

Masked Vision-Language Transformers for Scene Text Recognition

Jie Wu

wu.jie@westone.com.cn

Ying Peng

peng.ying06059@westone.com.cn

Shengming Zhang

shmizhang@gmail.com

Weigang Qi

qi.weigang@westone.com.cn

Jian Zhang

zhang.jian06086@westone.com.cn

Westone Information Industry INC.

Chengdu, China

Abstract

Scene text recognition (STR) enables computers to recognize and read the text in various real-world scenes. Recent STR models benefit from taking linguistic information in addition to visual cues into consideration. We propose a novel Masked Vision-Language Transformers (MVLT) to capture both the explicit and the implicit linguistic information. Our encoder is a Vision Transformer, and our decoder is a multi-modal Transformer. MVLT is trained in two stages: in the first stage, we design a STR-tailored pretraining method based on a masking strategy; in the second stage, we fine-tune our model and adopt an iterative correction method to improve the performance. MVLT attains superior results compared to state-of-the-art STR models on several benchmarks. Our code and model are available at <https://github.com/onealwj/MVLT>.

1 Introduction

Scene text recognition (STR) aims to read text from natural scenes, which is helpful in many practical artificial intelligence applications such as autonomous driving, instant translation, and natural scene understanding. STR has been studied extensively in the past two decades, however, the performance of which still struggles under scenarios with unpromising illuminations, occluded characters, complex deformations, etc. Recent studies [0, 8, 16, 17, 55, 59, 42] have made progress in dealing with such challenges by introducing textual semantics information except for visual cues. These language-aware methods are typically divided into two kinds: explicitly building an extra language model or implicitly extracting textual semantics from visual cues. The former [8, 16, 17, 42] rely on language models such as n-grams [17] or attention-based neural networks [42] to predict word-level text. The latter [0, 55, 59] attempt to guide the model to catch textual semantics according to visual context without using additional language models. For instance, [59] exploits a character-level

occluded strategy to make the visual model learn linguistic information along with visual features. Despite the effectiveness of both ways, each of them only captures textual semantics from a single orientation. Inspired by their success, we propose a method to learn both the explicit and the implicit textual semantics, promoting the ability of STR.

We propose a Masked Vision-Language Transformers (MVLT) for STR, with a Vision Transformer (ViT) [6] encoder and a creatively designed multi-modal Transformer [23, 34, 43] decoder. The potential of Transformers [57] has been proved in the area of NLP and CV, while the pretraining and fine-tuning pipelines of Transformer-based models further boost the performance of down-streaming tasks. To this end, we adopt a two-stage training strategy to train the model as follows:

In the first stage, we pretrain the model by borrowing the idea from masked autoencoders (MAE) [10]. MAE splits an image into several patches, randomly masks a percentage of them, and learns to reconstruct the masked ones. Different from MAE, our MVLT recognizes scene text in addition to reconstructing the masked patches. To recognize text, we seek to use linguistic information to assist the visual cues. Motivated by multi-modal Transformers, which combine image regions with language semantics through a Transformer and learn the interaction between the two, we build a multi-modal Transformer decoder to bring linguistic information into our model. For the input of our multi-modal decoder, the visual part consists of encoded patches and mask tokens; the textual part comes from the ground-truth text label of the corresponding image and is formed into a sequence of character embeddings. We build two sub-decoders to model explicit and implicit textual semantics, respectively. Similar to MLM in VisualBERT [23], which masks some tokens in the textual input and predicts them according to the visual and textual features, for one sub-decoder, we mask a part of characters in the ground-truth text label and learn to predict the masked characters. The visible characters serve as word-level textual cues, explicitly guiding the model to learn linguistic knowledge. To endow the model with the ability to predict the correct word-level text even without explicit textual input, which is closer to the application scenario of STR, for another sub-decoder, we mask all input characters. The model is now left with only visual cues and thus is pushed to learn linguistic information implicitly without an additional language model. It is worth noticing that the two sub-decoders share parameters during training for efficiency. Furthermore, as an extension of our MVLT, we propose a simple but efficient method to use real-world unlabeled data together with labeled one in pretraining.

In the second stage, we fine-tune the model, where the encoder takes the unmasked scene text image as input, and the decoder outputs the predicted text. To better use the pretrained knowledge, different from MAE or ViTSTR [2], which only fine-tune on the pretrained encoder, we fine-tune on both the encoder and the decoder. Meanwhile, we propose an iterative correction method to gradually modify the predicted text during iterations. In each iteration, the input of the decoder consists of the image feature that is output from encoder, and the text feature that is output from the previous iteration.

Our contributions are three-folds: (1) we propose a novel language-aware model for STR, allowing for learning both the explicit and the implicit semantics, which gets superior accuracy compared to previous works; (2) we design a masking-based strategy in pretraining the model and exploit an iterative correction method to correct text predictions in the fine-tuning stage. (3) we propose a method to further promote accuracy by leveraging labeled and unlabeled training data.

2 Related Work

One of the classical ways for STR is to extract visual features based on CNN, use RNN to perform sequence labeling, and adopt Connectionist Temporal Classification (CTC) [9] as a loss function [12, 32]. Recently, GTC [14] optimized the CTC-based method by using a graph convolutional network (GCN) [20] to learn the local correlations of features. Some works attempt to rectify the irregular text images [33, 41, 42]. For example, ASTER [33] adopts Thin-Plate-Spline (TPS) [5] transformation, and ScRN [41] adds symmetrical constraints in addition to the TPS. Recent works [1, 16, 22, 27, 31, 40] brought insightful ideas in dealing with challenging scenarios (occlusion, noise, etc.) by building language-aware models.

Language-aware methods. Some methods [8, 17, 27, 42] rely on external language models to extract the semantics. For example, SRN [22] builds a global semantic reasoning module, which learns semantics based on the predicted text from the visual model. This method is upgraded in ABINet [8], which develops a bidirectional cloze network to make better use of bidirectional linguistic information, and utilizes an iterative correction for the language model. Some other works [0, 35, 39] implicitly learn semantics without using language models. For example, SVTR [0] recognizes both the characters and the inter-character long-term dependence in a single visual model by using local and global mixing blocks. VST [35] extracts semantics from a visual feature map and performs a visual-semantic interaction using an alignment module. Most recently, some language-aware models explore more possibilities to boost the performance by considering spatial context [13], adding real-world images to train the model in a semi-supervised way [0], or developing a re-ranking method to get a better candidate output [30], while our work takes a step by extracting both the explicit and the implicit semantics.

Transformer-based STR. Recently, Transformer has shown its effectiveness in STR models [8, 28, 31, 39, 40, 42]. For instance, PIMNet [28] builds a bi-directional Transformer-based parallel decoder to iteratively capture context information. Lately, based on ViT structure [6], ViTSTR [2] uses a ViT encoder to perform STR without using the decoder. Additionally, ViTSTR is initialized by pretrained parameters of DeiT [56]. We also apply a ViT encoder in our model. However, unlike ViTSTR, we take the decoder into our architecture and propose a new pretraining method to better fit the STR task.

3 Methodology

Our proposed MVLT is trained with two stages, including a pretraining stage, as shown in Figure 1, and a fine-tuning stage, as shown in Figure 2 (a). In addition, we harness the unlabeled data to enhance the training of the model, as shown in Figure 2 (b) and (c).

3.1 Preliminary

Vision Transformer (ViT). To model 2D image into Transformer, which is designed to process 1D sequence, ViT [6] flattens an image by reshaping the image $x \in \mathbb{R}^{H \times W \times C}$ into a sequence of image patches $x \in \mathbb{R}^{N \times (P^2 \cdot C)}$, where (H, W) denotes the original scale of the image, C is the number of channels, and (P, P) is the resolution of each image patch.

Masked autoencoders (MAE). MAE [10] randomly divides N image patches into N_u unmasked ones and N_m masked ones. Correspondingly, $x = [x_u; x_m] \in \mathbb{R}^{N \times (P^2 \cdot C)}$, where x is

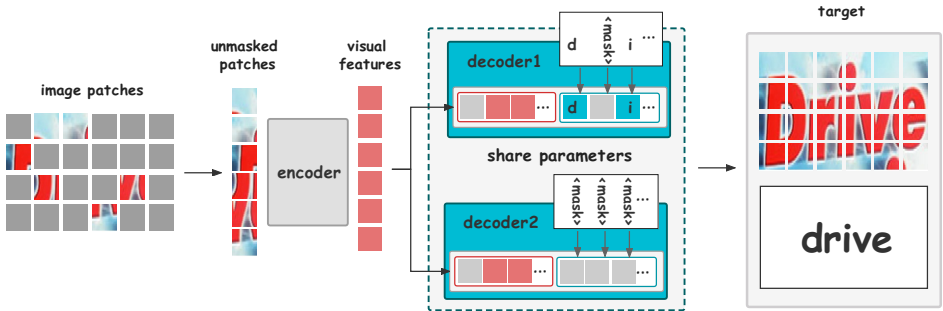


Figure 1: An illustration of MVLT in the pretraining stage.

the full set of image patches, x_u is the set of unmasked patches, and x_m is the set of masked patches. The encoder of MAE is a ViT, which only operates on x_u to learn the visual feature embeddings:

$$v_u = \text{encoder}(x_u), \quad (1)$$

where $v_u \in \mathbb{R}^{N_u \times D_1}$ and D_1 is the feature dimension in the encoder. The mask token v_m is introduced after the encoder, constituting the input of the decoder, together with v_u . The decoder reconstructs the image as follows:

$$\hat{v}_m, \hat{v}_u = \text{decoder}(v_m, v_u), \quad (2)$$

where $v_m \in \mathbb{R}^{N_u \times D_2}$ and D_2 is the feature dimension in the decoder. $\hat{v}_m \in \mathbb{R}^{N_m \times (P^2 \cdot C)}$ and $\hat{v}_u \in \mathbb{R}^{N_u \times (P^2 \cdot C)}$. The Mean Squared Error (MSE) loss is used to optimize MAE model.

3.2 Pretraining Stage

In pretraining, MVLT intends to reconstruct the masked image patches and recognize the scene text from the masked image. Like MAE, the reconstruction of image patches helps our model to learn an effective visual representation. However, recognizing texts is beyond the scope of MAE, and we propose a novel language-aware model to deal with it, as shown in Figure 1.

Masked Encoder. We use ViT as the encoder of MVLT, which is same as Eq. (1). Each image is split into N patches. With a mask ratio of 0.75, we divide the patches into a set of masked ones $x_m \in \mathbb{R}^{N_m \times (P^2 \cdot C)}$, and a set of the unmasked ones $x_u \in \mathbb{R}^{N_u \times (P^2 \cdot C)}$. The encoder embeds x_u by using it as the input of a linear layer and then adding the positional embeddings with the output of the linear layer. Next, the encoder uses a series of Transformer blocks to learn the visual embedding: $v_u = \text{encoder}(x_u)$.

Masked Decoder. We build a multi-modal Transformer as the decoder to exploit the visual cues and linguistic information in each image. The visual cues come from v_u . The linguistic information comes from the word-level text label that is character-wise mapped into a sequence of learnable character embeddings $t \in \mathbb{R}^{L \times D_2}$, where L is the length of the character embedding sequence. Similar to the image patches, we denote $t = [t_u, t_m]$, where $t_u \in \mathbb{R}^{L_u \times D_2}$ is the unmasked character embeddings, and $t_m \in \mathbb{R}^{L_m \times D_2}$ is a sequence of the

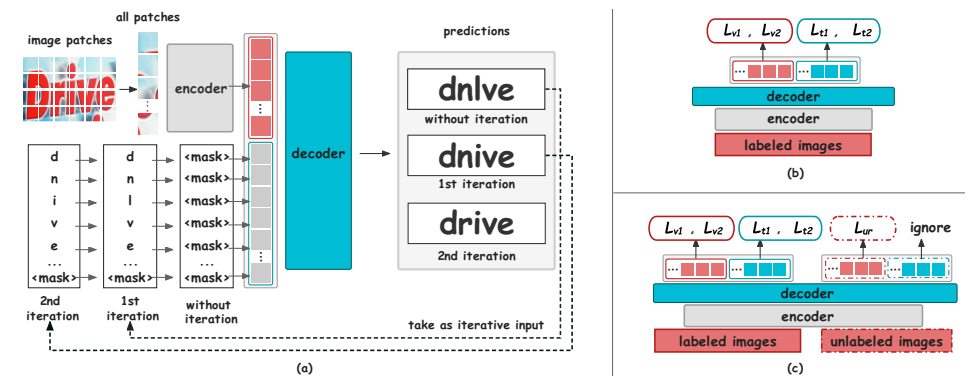


Figure 2: Sub-figure (a) shows our model in the fine-tuning stage. Sub-figure (b) illustrates our model using labeled data in the pretraining stage. Sub-figure (c) illustrates our model using both labeled and unlabeled data in the pretraining stage.

special "<mask>" token embeddings, L_u and L_m are the corresponding length. We add positional embeddings to (v_m, v_u, t_m, t_u) to build the decoder input. For symbolic simplification, we keep the symbols before and after adding the positional embeddings the same. We design two sub-decoders ($decoder_1$ and $decoder_2$), applying different masking strategies on t to learn textual semantics. Note that the two sub-decoders share parameters during pretraining. Our decoder is denoted as:

$$\hat{v}_m, \hat{v}_u, \hat{t}_m, \hat{t}_u = \text{decoder}(v_m, v_u, t_m, t_u), \quad (3)$$

where $\hat{v}_m \in \mathbb{R}^{N_m \times (P^2 \cdot C)}$, $\hat{v}_u \in \mathbb{R}^{N_u \times (P^2 \cdot C)}$, $\hat{t}_m \in \mathbb{R}^{L_m \times M}$, $\hat{t}_u \in \mathbb{R}^{L_u \times M}$, and M is the number of character's category.

Modeling explicit language semantics. We set the text mask ratio of $decoder_1$ to 0.2. For instance, if the text label consists of 10 characters, we randomly mask 2 of them. The length of unmasked character embeddings $L_u = 8$, and the length of masked ones $L_m = L - L_u$. With the unmasked character embeddings t_u serving as linguistic context, $decoder_1$ explicitly learns textual semantics.

Modeling implicit language semantics. For $decoder_2$, we set the mask ratio of the character embeddings as 1.0. With this setting, the length of unmasked character embeddings $L_u = 0$, and the masked ones $L_m = L$. As the characters are totally masked, $decoder_2$ only uses the visual information to reconstruct \hat{v}_m and predicts the word-level text label \hat{t}_m , which pushes it to learn implicit textual semantics from visual cues.

Pretraining objective. We use the MSE loss to optimize the reconstructed image patches and use the Cross Entropy Loss to optimize textual prediction:

$$\mathcal{L}_{\text{pretraining}} = \alpha \cdot \mathcal{L}_{v_1} + \beta \cdot \mathcal{L}_{v_2} + \gamma \cdot \mathcal{L}_{t_1} + \varepsilon \cdot \mathcal{L}_{t_2}, \quad (4)$$

where $\{\alpha, \beta, \gamma, \varepsilon\}$ are trade-off parameters. For $decoder_1$, $\mathcal{L}_{v_1} = \text{MSE}(\hat{v}_m, y_m)$, where $y_m \in \mathbb{R}^{N_m \times (P^2 \cdot C)}$ denotes the pixel values of the masked image patches. $\mathcal{L}_{t_1} = \text{Cross-Entropy}(\hat{t}_m, y_t)$, where $y_t \in \mathbb{R}^{L \times 1}$ is the ground-truth character index, and each element in $y_t \in [1, M]$. For $decoder_2$, $\mathcal{L}_{v_2} = \text{MSE}(\hat{v}_m, y_m)$, and $\mathcal{L}_{t_2} = \text{Cross-Entropy}(\hat{t}_m, y_t)$.

3.3 Fine-tuning Stage

We fine-tune the above pre-trained encoder and decoder to further promote the performance of our model on the STR task. Motivated by the iterative strategy of ABINet [8], we design an iterative correction method that fits the architecture of our model, as shown in Figure 2 (a). This iterative correction method is an alternative during fine-tuning.

Encoder. The full set of image patches $x \in \mathbb{R}^{N \times (P^2 \cdot C)}$ is taken as the input of the encoder without utilizing the masking operation:

$$v = \text{encoder}(x), \quad (5)$$

where $v \in \mathbb{R}^{N \times D_1}$ and N is the number of image patches.

Decoder. In the pretraining stage, our encoder takes the output of the encoder (visual-related feature) and a sequence of character embeddings (linguistic-related feature) as input. However, in the fine-tuning stage, when the iterative correction method is not used or before the first iteration, only the visual-related feature is visible by the decoder, without the linguistic-related feature. Thus, to be consistent with the input of the pretrained decoder, we use a sequence of "<mask>" token embeddings as the character embedding, $t_m \in \mathbb{R}^{L \times D_2}$. Different from the pretraining stage, where the output of the decoder consists of reconstructed image patches and the predicted text, during the fine-tuning stage, the predicted image patches are ignored, and only the predicted text is kept because the STR task only focuses on the text prediction. The input and output of the decoder are formalized as follows:

$$\hat{t} = \text{decoder}(v, t_m), \quad (6)$$

where $\hat{t} \in \mathbb{R}^{L \times M}$ is the logits of predicted text.

Iterative correction. We regard the output \hat{t} as a raw text prediction that will be corrected during iterations. The prediction probability of each character is computed based on \hat{t} . A linear projection layer takes the prediction probabilities as the input. It outputs a sequence of character-related feature representations, which is set as the new character embedding input at the current iteration. For example, we perform K times of iterative corrections by the follows: for the 1st iteration, we pass $\hat{t}_{\text{itr}=0} = \hat{t}$ into a softmax layer and then a linear layer to get $t_{\text{itr}=1}$, which is taken as the new character embedding to replace the totally masked character embeddings t_m in Eq. (6). Then, the k -th ($k \in [1, K]$) iteration process is formalized as:

$$\text{prob}_{\text{itr}=k} = \text{softmax}(\hat{t}_{\text{itr}=k-1}), \quad (7)$$

$$t_{\text{itr}=k} = \text{linear}(\text{prob}_{\text{itr}=k}), \quad (8)$$

$$\hat{t}_{\text{itr}=k} = \text{decoder}(v, t_{\text{itr}=k}), \quad (9)$$

The output of the K -th iteration is regarded as the final corrected text prediction.

Fine-tuning objective. Different from the pretraining objective, we only focus on optimizing the textual prediction by leveraging the Cross Entropy Loss:

$$\mathcal{L}_{\text{fine-tuning}} = \frac{1}{2} \text{Cross-Entropy}(\hat{t}_{\text{itr}=0}, y_t) + \frac{1}{2(K-1)} \sum_{j=1}^K \text{Cross-Entropy}(\hat{t}_{\text{itr}=j}, y_t). \quad (10)$$

3.4 Using Unlabeled Real Dataset

Our pretraining objective includes: (a) reconstructing the masked image patches; and (b) recognizing the scene texts. Since (b) requires leveraging labeled scene text to attain supervised training, we use labeled synthetic datasets for training. However, there is a mass of real-world data, which is unlabeled and costs much to label. Fortunately, we find it easy to use unlabeled data to enhance learning features during pretraining. Specifically, as shown in Figure 2 (c), we concatenate the unlabeled image data and the labeled image data along the batch dimension as a new input batch. For a batch of output $O = \{O_{syn}, O_{ur}\}$, $O_{syn} = \{(\hat{v}_m, \hat{v}_u, \hat{t}_m, \hat{t}_u)_1, \dots, (\hat{v}_m, \hat{v}_u, \hat{t}_m, \hat{t}_u)_{N_1}\}$ is the output corresponding to the labeled synthetic datasets, and $O_{ur} = \{(\hat{v}'_m, \hat{v}'_u, \hat{t}'_m, \hat{t}'_u)_1, \dots, (\hat{v}'_m, \hat{v}'_u, \hat{t}'_m, \hat{t}'_u)_{N_2}\}$ is the output corresponding to the unlabeled real datasets, where N_1 is the number of labeled synthetic data in the batch, and N_2 is the number of unlabeled real data in the batch. O_{syn} is optimized on the loss function of Eq. (4), while O_{ur} is optimized only on a MSE loss $\mathcal{L}_{ur} = \text{MSE}(\hat{v}'_m, y'_m)$, \hat{v}'_m is the output corresponding to the masked patches of the real data, and y'_m is the target pixels of the real data. Through this semi-supervised pretraining process, we can learn real-world domain knowledge, thus reducing the gap between the pretraining data and the real-world scenario.

4 Experiment

Datasets. To conduct supervised training, we use two synthetic datasets, MJSynth (MJ) [14], [18] and SynthText (ST) [10], as the labeled training dataset of MVLT. We use all of the 14 real datasets that collected in [9], and remove all the image labels to build our unlabeled real dataset. We denote the unlabeled real dataset as UR for a simplified description. We use the same test dataset as ABINet [8], including six standard benchmarks, which consists of 857 images from ICDAR2013 (IC13) [19], 1,811 images from ICDAR2015 (IC15) [20], 647 images from Street View Text (SVT) [38], 645 images from SVT-Perspective (SVTP) [26], 3,000 images from IIIT 5-K Words (IIIT) [25], and 288 images from CUTE80 (CUTE) [29].

Model detail. Images are scaled to 112×448 , with a resolution of 14×14 for each patch. Although the image size is different from ViT [9], we keep the number of patches the same. Our encoder uses the same settings as ViT-B in MAE [10]. We use a lightweight decoder, which has depth 4, width 512, and 8 attention heads. Thus the dimension D_1 and D_2 are set to 768 and 512, respectively. The length of the character embedding sequence, L , is set to 27, because, according to our observation, the vast majority of the word is shorter than 27 characters. In the pretraining stage, to build the input of $decoder_1$, we pad the text labels that are shorter than L by a special token "`<mask>`". The input of $decoder_2$ is built from L "`<mask>`" tokens.

Training detail. The trade-off parameters α and β are set to 0.5, γ and ε are set to 0.01. In both the pretraining and the fine-tuning stage, an AdamW [24] optimizer and a cosine learning rate decay scheduler are applied. In the pretraining stage, we set the initial learning rate to $1.5e-4$ and weight decay to 0.05. When using only the labeled data, the batch size is set to 4,096. When using both the labeled and the unlabeled data, the batch size is 6,144, with 4,096 labeled images and 2,048 unlabeled ones. We conduct a total of 120,000 iterations, with 8,000 warm-up iterations. The optimizer momentum is set to $\beta_1=0.9$ and $\beta_2=0.95$. We do not use grad clip during pretraining. In the fine-tuning stage, the batch size is 1,024, the initial learning rate is $1e-5$, and the weight decay is 0.05. We conduct a total of 20,000

Table 1: Accuracy results of our MVLT and SOTA methods on six regular and irregular STR datasets. "UR" denotes the unlabeled real dataset.

Method	Datasets	Regular Text			Irregular Text		
		IC13	SVT	IIIT	IC15	SVTP	CUTE
ASTER [15]	MJ+ST	91.8	89.5	93.4	76.1	78.5	79.5
ESIR [16]	MJ+ST	91.3	90.2	93.3	76.9	79.6	83.3
ScRN [17]	MJ+ST	93.9	88.9	94.4	78.7	80.8	87.5
PIMNet [18]	MJ+ST	93.4	91.2	95.2	81.0	84.3	84.4
SAR [19]	MJ+ST	94.0	91.2	95.0	78.8	86.4	89.6
SRN [20]	MJ+ST	95.5	91.5	94.8	82.7	85.1	87.8
GTC [21]	MJ+ST	94.3	92.9	95.5	82.5	86.2	92.3
VisionLAN [35]	MJ+ST	95.7	91.7	95.8	83.7	90.7	88.5
PREN2D [10]	MJ+ST	96.4	94.0	95.6	83.0	87.6	91.7
S-GTR [9]	MJ+ST	96.8	94.1	95.8	84.6	87.9	92.3
ABINet [8]	MJ+ST	97.4	93.5	96.2	86.0	89.3	89.2
MVLT	MJ+ST	97.3	94.7	96.8	87.2	90.9	91.3
MVLT*	MJ+ST+UR	98.0	96.3	97.4	89.0	92.7	95.8

iterations, with 8,000 warm-up iterations. The optimizer momentum is set to $\beta_1=0.9$ and $\beta_2=0.999$. The grad clip is set to 2.0. The layer-wise learning rate decay is set to 0.75. We perform 3 times of iterative corrections in fine-tuning the model and 3 times of iterative corrections in testing the model. We use 8 NVIDIA RTX A6000 GPUs, with 48GB memory, to conduct the experiments, and use gradient accumulation to maintain a large effective batch size. The pretraining stage takes around 3.5 days, and the fine-tuning stage takes around 5 hours.

Data augmentation. In the pretraining stage, we apply RandomResizedCrop to augment data, which is similar to MAE. Specifically, we set the scale to (0.85, 1.0) and the ratio to (3.5, 5.0). In the fine-tuning stage, we use the same data augmentation method as ABINet, including rotation, affine, and perspective.

4.1 Performance Analysis

Comparisons with state-of-the-arts. Table 1 summarizes the comparison result of our proposed method with eleven SOTA. MVLT is trained on ST and MJ, and MVLT* uses extra unlabeled data. MVLT achieves the highest accuracy in most of the datasets. The result proves the effectiveness of our design of model architecture and our two-stage training procedure. Benefiting from learning linguistic knowledge, MVLT is more tolerant to irregular texts, leading to a higher performance even compared with the methods that use rectification modules, such as ASTER, ESIR, and GTC. Compared with other language-aware models, including VisionLAN, ABINet, PREN2D, etc., our model performs better on the vast majority of the test datasets, further justifying the capture of textual semantics. To perform a fair comparison, MVLT uses the same data and the data augmentation method as ABINet in the fine-tuning stage. MVLT outperforms ABINet with 1.2%, 0.6%, 1.2%, 1.6%, and 2.1% on SVT, IIIT, IC15, SVTP, and CUTE datasets, respectively.

Using unlabeled real dataset. As shown in Table 1, MVLT* outperforms MVLT on most test datasets. Specifically, MVLT* largely promotes the result of recognizing irregular texts, showing that using unlabeled real data is able to enhance the ability of STR, leading to a more practicable model facing the real-world scenario.

Iterative correction. Iterative correction is applied during training the model in the fine-




Figure 3: Accuracy of iterative correction in the training and testing process.

tuning stage, as well as testing. Figure 3 shows the results with different settings of the iterate times, which suggests that the iterative correction method can help obtain higher accuracy.

4.2 Ablation Study

Table 2: The results of ablation study. "Iter" represents using iterative correction method in training the model in the fine-tuning stage, and in testing on the test datasets.

\mathcal{L}_{v_1}	\mathcal{L}_{t_1}	\mathcal{L}_{v_2}	\mathcal{L}_{t_2}	Iter	Regular Text			Irregular Text			Total
					IC13	SVT	IIIT	IC15	SVTP	CUTE	
✓	✓	✓	✓	✓	97.3	94.7	96.8	87.2	90.9	91.3	93.5
✓	✓	✓	✓	✓	97.0	94.1	96.8	86.6	89.6	90.6	93.2
		✓	✓	✓	96.8	94.4	96.3	87.0	89.9	90.6	93.1
		✓	✓		96.7	94.6	96.3	86.7	90.2	91.0	93.1
✓	✓			✓	96.4	92.4	96.4	86.7	89.3	91.0	92.8
✓	✓				95.9	92.7	96.0	86.3	88.7	89.6	92.4
✓					96.4	91.3	96.5	85.5	88.8	89.9	92.3

The statistics in Table 2 show that: 1) Using only the visual-related loss \mathcal{L}_{v_1} is not as effective as using both the visual-related and the textual-related losses. 2) Losing either the loss related to learning explicit semantics or the loss related to learning implicit semantics will result in a decrease in accuracy. 3) Learning the implicit textual semantics leads to a higher leap in accuracy than learning the explicit semantics. 4) The iterative correction is especially useful when the model is trained with explicit textual semantics.

4.3 Visualization and Analysis

Figure 4 and Figure 5 display several visualization results after the pretraining and fine-tuning stage. The images are from test datasets. Figure 4 suggests that the model has already acquired visual and linguistic semantic knowledge through the pretraining stage. Figure 5 shows a stronger model after fine-tuning, which is capable of dealing with more challenging real-world images. More examples are shown in the supplementary material.



Figure 4: Visualization of pretrained MVLT. For each example, we show the masked image (left), our image reconstruction and text prediction from $decoder_1$ (mid-left) and $decoder_2$ (mid-right), and the ground-truth (right).



Figure 5: Visualization of fine-tuned MVLT. We show successfully recognized difficult examples with occlusion (1st row), complex font styles (2nd row), blur (3rd row), and deformation (4th row).

5 Conclusion

We propose a Masked Vision-Language Transformers (MVLT) for STR, getting superior results compared to the state-of-the-art models. For pretraining the model, we design a masking strategy to lead the model in learning both the explicit and implicit textual semantics. Experiment results have proved the effectiveness of our model in capturing semantics and the usefulness of both kinds of semantics. During training and testing our model in the fine-tuning stage, we apply an iterative correction method, which boosts the performance of STR. Furthermore, the use of unlabeled real data improves the applicability of our model in real-world scenarios.

References

- [1] Aviad Aberdam, Roy Ganz, Shai Mazor, and Ron Litman. Multimodal semi-supervised learning for text recognition. *arXiv preprint arXiv:2205.03873*, 2022.
- [2] Rowel Atienza. Vision transformer for fast and efficient scene text recognition. In Josep Lladós, Daniel Lopresti, and Seiichi Uchida, editors, *16th International Conference on Document Analysis and Recognition, ICDAR 2021, Lausanne, Switzerland, September 5-10, 2021, Proceedings, Part I*, volume 12821 of *Lecture Notes in Computer Science*, pages 319–334. Springer, 2021.

- [3] Jeonghun Baek, Yusuke Matsui, and Kiyoharu Aizawa. What if we only use real datasets for scene text recognition? toward scene text recognition with fewer labels. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 3113–3122. Computer Vision Foundation / IEEE, 2021.
- [4] Ayan Kumar Bhunia, Aneeshan Sain, Amandeep Kumar, Shuvojit Ghose, Pinaki Nath Chowdhury, and Yi-Zhe Song. Joint visual semantic reasoning: Multi-stage decoder for text recognition. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 14920–14929. IEEE, 2021.
- [5] Fred L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Mach. Intell.*, 11(6):567–585, 1989.
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.
- [7] Yongkun Du, Zhineng Chen, Caiyan Jia, Xiaoting Yin, Tianlun Zheng, Chenxia Li, Yuning Du, and Yu-Gang Jiang. SVTR: scene text recognition with a single visual model. In Luc De Raedt, editor, *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI 2022, Vienna, Austria, 23-29 July 2022*, pages 884–890. ijcai.org, 2022.
- [8] Shancheng Fang, Hongtao Xie, Yuxin Wang, Zhendong Mao, and Yongdong Zhang. Read like humans: Autonomous, bidirectional and iterative language modeling for scene text recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 7098–7107. Computer Vision Foundation / IEEE, 2021.
- [9] Alex Graves, Santiago Fernández, Faustino J. Gomez, and Jürgen Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In William W. Cohen and Andrew W. Moore, editors, *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006*, volume 148 of *ACM International Conference Proceeding Series*, pages 369–376. ACM, 2006.
- [10] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. Synthetic data for text localisation in natural images. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 2315–2324. IEEE Computer Society, 2016.
- [11] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022.

- [12] Pan He, Weilin Huang, Yu Qiao, Chen Change Loy, and Xiaoou Tang. Reading scene text in deep convolutional sequences. In Dale Schuurmans and Michael P. Wellman, editors, *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 3501–3508. AAAI Press, 2016.
- [13] Yue He, Chen Chen, Jing Zhang, Juhua Liu, Fengxiang He, Chaoyue Wang, and Bo Du. Visual semantics allow for textual reasoning better in scene text recognition. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 888–896. AAAI Press, 2022.
- [14] Wenyang Hu, Xiacong Cai, Jun Hou, Shuai Yi, and Zhiping Lin. GTC: guided training of CTC towards efficient and accurate scene text recognition. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 11005–11012. AAAI Press, 2020.
- [15] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Synthetic data and artificial neural networks for natural scene text recognition. In *Workshop on Deep Learning, NIPS*, 2014.
- [16] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Deep features for text spotting. In David J. Fleet, Tomás Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part IV*, volume 8692 of *Lecture Notes in Computer Science*, pages 512–528. Springer, 2014.
- [17] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep structured output learning for unconstrained text recognition. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [18] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Reading text in the wild with convolutional neural networks. *International journal of computer vision*, 116(1):1–20, 2016.
- [19] Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluís Gomez i Bigorda, Sergi Robles Mestre, Joan Mas, David Fernández Mota, Jon Almazán, and Lluís-Pere de las Heras. ICDAR 2013 robust reading competition. In *12th International Conference on Document Analysis and Recognition, ICDAR 2013, Washington, DC, USA, August 25-28, 2013*, pages 1484–1493. IEEE Computer Society, 2013.
- [20] Dimosthenis Karatzas, Lluís Gomez-Bigorda, Angelos Nicolaou, Suman K. Ghosh, Andrew D. Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, Faisal Shafait, Seiichi Uchida, and Ernest Valveny. ICDAR 2015 competition on robust reading. In *13th International Conference on Document Analysis and Recognition, ICDAR 2015, Nancy, France, August 23-26, 2015*, pages 1156–1160. IEEE Computer Society, 2015.

- [21] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017.
- [22] Hui Li, Peng Wang, Chunhua Shen, and Guyu Zhang. Show, attend and read: A simple and strong baseline for irregular text recognition. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 8610–8617. AAAI Press, 2019.
- [23] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019.
- [24] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018.
- [25] Anand Mishra, Karteek Alahari, and C. V. Jawahar. Scene text recognition using higher order language priors. In Richard Bowden, John P. Collomosse, and Krystian Mikolajczyk, editors, *British Machine Vision Conference, BMVC 2012, Surrey, UK, September 3-7, 2012*, pages 1–11. BMVA Press, 2012.
- [26] Trung Quy Phan, Palaiahnakote Shivakumara, Shangxuan Tian, and Chew Lim Tan. Recognizing text with perspective distortion in natural scenes. In *IEEE International Conference on Computer Vision, ICCV 2013, Sydney, Australia, December 1-8, 2013*, pages 569–576. IEEE Computer Society, 2013.
- [27] Zhi Qiao, Yu Zhou, Dongbao Yang, Yucan Zhou, and Weiping Wang. SEED: semantics enhanced encoder-decoder framework for scene text recognition. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 13525–13534. Computer Vision Foundation / IEEE, 2020.
- [28] Zhi Qiao, Yu Zhou, Jin Wei, Wei Wang, Yuan Zhang, Ning Jiang, Hongbin Wang, and Weiping Wang. Pimnet: A parallel, iterative and mimicking network for scene text recognition. In Heng Tao Shen, Yueting Zhuang, John R. Smith, Yang Yang, Pablo Cesar, Florian Metze, and Balakrishnan Prabhakaran, editors, *MM '21: ACM Multimedia Conference, Virtual Event, China, October 20 - 24, 2021*, pages 2046–2055. ACM, 2021.
- [29] Anhar Risnumawan, Palaiahnakote Shivakumara, Chee Seng Chan, and Chew Lim Tan. A robust arbitrary text detection system for natural scene images. *Expert Syst. Appl.*, 41(18):8027–8048, 2014.
- [30] Ahmed Sabir, Francesc Moreno-Noguer, and Lluís Padró. Visual re-ranking with natural language understanding for text spotting. In C. V. Jawahar, Hongdong Li, Greg Mori, and Konrad Schindler, editors, *Computer Vision - ACCV 2018 - 14th Asian Conference on Computer Vision, Perth, Australia, December 2-6, 2018, Revised Selected Papers, Part III*, volume 11363 of *Lecture Notes in Computer Science*, pages 68–82. Springer, 2018.

- [31] Fenfen Sheng, Zhineng Chen, and Bo Xu. NRTR: A no-recurrence sequence-to-sequence model for scene text recognition. In *2019 International Conference on Document Analysis and Recognition, ICDAR 2019, Sydney, Australia, September 20-25, 2019*, pages 781–786. IEEE, 2019.
- [32] Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(11):2298–2304, 2017.
- [33] Baoguang Shi, Mingkun Yang, Xinggang Wang, Pengyuan Lyu, Cong Yao, and Xiang Bai. ASTER: an attentional scene text recognizer with flexible rectification. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(9):2035–2048, 2019.
- [34] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. VL-ERT: pre-training of generic visual-linguistic representations. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- [35] Xin Tang, Yongquan Lai, Ying Liu, Yuanyuan Fu, and Rui Fang. Visual-semantic transformer for scene text recognition. *arXiv preprint arXiv:2112.00948*, 2021.
- [36] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 10347–10357. PMLR, 2021.
- [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017.
- [38] Kai Wang, Boris Babenko, and Serge J. Belongie. End-to-end scene text recognition. In Dimitris N. Metaxas, Long Quan, Alberto Sanfeliu, and Luc Van Gool, editors, *IEEE International Conference on Computer Vision, ICCV 2011, Barcelona, Spain, November 6-13, 2011*, pages 1457–1464. IEEE Computer Society, 2011.
- [39] Yuxin Wang, Hongtao Xie, Shancheng Fang, Jing Wang, Shenggao Zhu, and Yongdong Zhang. From two to one: A new scene text recognizer with visual language modeling network. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 14174–14183. IEEE, 2021.
- [40] Ruijie Yan, Liangrui Peng, Shanyu Xiao, and Gang Yao. Primitive representation learning for scene text recognition. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 284–293. Computer Vision Foundation / IEEE, 2021.

- [41] Mingkun Yang, Yushuo Guan, Minghui Liao, Xin He, Kaigui Bian, Song Bai, Cong Yao, and Xiang Bai. Symmetry-constrained rectification network for scene text recognition. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 9146–9155. IEEE, 2019.
- [42] Deli Yu, Xuan Li, Chengquan Zhang, Tao Liu, Junyu Han, Jingtuo Liu, and Errui Ding. Towards accurate scene text recognition with semantic reasoning networks. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 12110–12119. Computer Vision Foundation / IEEE, 2020.
- [43] Fei Yu, Jiji Tang, Weichong Yin, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. Ernie-vil: Knowledge enhanced vision-language representations through scene graphs. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pages 3208–3216. AAAI Press, 2021.
- [44] Fangneng Zhan and Shijian Lu. ESIR: end-to-end scene text recognition via iterative image rectification. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 2059–2068. Computer Vision Foundation / IEEE, 2019.