mathbf{y}'_{ins} , where $\mathbf{y}'_{ins} = \mathcal{E}(\mathbf{y}_{ins}, \mathbf{p})$, $\mathbf{p} \sim \pi_{\theta}^{plh}$. Moreover, $d_{\tilde{\pi}_{ins}}$ is a text distribution induced by policy $\tilde{\pi}_{ins}$ for the sequence generation with deleted noise:

$$d_{\tilde{\pi}_{ins}} = \begin{cases} \{\mathcal{E}(\mathbf{y}^*, \tilde{\mathbf{d}}), \tilde{\mathbf{d}} \sim \pi^R\}, \beta < \mu \\ \{\mathcal{E}(\mathbf{y}^0, \mathbf{d}^*), \mathbf{d}^* \sim \pi_{\theta^*}^{del}\}, \beta \geq \mu, \end{cases} \quad (9)$$

where factor $\beta \sim \text{Uniform}[0, 1]$. We also adopt mixture manner to construct \mathbf{y}_{ins} for insertion learning. For $\beta < \mu$, we randomly delete some symbols on ground-truth \mathbf{y}^* to produce \mathbf{y}_{ins} . For $\beta \geq \mu$, expert deletion policy is employed to generate \mathbf{y}_{ins} based on initial sequence \mathbf{y}^0 . The training procedure is illustrated in Fig. 2 (c) with green color, and the final loss function is formulated as:

$$L = \lambda_1 L_v + \lambda_2 L_{del} + \lambda_3 L_{ins}, \quad (10)$$

where λ_1 , λ_2 and λ_3 are weights for visual prediction, *deletion* and *insertion*.

Notably, the visual model is pre-trained with only images for better initialization by L_v . And the textual model and VLT blocks can also be pre-trained with only texts. Specifically, the input of VLT is $\mathbf{H}^{(0)} = \bar{\mathbf{T}}$ without image feature, which is used for *deletion* and *insertion* learning via L_{del} and L_{ins} . Based on these pre-trained models, LevOCR is further trained by Eq. (10).

Note that the input sequences (*i.e.* $\tilde{\mathbf{y}}$) for *deletion* and *insertion* are indeed different. Typically, $\tilde{\mathbf{y}}$ is not always a ‘‘true’’ word. For instance, $\tilde{\mathbf{y}}$ could be the output of placeholder insertion for token prediction that includes placeholders or even be an empty sequence. Therefore, different input sequences $\tilde{\mathbf{y}}$ should be fed into textual model individually and encoded as the unique text features \mathbf{T} and aggregated features \mathbf{H} for the specific action (*deletion*, *placeholder insertion* and *token prediction*) learning in training phase.

3.5 Inference Phase

We alternatively employ *deletion* and *insertion* to refine text at inference process, until two policies converge (either nothing to delete or insert, or reaching maximum iterations). Concretely, given a image \mathbf{x} , we first obtain the visual feature \mathbf{V} and the initial sequence \mathbf{y}^0 (*e.g.* visual prediction). And then \mathbf{y}^0 is fed into textual module, generating refined text texture \mathbf{T} . Furthermore, the visual feature and textual one are interacted and fused in VLT to produce aggregated feature \mathbf{H} . Finally, action decoders greedily choose the action with the maximum probability at each position by Eq. (2)(3)(4) in parallel. Then, *deletion* and *insertion* are executed in turn with new corresponding features \mathbf{T} and \mathbf{H} . Note that LevOCR can not only accomplish text refinement on initial visual predictions, but also perform text generation with empty sequence \mathbf{y}^0 .

4 Experiment

4.1 Datasets

For fair comparison, we follow previous settings [2, 8] to train LevOCR on two synthetic datasets MJSynth (MJ) [16, 17] and SynthText(ST) [11] without fine-tuning on other datasets. Extensive experiments are conducted on six standard Latin scene text benchmarks, including 3 regular text datasets (IC13 [19], SVT [43], IIIT [30]) and 3 irregular ones (IC15 [18], SVTP [31], CUTE [34]).

ICDAR 2013 (IC13) [19] includes 1095 cropped word images for testing. We evaluate on 857 images with alphanumeric characters and more than 2 characters. **Street View Text (SVT)** [43] contains 647 testing images collected from Google Street View. **IIIT 5K-Words (IIIT5k)** [30] is crawled from Google image search, and consists of 3000 testing images. **ICDAR 2015 (IC15)** [18] includes word patches cropped from incidental scene images captured by Google Glasses. **Street View Text-Perspective (SVTP)** [31] consists of 639 images collected from Google Street View, and many images are heavily distorted. **CUTE80 (CUTE)** [34] contains 80 natural scenes images for curved text recognition. 288 images are cropped from these images for testing.

4.2 Implementation Details

The textual model and VLT block consist of 6 stacked transformer units with 8 heads for each layer, respectively. The number of hidden units in FC-layer of transformer block is 2048, and the dimension D of visual and textual feature is set to 512. Besides, The size of symbol dictionary $|\mathcal{V}|$ is 40, including 0-9, $a-z$, $\langle s \rangle$, $\langle /s \rangle$, pad and $\langle p \rangle$. The max length of output sequence N is set to 28, and the max number of placeholders M is set to 28. The mixture factor μ is set to 0.5. λ_1 , λ_2 and λ_3 are set to 1. The input images are resized to 32×128 . Common data augmentation is employed, such as rotation, affine, perspective distortion, image quality deterioration and color jitter. Our approach is trained on NVIDIA Tesla V100 GPUs with batch size 128. Adadelata [49] optimizer is adopted with the initial learning rate 0.1, and the max training epoch is 10.

4.3 Text Refinement and Text Generation

The initial sequence \mathbf{y}^0 is pivotal in LevOCR, which not only determines the final performance, but also endows LevOCR with the ability of highly flexible text editing. To empirically verify the ability of text refinement and generation, we construct 4 kinds of initial sequence: (1) \mathbf{y}_{VP}^0 : the visual prediction is directly adopted and LevOCR essentially aims to text refinement. (2) \mathbf{y}_{Emp}^0 : empty sequence is simply used and thus the inference stage falls back to a text generation task. (3) \mathbf{y}_{Rand}^0 : ground-truth is corrupted by random noise, where we replace one character for 30% text, add one character for 30% text, delete one character for 40% text, and remain the digit text unchangeable. (4) \mathbf{y}_{GT}^0 : ground-truth. These results with different initial sequences are reported in Tab. 1.

Table 1. The accuracies of LevOCR with different initial sequences \mathbf{y}^0 and max iterations on 6 public benchmarks.

Initial Sequence	Iteration	IC13	SVT	IIIT	IC15	SVTP	CUTE	AVG
LevOCR _{VP}	-	95.10	90.57	95.23	83.98	83.41	87.84	90.65
LevOCR + \mathbf{y}_{VP}^0	#1	96.50	92.43	96.50	86.09	87.75	91.32	92.55
	#2	96.73	92.89	96.63	86.42	88.06	91.67	92.79
	#3	96.85	92.89	96.63	86.42	88.06	91.67	92.81
LevOCR + \mathbf{y}_{Emp}^0	#1	95.45	90.42	95.00	83.99	83.72	88.89	90.64
	#2	96.73	92.43	96.33	86.03	87.60	92.01	92.51
	#3	96.97	92.74	96.37	86.09	88.06	92.01	92.63
LevOCR + \mathbf{y}_{Rand}^0	#1	84.83	85.32	87.90	82.05	82.33	84.38	85.20
	#2	84.83	85.63	87.93	82.05	82.33	84.38	85.25
	#3	84.83	85.63	87.93	82.05	82.33	84.38	85.25
LevOCR + \mathbf{y}_{GT}^0	#1	99.07	98.15	98.67	92.16	97.83	97.22	96.91
	#2	99.07	98.15	98.67	92.16	97.83	97.22	96.91
	#3	99.07	98.15	98.67	92.16	97.83	97.22	96.91

Table 2. The accuracies of LevOCR with different backbones.


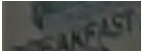
Methods	Backbone	IC13	SVT	IIIT	IC15	SVTP	CUTE	AVG
LevOCR _{VP} w/o LevT	CNN	95.21	90.41	95.30	83.26	83.41	88.88	90.53
LevOCR w/o LevT		95.21	90.42	95.43	83.43	84.03	89.23	90.70
LevOCR _{VP}	ViT	94.86	89.18	93.6	82.38	84.18	82.98	89.30
LevOCR		96.15	91.80	95.63	85.81	88.06	86.81	91.87
LevOCR _{VP}	CNN	95.10	90.57	95.23	83.98	83.41	87.84	90.65
LevOCR		96.85	92.89	96.63	86.42	88.06	91.67	92.81

The performance of \mathbf{y}_{Rand}^0 and \mathbf{y}_{GT}^0 can be regarded as the lower bound and upper bound of LevOCR, respectively. Based on the noisy initial sequences, \mathbf{y}_{Rand}^0 gets 85.25% average accuracy. This implies that almost 85% noisy texts can be corrected. Moreover, \mathbf{y}_{GT}^0 achieves the best accuracy 96.91% than others, proving that better linguistic context can be captured in better initial text sequence, leading to better performance of LevOCR. Additionally, the first line LevOCR_{VP} records the accuracy of pure visual prediction of LevOCR. Both \mathbf{y}_{VP}^0 and \mathbf{y}_{Emp}^0 can obtain obvious improvements over LevOCR_{VP}, showing the effectiveness of text refinement and generation of LevOCR. Generally, the accuracies of 4 constructions (*i.e.* $\mathbf{y}_{GT}^0 > \mathbf{y}_{VP}^0 > \mathbf{y}_{Emp}^0 > \mathbf{y}_{Rand}^0$) suggest that choosing a better \mathbf{y}^0 is crucial, and LevOCR indeed adopts linguistic information, not just relies on visual features. These results sufficiently demonstrate the powerful abilities of LevOCR in text refinement and text generation.

4.4 Iterative Refinement

We investigate the influence of various maximum refinement iterations (*i.e.*, 1, 2, 3) and the results are also shown in Tab. 1. Notably, even if the maximum iteration is set to 3, the refinement process might be stopped after 2 iterations

Table 3. The detailed iterative process of LevOCR with different initial sequences on 6 public benchmarks. “Iteration” represents the current number of refinement iteration. “Input” represents the initial sequence or the output of last iteration. “Delete”, “PLHins” and “TokenPred” represent the action output of deletion, placeholder insertion and token prediction, respectively. And $\langle 3p \rangle$ indicates that three consecutive placeholders $\langle p \rangle$ are inserted into the sequence. Best viewed in colors.

Image	GT	Initial Sequence	Iteration: Input-Delete-PLHins-TokenPred
1: IC13 	service	$y_{VP}^0 = \text{service}$	#1: service- sere-ser $\langle 3p \rangle$ e-service
		$y_{Emp}^0 = \langle s \rangle \langle /s \rangle$	#1: $\langle s \rangle \langle /s \rangle$ - $\langle s \rangle \langle /s \rangle$ - $\langle 7p \rangle$ -serbice #2: serbice- serbice-serb $\langle p \rangle$ ice- serblice #3: serblice- serice-ser $\langle p \rangle$ ice-service
		$y_{Rand}^0 = \text{servcce}$	#1: servcce- service-serv $\langle p \rangle$ ce-service
2: SVT 	public	$y_{VP}^0 = \text{publif}$	#1: publif- publi-publi $\langle p \rangle$ -public
		$y_{Emp}^0 = \langle s \rangle \langle /s \rangle$	#1: $\langle s \rangle \langle /s \rangle$ - $\langle s \rangle \langle /s \rangle$ - $\langle 6p \rangle$ -publif #2: publif- publi-publi $\langle p \rangle$ -public
		$y_{Rand}^0 = \text{publcc}$	#1: publcc- public-public $\langle 2p \rangle$ -publicip #2: publicip- publi-publi $\langle p \rangle$ -public
3: IIIT 	solaris	$y_{VP}^0 = \text{solris}$	#1: solris- soris-so $\langle 2p \rangle$ ris-solaris
		$y_{Emp}^0 = \langle s \rangle \langle /s \rangle$	#1: $\langle s \rangle \langle /s \rangle$ - $\langle s \rangle \langle /s \rangle$ - $\langle 6p \rangle$ -solris #2: solris- soris-so $\langle 2p \rangle$ ris-solaris
		$y_{Rand}^0 = \text{solacis}$	#1: solacis- solais-sola $\langle p \rangle$ is-solaris
4: IC15 	breakfast	$y_{VP}^0 = \text{breakeast}$	#1: breakeast- breakast-break $\langle p \rangle$ ast-breakfast
		$y_{Emp}^0 = \langle s \rangle \langle /s \rangle$	#1: $\langle s \rangle \langle /s \rangle$ - $\langle s \rangle \langle /s \rangle$ - $\langle 9p \rangle$ -breakeast #2: breakeast- breakast-break $\langle p \rangle$ ast-breakfast
		$y_{Rand}^0 = \text{breaufast}$	#1: breaufast- breakfast-brea $\langle p \rangle$ fast-breakfast
5: SVTP 	house	$y_{VP}^0 = \text{houce}$	#1: houce- hou-hou $\langle 2p \rangle$ -house
		$y_{Emp}^0 = \langle s \rangle \langle /s \rangle$	#1: $\langle s \rangle \langle /s \rangle$ - $\langle s \rangle \langle /s \rangle$ - $\langle 5p \rangle$ -houce #2: houce-hou- hou $\langle 2p \rangle$ -house
		$y_{Rand}^0 = \text{houk}$	#1: houk- houk-houk $\langle 2p \rangle$ -houke2 #2: houke2- hou2-hou $\langle p \rangle$ 2-hous2 #3: hous2- hou-hou $\langle 2p \rangle$ -house
6: CUTE 	vacation	$y_{VP}^0 = \text{vagation}$	#1: vagation- vaation-va $\langle p \rangle$ ation-vacation
		$y_{Emp}^0 = \langle s \rangle \langle /s \rangle$	#1: $\langle s \rangle \langle /s \rangle$ - $\langle s \rangle \langle /s \rangle$ - $\langle 8p \rangle$ -vacation
		$y_{Rand}^0 = \text{vacction}$	#1: vacction- vaction-vac $\langle p \rangle$ tion-vacation

when the sequence remains the same as the first iteration. Thus the results may be the same under different maximum iterations. Apparently, the results at #1 are already promising, and more iterations lead to further improvements. Therefore, there is no sophisticated setting for the maximum iteration and the adaptive convergence enables access to high efficiency. Thus, the maximum iteration is set to 3 for convenience. It is crucial to note that the accuracy of y_{Emp}^0 at #1 is close to LevOCR_{VP}, which is only based on visual feature. While the accuracy of y_{Emp}^0 is greatly increased 2.1% at #2 iteration over #1, demonstrating that LevOCR indeed utilizes language knowledge for further text refinement.

4.5 Effectiveness of Levenshtein Transformer

To further explore the significance of Levenshtein Transformer pipeline for text refinement, we construct some variants of LevOCR and these results are shown in Tab. 2. We remove the *deletion* and *insertion* losses in Eq.(10) and replace them with only one classification loss, the performance of LevOCR without LevT is not significantly improved and approximate to pure image prediction. While the proposed LevOCR with LevT obtains obvious 2.4% improvement than pure visual prediction. Besides CNN feature, vision transformer (ViT) [7] with 4×4 patches is utilized for image feature extraction. And the proposed *deletion* and *insertion* actions are learned for text refinement. ViT-based LevOCR can also achieve stable improvement 2.9% than visual prediction. These results demonstrate that LevT is a perfect fit for text refinement and the proposed LevOCR works well with CNN and ViT backbones, resulting in promising generalization.

4.6 Qualitative Analyses

It is essential to qualitatively dissect the iterative process of LevOCR. We select 6 exemplar images for qualitative exhibition, of which the visual predictions are incorrect. Generally, these challenging images are with various types of noises, such as motion blur, Gaussian blur, irregular font, occlusion, curved shape, perspective distortion and low resolution. The iterative refinement processes are elaborately reported in Tab. 3. And different 3 initial sequences as in Sec. 4.3 are adopted to show the amending capability powered by *deletion* and *insertion*.

Deletion action targets at removing wrong letters. As in Img.6, the ‘C’ is recognized as ‘G’ by visual model due to the curved shape. By leveraging both visual and linguistic information, ‘G’ in \mathbf{y}_{VP}^0 tends to be deleted. Similarly, deletion is triggered in \mathbf{y}_{Rand}^0 to remove a redundant ‘C’. Considering placeholder insertion in empty sequence \mathbf{y}_{Emp}^0 , we clearly observe that the inserted number of placeholder equals to the length of GT, indicating that LevOCR can fully comprehend image. When the length of input text is shorter than GT, placeholder insertion try to add the rest of placeholders at the right position. Consequently, placeholder insertion and deletion endows LevOCR with the ability of directly altering sequence length, which is different from previous methods (such as SRN [47] and ABINet [8]).

As for token prediction, it heavily relies on the output of placeholder insertion. When both predicted placeholders and the rest characters are correct, token prediction can directly make right decision (such as Img.4 and Img.6). However, when the output of deletion or placeholder insertion is incorrect, the refinement might collapse. As shown in Img.5 with \mathbf{y}_{Rand}^0 , ‘k’ is not removed and 2 more placeholders are inserted. Fortunately, Img.5 can be successfully corrected after 2 more iterations. This phenomenon demonstrates that *deletion* and *insertion* are well learned by imitation learning, and adversarial learning guarantees that these two actions are complementary and inter-inhibitive. Additionally, token prediction supports inserting not only one character, but also word piece (multiple characters), such as “vic” in Img.1, “la” in Img.3 and “se” in Img.5. These results clearly confirm the superiority of our method.

Table 4. The accuracy comparisons with SOTA methods on 6 public benchmarks. The underlined and bold values represent the second and the best results, respectively.

Methods	Datasets	Regular Text			Irregular Text			Average
		IC13	SVT	IIIT	IC15	SVTP	CUTE	
TBRA [2]	MJ+ST	93.6	87.5	87.9	77.6	79.2	74.0	84.6
ViTSTR [1]	MJ+ST	93.2	87.7	88.4	78.5	81.8	81.3	85.6
ESIR [50]	MJ+ST	91.3	90.2	93.3	76.9	79.6	83.3	87.1
SAM [24]	MJ+ST	95.3	90.6	93.9	77.3	82.2	87.8	88.3
SE-ASTER [33]	MJ+ST	92.8	89.6	93.8	80.0	81.4	83.6	88.3
TextScanner [40]	MJ+ST	92.9	90.1	93.9	79.4	84.3	83.3	88.5
DAN [44]	MJ+ST	93.9	89.2	94.3	74.5	80.0	84.4	87.2
ScRN [46]	MJ+ST	93.9	88.9	94.4	78.7	80.8	87.5	88.4
RobustScanner [48]	MJ+ST	94.8	88.1	95.3	77.1	79.5	90.3	88.4
PIMNet [32]	MJ+ST	95.2	91.2	95.2	83.5	84.3	84.4	90.5
SATRN [22]	MJ+ST	94.1	91.3	92.8	79.0	86.5	87.8	88.6
MASTER [29]	MJ+ST	95.3	90.6	95.0	79.4	84.5	87.5	89.5
SRN [47]	MJ+ST	95.5	91.5	94.8	82.7	85.1	87.8	90.4
ABINet [8]	MJ+ST	97.4	93.5	<u>96.2</u>	<u>86.0</u>	89.3	89.2	<u>92.6</u>
LevOCR	MJ+ST	<u>96.85</u>	<u>92.89</u>	96.63	86.42	<u>88.06</u>	91.67	92.81

4.7 Comparisons with State-of-the-Arts

We compare our LevOCR against thirteen state-of-the-art scene text recognition methods on 6 public benchmarks, and the recognition results are illustrated in Tab. 4. For fair comparison, we only choose the methods that trained on synthetic datasets MJ and ST, and no lexicon is employed for evaluation. Specifically, LevOCR outperforms SATRN [22] and MASTER [29] that are transformer-based encoder-decoder models, showing the effectiveness of BERT-based framework for text recognition. In addition, SRN [47] and ABINet [8] achieve impressive performance by explicitly modelling linguistic information in their methods. The proposed LevOCR gains 2.4% improvement over SRN on average accuracy. Meanwhile, LevOCR surpasses the performance of ABINet on IIIT, IC15 and CUTE datasets and thus achieves the state-of-the-art performance on average accuracy. These results demonstrate the effectiveness of LevOCR.

4.8 Interpretability of LevOCR

Four intuitive examples to demonstrate the good interpretability of LevOCR are shown in Fig. 3. As in Fig. 3 (a), the second character “n” is mis-classified as “m” in the image with GT string “snout”. LevOCR is able to identify that ‘m’ is wrong and thus should be *deleted* (the probability of deletion is 0.87). Beyond that, one can easily examine the reason why LevOCR makes this decision through attention visualization and quantitative comparison. In this case, LevOCR relies more on visual information than on textual information (0.737 *vs.* 0.263) and the pixels near ‘n’ in the image contribute the most for this decision.

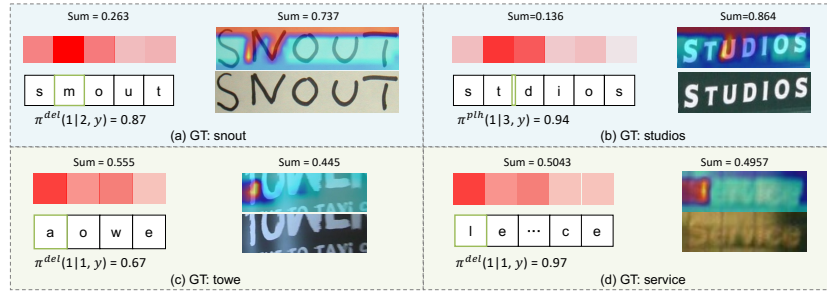


Fig. 3. Illustration of the interpretability of LevOCR. The details (attention maps and intermediate results) of the prediction of LevOCR are depicted. The sum attention weights of different modalities in the top layer are visualized. Best viewed in colors.

Similarly, in the case of Fig. 3 (b), where the GT string is “studios” and ‘u’ is intentionally removed, LevOCR manages to perceive that the third character is missing and thus should be *brought back* (the probability of inserting one character between ‘t’ and ‘d’ is 0.94).

Additionally, when the initial prediction is incorrect and the image is severely corrupted or blurry, LevOCR will pay more attention to textual information in decision making, as shown in Fig. 3 (c) and (d). Through the visualization and numbers produced by LevOCR, we can intuitively comprehend the character-level supporting evidences used by LevOCR in giving this prediction. Notably, good interpretability is a unique characteristic that distinguish the proposed LevOCR from other scene text recognition algorithms. Based on it, we can understand the underlying reason for a specific decision of LevOCR and diagnose its defects. We could even gain insights for designing better systems in the future.

4.9 Limitation

The *deletion* and *insertion* operations in LevOCR are relatively time-consuming, where one refinement iteration costs about 36 ms. Considering the time of visual feature extraction (11 ms) and alternative process, the elapsed time is about 47/83/119 ms for 1/2/3 iteration. The model size (109×10^6 parameters) is relatively large. Therefore, a more efficient architecture will be our future work.

5 Conclusion

In this paper, we have presented an effective and explainable algorithm LevOCR for scene text recognition. Based on Vision-Language Transformer (VLT), LevOCR is able to sufficiently exploit the information from both the vision modality and the language modality to make decisions in an iterative fashion. LevOCR can give fine-granularity (character-level) predictions and exhibits a special property of good interpretability. Extensive experiments verified the effectiveness and advantage of the proposed LevOCR algorithm.

References

1. Atienza, R.: Vision transformer for fast and efficient scene text recognition. In: ICDAR. vol. 12821, pp. 319–334 (2021)
2. Baek, J., Kim, G., Lee, J., Park, S., Han, D., Yun, S., Oh, S.J., Lee, H.: What is wrong with scene text recognition model comparisons? dataset and model analysis. In: ICCV. pp. 4714–4722 (2019)
3. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV. vol. 12346, pp. 213–229 (2020)
4. Chen, X., Jin, L., Zhu, Y., Luo, C., Wang, T.: Text recognition in the wild: A survey. *ACM Computing Surveys (CSUR)* **54**(2), 1–35 (2021)
5. Chen, Y.C., Li, L., Yu, L., El Kholly, A., Ahmed, F., Gan, Z., Cheng, Y., Liu, J.: Uniter: Universal image-text representation learning. In: ECCV. pp. 104–120 (2020)
6. Cheng, Z., Bai, F., Xu, Y., Zheng, G., Pu, S., Zhou, S.: Focusing attention: Towards accurate text recognition in natural images. In: CVPR. pp. 5086–5094 (2017)
7. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houshy, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2021)
8. Fang, S., Xie, H., Wang, Y., Mao, Z., Zhang, Y.: Read like humans: Autonomous, bidirectional and iterative language modeling for scene text recognition. In: CVPR. pp. 7098–7107 (2021)
9. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A.C., Bengio, Y.: Generative adversarial nets. In: NeurIPS. pp. 2672–2680 (2014)
10. Gu, J., Wang, C., Zhao, J.: Levenshtein transformer. In: NeurIPS. pp. 11179–11189 (2019)
11. Gupta, A., Vedaldi, A., Zisserman, A.: Synthetic data for text localisation in natural images. In: CVPR. pp. 2315–2324 (2016)
12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
13. He, P., Huang, W., Qiao, Y., Loy, C.C., Tang, X.: Reading scene text in deep convolutional sequences. In: AAAI. pp. 3501–3508 (2016)
14. He, Y., Chen, C., Zhang, J., Liu, J., He, F., Wang, C., Du, B.: Visual semantics allow for textual reasoning better in scene text recognition. In: AAAI. pp. 888–896 (2022)
15. Hu, W., Cai, X., Hou, J., Yi, S., Lin, Z.: GTC: guided training of CTC towards efficient and accurate scene text recognition. In: AAAI. pp. 11005–11012 (2020)
16. Jaderberg, M., Simonyan, K., Vedaldi, A., Zisserman, A.: Synthetic data and artificial neural networks for natural scene text recognition. NIPS Deep Learning Workshop (2014)
17. Jaderberg, M., Simonyan, K., Vedaldi, A., Zisserman, A.: Reading text in the wild with convolutional neural networks. *Int. J. Comput. Vis.* **116**(1), 1–20 (2016)
18. Karatzas, D., Gomez-Bigorda, L., Nicolaou, A., Ghosh, S.K., Bagdanov, A.D., Iwamura, M., Matas, J., Neumann, L., Chandrasekhar, V.R., Lu, S., Shafait, F., Uchida, S., Valveny, E.: ICDAR 2015 competition on robust reading. In: ICDAR. pp. 1156–1160 (2015)

19. Karatzas, D., Shafait, F., Uchida, S., Iwamura, M., i Bigorda, L.G., Mestre, S.R., Mas, J., Mota, D.F., Almazán, J., de las Heras, L.: ICDAR 2013 robust reading competition. In: ICDAR. pp. 1484–1493 (2013)
20. Kim, W., Son, B., Kim, I.: Vilt: Vision-and-language transformer without convolution or region supervision. In: ICML. vol. 139, pp. 5583–5594 (2021)
21. Lee, C., Osindero, S.: Recursive recurrent nets with attention modeling for OCR in the wild. In: CVPR. pp. 2231–2239 (2016)
22. Lee, J., Park, S., Baek, J., Oh, S.J., Kim, S., Lee, H.: On recognizing texts of arbitrary shapes with 2d self-attention. In: CVPR Workshops. pp. 2326–2335 (2020)
23. Li, M., Lv, T., Cui, L., Lu, Y., Florêncio, D.A.F., Zhang, C., Li, Z., Wei, F.: Trocr: Transformer-based optical character recognition with pre-trained models. CoRR **abs/2109.10282** (2021)
24. Liao, M., Lyu, P., He, M., Yao, C., Wu, W., Bai, X.: Mask textspotter: An end-to-end trainable neural network for spotting text with arbitrary shapes. IEEE Trans. Pattern Anal. Mach. Intell. **43**(2), 532–548 (2021)
25. Liao, M., Zhang, J., Wan, Z., Xie, F., Liang, J., Lyu, P., Yao, C., Bai, X.: Scene text recognition from two-dimensional perspective. In: AAAI. pp. 8714–8721 (2019)
26. Liu, H., Wang, B., Bao, Z., Xue, M., Kang, S., Jiang, D., Liu, Y., Ren, B.: Perceiving stroke-semantic context: Hierarchical contrastive learning for robust scene text recognition. In: AAAI. pp. 1702–1710 (2022)
27. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. CoRR **abs/2103.14030** (2021)
28. Long, S., He, X., Yao, C.: Scene text detection and recognition: The deep learning era. IJCV **129**(1), 161–184 (2021)
29. Lu, N., Yu, W., Qi, X., Chen, Y., Gong, P., Xiao, R., Bai, X.: MASTER: multi-aspect non-local network for scene text recognition. PR **117**, 107980 (2021)
30. Mishra, A., Alahari, K., Jawahar, C.V.: Scene text recognition using higher order language priors. In: BMVC. pp. 1–11 (2012)
31. Phan, T.Q., Shivakumara, P., Tian, S., Tan, C.L.: Recognizing text with perspective distortion in natural scenes. In: ICCV. pp. 569–576 (2013)
32. Qiao, Z., Zhou, Y., Wei, J., Wang, W., Zhang, Y., Jiang, N., Wang, H., Wang, W.: Pimnet: A parallel, iterative and mimicking network for scene text recognition. In: ACM MM. pp. 2046–2055 (2021)
33. Qiao, Z., Zhou, Y., Yang, D., Zhou, Y., Wang, W.: SEED: semantics enhanced encoder-decoder framework for scene text recognition. In: CVPR. pp. 13525–13534 (2020)
34. Risnumawan, A., Shivakumara, P., Chan, C.S., Tan, C.L.: A robust arbitrary text detection system for natural scene images. Expert Syst. Appl. **41**(18), 8027–8048 (2014)
35. Sheng, F., Chen, Z., Xu, B.: NRTR: A no-recurrence sequence-to-sequence model for scene text recognition. In: ICDAR. pp. 781–786 (2019)
36. Shi, B., Bai, X., Yao, C.: An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. IEEE TPAMI **39**(11), 2298–2304 (2017)
37. Shi, B., Yang, M., Wang, X., Lyu, P., Yao, C., Bai, X.: ASTER: an attentional scene text recognizer with flexible rectification. IEEE TPAMI **41**(9), 2035–2048 (2019)
38. Su, W., Zhu, X., Cao, Y., Li, B., Lu, L., Wei, F., Dai, J.: VL-BERT: pre-training of generic visual-linguistic representations. In: ICLR (2020)

39. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NeurIPS. pp. 5998–6008 (2017)
40. Wan, Z., He, M., Chen, H., Bai, X., Yao, C.: Textscanner: Reading characters in order for robust scene text recognition. In: AAAI. pp. 12120–12127 (2020)
41. Wan, Z., Xie, F., Liu, Y., Bai, X., Yao, C.: 2d-ctc for scene text recognition. arXiv preprint arXiv:1907.09705 (2019)
42. Wan, Z., Zhang, J., Zhang, L., Luo, J., Yao, C.: On vocabulary reliance in scene text recognition. In: CVPR. pp. 11422–11431 (2020)
43. Wang, K., Babenko, B., Belongie, S.J.: End-to-end scene text recognition. In: ICCV. pp. 1457–1464 (2011)
44. Wang, T., Zhu, Y., Jin, L., Luo, C., Chen, X., Wu, Y., Wang, Q., Cai, M.: Decoupled attention network for text recognition. In: AAAI. pp. 12216–12224 (2020)
45. Wang, Y., Xie, H., Fang, S., Wang, J., Zhu, S., Zhang, Y.: From two to one: A new scene text recognizer with visual language modeling network. In: ICCV. pp. 1–10 (2021)
46. Yang, M., Guan, Y., Liao, M., He, X., Bian, K., Bai, S., Yao, C., Bai, X.: Symmetry-constrained rectification network for scene text recognition. In: ICCV. pp. 9146–9155 (2019)
47. Yu, D., Li, X., Zhang, C., Liu, T., Han, J., Liu, J., Ding, E.: Towards accurate scene text recognition with semantic reasoning networks. In: CVPR. pp. 12110–12119 (2020)
48. Yue, X., Kuang, Z., Lin, C., Sun, H., Zhang, W.: Robustscanner: Dynamically enhancing positional clues for robust text recognition. In: ECCV. vol. 12364, pp. 135–151 (2020)
49. Zeiler, M.D.: ADADELTA: an adaptive learning rate method. CoRR **abs/1212.5701** (2012)
50. Zhan, F., Lu, S.: ESIR: end-to-end scene text recognition via iterative image rectification. In: CVPR. pp. 2059–2068 (2019)
51. Zhang, X., Zhu, B., Yao, X., Sun, Q., Li, R., Yu, B.: Context-based contrastive learning for scene text recognition. In: AAAI. pp. 888–896 (2022)
52. Zhu, Y., Yao, C., Bai, X.: Scene text detection and recognition: Recent advances and future trends. *Frontiers of Computer Science* **10**(1), 19–36 (2016)