

GTC: Guided Training of CTC towards Efficient and Accurate Scene Text Recognition

Wenyang Hu,^{1*} Xiaocong Cai,² Jun Hou,² Shuai Yi,² Zhiping Lin¹

¹Nanyang Technological University

²SenseTime Group Ltd.

huwe0013@e.ntu.edu.sg, {caixiaocong, houjun, yishuai}@sensetime.com, ezplin@ntu.edu.sg

Abstract

Connectionist Temporal Classification (CTC) and attention mechanism are two main approaches used in recent scene text recognition works. Compared with attention-based methods, CTC decoder has a much shorter inference time, yet a lower accuracy. To design an efficient and effective model, we propose the guided training of CTC (GTC), where CTC model learns a better alignment and feature representations from a more powerful attentional guidance. With the benefit of guided training, CTC model achieves robust and accurate prediction for both regular and irregular scene text while maintaining a fast inference speed. Moreover, to further leverage the potential of CTC decoder, a graph convolutional network (GCN) is proposed to learn the local correlations of extracted features. Extensive experiments on standard benchmarks demonstrate that our end-to-end model achieves a new state-of-the-art for regular and irregular scene text recognition and needs 6 times shorter inference time than attention-based methods.

Introduction

Scene text recognition has been studied in academia and industry for many years, as it plays an important role in various real-world applications such as vehicle license plate recognition, identity authentication and content analysis. In recent years, many methods proposed (Bissacco et al. 2013; Shi, Bai, and Yao 2016; Shi et al. 2018) to recognize text in the wild. However, due to different sizes, fonts, colors and character placements of scene texts, scene text recognition is still a challenging task.

Current Recognition Framework Generally, the framework of scene text recognition models is an encoder-decoder structure. Recent methods mainly use two techniques to train the sequence recognition model, namely Connectionist Temporal Classification (CTC) and attention mechanism. Inspired by speech recognition, CTC is introduced to align the frame-wise probability with labels. In CTC-based methods (Liu et al. 2016; Shi, Bai, and Yao 2016), CNN is used

*This work was done when Wenyang was an intern at SenseTime Group Ltd.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

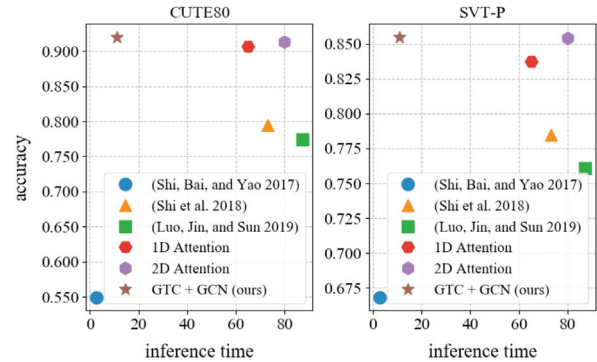


Figure 1: Accuracy versus inference time (ms/image) trade-off tables of different approaches. 1D Attention and 2D Attention represent two attention-based methods. Our method is much faster and more effective.

to extract feature sequence and the Recurrent Neural Network (RNN) is used to model the feature sequence. They are trained with CTC loss and can make fast prediction using parallel decoding. Attention-based methods use the attention-mechanism to capture the dependencies of each character in a text line, which can learn a better alignment and deeper feature representations than CTC-based methods. Also some rectification methods have been proposed for text image preprocessing. The rectify module transforms the input text image and rectifies character alignments based on Thin-Plate Spline transformation. The rectified images are then passed to the encoder-decoder structure for recognition. This module can be added to CTC or attention-based methods and it is trained in an end-to-end fashion to learn adaptive transformations.

Motivations Though previous approaches give promising results on either regular or irregular scene text recognition, they still have limitations of balancing the trade-off between recognition accuracy and inference time (see Figure 1). As attention-based methods make predictions depending on the features in previous time steps, this non-parallel decoding scheme will slow down the inference process a lot. Although

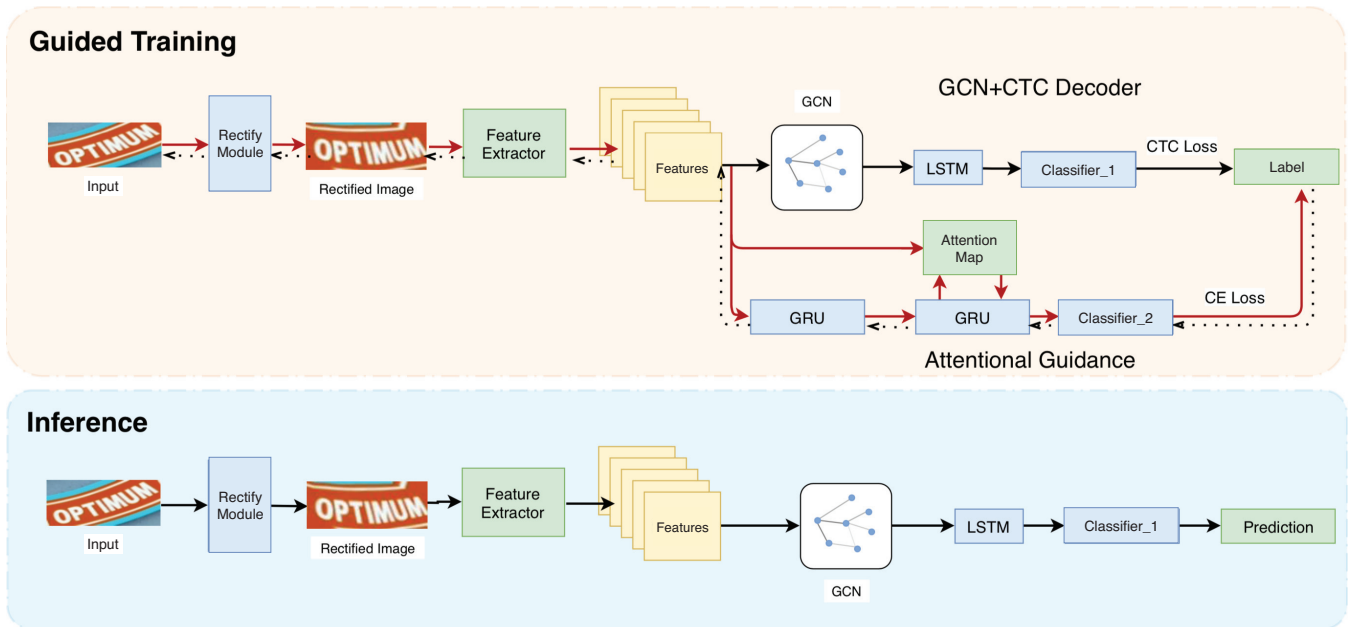


Figure 2: Overview of the proposed method. Different colors of the arrow indicate different computational graphs. Note that the gradients calculated from CTC loss are only used to update the weights of GCN+CTC decoder in the back-propagation process. The ground truth label is the same for both the guidance and GCN+CTC decoder.

CTC-based methods are relatively efficient, they are not as effective as attention-based methods, where CTC loss misleads the training of its feature alignments and feature representations. With the intention of designing an efficient and effective scene text recognizer, we aim to optimize the CTC model. To overcome the limitations of CTC, we have two motivations: (1) learning better feature representations from a more effective guidance, and (2) building correlations among the local features.

Our Work We propose the guided training of CTC (GTC), a novel method for general CTC optimization. The proposed method uses an end-to-end trainable framework as shown in Figure 2. The rectify module is a simple transformation network that applies rectification to text images. The CTC decoder is used for both training and evaluation, which leads to efficient inference. For guidance, it is made up of an effective attention decoder which will not be used in the inference. The encoder from a powerful network can learn a better alignment and feature representations, where the feature maps are easier to be decoded. Thus, the CTC model learns from the guidance and becomes more effective. The gradients calculated from cross entropy loss will be used to directly optimize the rectify module, ResNet CNN and the attention decoder, which gives a powerful guidance. The choice of guidance is flexible, which makes this method very general.

As CTC decoding allows repetitions of characters and blank labels, one label can be predicted in multiple time steps. We assume neighbouring time steps have supplementary features and there are certain correlations inside the feature sequence. For example, the letter ‘H’ can be misclassi-

fied as ‘I’ if only half part of ‘H’ is considered. To merge the supplementary features of the feature sequence, we further propose a novel GCN module to capture the dependency of each sequence slice and merge the features belonging to the same label based on the correlation. This further improves the robustness and accuracy of the CTC decoder while maintaining a fast inference speed.

In summary, the contributions of this paper are three-fold:

- 1) We design a novel GTC method for scene text recognition, which is flexible and effective. This method can be adopted to various CTC-based methods and makes the CTC model more effective by learning from powerful guidance, which may become a general method to improve the effectiveness of CTC-based methods. We also use different guidance to show that GTC is a general method for CTC model optimization.
- 2) This is the first attempt to apply graphs in scene text recognition and build sequence correlations by using GCN, which improves the accuracy and robustness of CTC decoder. Through extensive experiments, we show that the GCN module is effective.
- 3) Our proposed method has a 6 times shorter inference time than attention-based methods and achieves a new state-of-the-art on most regular and irregular scene text datasets, which is more efficient and effective compared to other works.

Related Work

Regular Scene Text Recognition Early works, e.g. (Bisacco et al. 2013), treated scene text recognition as a character segmentation and recognition problem. However, the

complicated background and different alignments of scene texts make the character segmentation the most challenging part to be trained. Besides, (Jaderberg et al. 2015) also used a multi-classification approach to directly predict each word, but this method is heavily constrained by the dictionary size. (Lee and Osindero 2016) proposed a recurrent model R^2AM with an attention decoder for regular scene text recognition. Inspired by the sequence-to-sequence alignment in speech recognition, (Shi, Bai, and Yao 2016) introduced CTC decoder into scene text recognition with a Bidirectional Long Short-Term Memory (BiLSTM) to model the feature sequence, which is known as CRNN. A Gated Recurrent Convolutional Neural Network (GRCNN) was presented by (Wang and Hu 2017) which is trained with CTC loss for regular text recognition. Alternatively, (Liu et al. 2018) proposed a binary convolutional encoder-decoder network (B-CEDNet) trained with cross-entropy loss, and achieved fast inference. However, it is designed for regular scene text recognition and requires pixel-wise labels for training. Inspired by CTC, (Bai et al. 2018) proposed a “Edit Probability” to optimize the training process, as missing or superfluity of characters may mislead CTC training. (Zhang et al. 2019) also introduced a domain adaption method to varying length text recognition. The major approach for recent regular text recognition methods is still CTC-based, which enforces the alignment between feature sequence and labels. However, CTC fails on predicting irregular scene text, as the curvatures mislead the alignments.

Irregular Scene Text Recognition Recognizing irregular scene text has attracted increasing attention in recent years, as it is a more challenging problem. Due to the distortions and curvatures of irregular texts, (Shi et al. 2018) and (Shi et al. 2016) firstly rectified the irregular texts, based on Spatial Transformer Network (STN) (Jaderberg et al. 2016), which makes text images more regular and easier to be recognized. (Luo, Jin, and Sun 2019) and (Zhan and Lu 2019) focused on improving the rectification pipeline to get better transformation results. These methods all use attention-based decoders. CTC-based decoder was also used in (Liu et al. 2016), with a STN rectifying the input image. Instead of rectifying images, some works recognize the irregular text directly. Due to the misalignment between the feature map and the attention region, (Cheng et al. 2017) used a Focusing Network to adjust the “attention drift”. (Liu, Chen, and Wong 2018) also proposed a Character-Aware Neural Network (Char-Net) to rectify the individual characters. However, both methods require character-level annotations. (Cheng et al. 2018) used a multi-direction approach to encode features before the attention network. Alternatively, there are other methods which extend the attention-mechanism into 2D feature maps, as the 2D space captures more spatial dependency. (Yang et al. 2017) and (Liao et al. 2019) recently use 2D local features in their attention networks, but these two works also require character-level annotations. (Li et al. 2019) proposed a 2D attention decoder to significantly improve the performance on irregular text recognition, although the inference time is relatively longer. (Wang et al. 2019) proposed a Transformer-based decoder which is also connected to a 2D feature map. This method achieves parallel training,

but non-parallel decoding. Among the mentioned works, the attention-based (Attention or STN+Attention) methods generally have higher accuracy on irregular scene text datasets, but the attention decoder slows down the inference. Although STN+CTC maintains a faster inference time, the capacity of CTC decoder constrains its performance.

Different from the mentioned approaches, this paper uses the guided training to optimize the CTC model with better representations of images. We also build spatial and contextual correlations among the sequence by using a designed GCN layer. To our best knowledge, this may be the first work of scene text recognition that uses a guided training method and a GCN layer.

Although (Cheng et al. 2017) used extensive experiments to show that the direct combination of CTC and attention network does not work well on scene text recognition, they did not give explanations. The reason behind is that CTC degrades the learning of feature representations. We use solid experiments to show that CTC model can achieve a much better performance by learning from an powerful guidance. Different from (Kim, Hori, and Watanabe 2017) which used a shared encoder by CTC-attention and uses attention decoder for evaluation in speech recognition, encoder and rectification model in our network are solely optimized by the gradients calculated from the guidance. At the test time, only the CTC decoder is used to make predictions and this achieves a much shorter inference time than the attention method. The guidance is only used in the training phase, which supervises the CTC model to learn a better context alignment and feature representations. Therefore, GTC makes the CTC model more efficient and effective.

Methodology

Overview

Guided training is proposed to overcome the limits of CTC itself. In the inference of CTC model, only the maximum probability of each time step is chosen as the final prediction. However, in training, different probabilities in a single time step contribute to the loss of different CTC paths. The CTC training will make feature representations to tolerate some prediction error. If we have a ground truth label ‘AB’, for a 3 time-steps output, the CTC path (pseudo label) can be ‘A-B’ or ‘-AB’ or ‘AB-’ or ‘AAB’ or ‘ABB’. As labels for CTC loss calculation are ambiguous, it is confusing to learning feature representations in each time step. Missing or superfluous characters may degrade the learning of its feature alignments and feature representations. Though we used the same encoders for both attention model and CTC model, we found in experiments that the encoder of CTC model has poor feature representations. We assert that the performance of CTC encoder is actually constrained by CTC loss itself. Therefore, a guidance can provide better feature representations for CTC model.

The proposed GTC is described in this section, where a general attention decoder is used as a guidance and our GCN+CTC decoder is used for training and inference. As shown in Figure 2, our network consists of four parts. The first part is a STN which is the same as in (Shi et al. 2016).

It transforms input images into normalized images. The second part is a ResNet backbone for feature extraction, which is widely used in scene text recognition (Li et al. 2019; Wang et al. 2019; Zhan and Lu 2019). The third part is an attentional guidance which uses attention mechanism to output the text sequence. The fourth part is a GCN powered CTC decoder which strengthens the correlations of feature sequence. The STN, ResNet-CNN and the attentional guidance are solely trained with cross entropy loss, while the GCN+CTC decoder is trained with CTC loss.

Spatial Transformer Network

As many text images in natural scenes appear with curved texts and different perspectives, the transformation module is adopted for robust and accurate recognition, which applies spatial transformation to text images and normalizes the character region. It is a differentiable module and consists of a localization network and a grid generator. The localization network will predict transformation parameters and use them to create a grid. The grid and the input image will be sampled by the generator to generate the transformed output.

Feature Extractor

We choose ResNet50 (He et al. 2016) as our network’s backbone, which is shown in Table 1. To extract more precise features, we change the original residual block stride from 2 to 1. We also add two max-pooling layers for down-sampling the feature map. The extracted feature sequence $h^{1:T}$ has a fixed height and a varying length, which will be used for decoding.

Attentional Guidance

Inspired by machine translation, the sequence-to-sequence model is used to translate a feature sequence into a character sequence, which aligns outputs and labels. The attention mechanism in such a model has the ability to capture output dependencies and focus on character region at each time step. For fair comparisons, we choose a general attention decoder as in (Shi et al. 2016; Cheng et al. 2017; Zhan and Lu 2019) to demonstrate the effectiveness of GTC. We adopt the attentional sequence-to-sequence decoder at the top of the ResNet backbone. It is based on an RNN producing a target sequence of length T , denoted by (y_1, y_2, \dots, y_T) .

The attention decoder either predicts a character or an end-of-sequence token ‘EOS’. It stops predicting when it predicts an ‘EOS’. The Gated Recurrent Cell (GRU) is adopted to learn the attention dependency. At time-step t , output y_t is,

$$y_t = \text{Softmax}(W^T s_t), \quad (1)$$

where s_t is a hidden state of the GRU cell and W is a trainable parameter.

The hidden state s_t is updated via the recurrent process of GRU:

$$s_t = \text{GRU}(y_{prev}, g_t, s_{t-1}), \quad (2)$$

where y_{prev} is the embedding vector of the previous output y_{t-1} . During training, y_{t-1} is replaced by the ground truth

Table 1: The configuration of our ResNet50 feature extractor. ‘‘Conv’’ means convolutional layers, provided with its kernel size, output channels, stride and padding. The stride for all ‘Residual Block’ is set to 1. The configuration for ‘‘Max-pooling’’ and ‘‘Average-pooling’’ represents its kernel size, stride, and padding. The overall down-sampling ratio is $W: 1/4, H: 1/16$.

Layer Name	Configuration
Conv	$7 \times 7, 64, 2 \times 2, 3$
Max-pooling	$3 \times 3, 2 \times 2, 1$
Residual Block	$\begin{bmatrix} \text{Conv} : 1 \times 1, 64 \\ \text{Conv} : 3 \times 3, 64 \\ \text{Conv} : 1 \times 1, 256 \end{bmatrix} \times 3$
Residual Block	$\begin{bmatrix} \text{Conv} : 1 \times 1, 128 \\ \text{Conv} : 3 \times 3, 128 \\ \text{Conv} : 1 \times 1, 512 \end{bmatrix} \times 4$
Max-pooling	$2 \times 1, 2 \times 1$
Residual Block	$\begin{bmatrix} \text{Conv} : 1 \times 1, 256 \\ \text{Conv} : 3 \times 3, 256 \\ \text{Conv} : 1 \times 1, 1024 \end{bmatrix} \times 3$
Max-pooling	$2 \times 1, 2 \times 1$
Residual Block	$\begin{bmatrix} \text{Conv} : 1 \times 1, 512 \\ \text{Conv} : 3 \times 3, 512 \\ \text{Conv} : 1 \times 1, 2048 \end{bmatrix} \times 3$
Average-pooling	$4 \times 1, 1 \times 1$

sequence. g_t represents the glimpse vector calculated as:

$$g_t = \sum_{i=1}^T (\alpha_{ti} h_i), \quad (3)$$

where h_i is the feature sequence vector of $h^{1:T}$ at the time-step i . α_t is the attention weight vector as follows:

$$\alpha_t = \text{Attention}(s_{t-1}, h_i), \quad (4)$$

which is described in (Luong, Pham, and Manning 2015).

The attentional guidance is trained with the cross entropy loss and the prediction results will not be used in the evaluation.

GCN+CTC Decoder

Given a sequence of probability distributions $y^{1:T}$ of length T , it seeks multiple paths $M^{-1}(l)$ that produce the same label sequence l , allowing repetitions of consecutive characters or blank labels \emptyset . y^t represents the probability distribution at time-step t over the classification labels L , where $\emptyset \in L$. M defines the operation of mapping all possible paths π to the target label. For example, it maps the path ‘-hh-e-ll-l-oo-’ into ‘hello’. CTC trains the network to optimize the summation of probabilities over all paths:

$$p(l|h^{1:T}) = \sum_{\pi \in M^{-1}(l)} p(\pi|h^{1:T}), \quad (5)$$

where the probability of one possible path π is calculated as:

$$p(\pi|h^{1:T}) = \prod_{t=1}^T y_{\pi_t}^t, \forall \pi \in M^{-1}(l). \quad (6)$$

Based on Equations (5) and (6), CTC trains the network to optimize the loss function:

$$Loss_{CTC} = -\log p(l|h^{1:T}). \quad (7)$$

In CRNN, BiLSTM is used to extract sequence feature by reading the text line from both directions. However, it lack the ability of focusing on local regions, as characters appear in separate locations. Graph Convolutional Networks (GCNs) (Kipf and Welling 2016) are an efficient variant of CNNs on graph data, where edges of the graph represent implicit connections within the data. Given a relation defined by a graph, graph convolutions pass the messages from a node to its neighbors. We propose a special GCN layer before the BiLSTM, where a similarity adjacency matrix and a distance matrix are combined to describe the spatial contextual correlations.

Given the feature map $h_{1:T}$ from ResNet CNN, the adjacency matrix is learned by computing pairwise similarity between every two sequence slices:

$$A_S(i, j) = f(c_i, c_j), \quad (8)$$

The similarity projection function is defined as:

$$f(c_i, c_j) = \frac{c_i \cdot c_j}{\|c_i\| \|c_j\|}, \quad (9)$$

where c_i is a linear transformation result of h_i . The formula basically calculates pairwise cosine similarities. In addition to using similarity relations to focus on similar features, a distance matrix is also used to constrain the similarity to focus on neighboring features. The distance matrix is defined as:

$$A_D(i, j) = \frac{\exp(-d_{ij} + \beta)}{\exp(-d_{ij} + \beta) + 1}, \quad (10)$$

where $d_{ij} = |i - j|$ and β is a scale factor. Therefore, the final output of our GCN+CTC is calculated as:

$$X = (A_S * A_D)HW_g, \quad (11)$$

where W_g is an optional weight matrix. The X is then passed to the BiLSTM for sequence modelling.

$$logits = Seq(X)W_c, \quad (12)$$

where W_c is a weight matrix for classification and Seq is the BiLSTM with the hidden size of 512. The logits and label l are finally used to calculate CTC loss to train the GCN+CTC decoder.

In summary, GTC uses a more powerful model to guide CTC decoder, where the gradients calculated from CTC loss will not be used to update the rectify module, ResNet CNN, feature maps or the attentional guidance. CTC Decoder updates itself through the training process of the attentional guidance, where it learns to predict from better feature representations and better alignments. GCN builds correlations among features and further improves the performance.

Experiments

We conduct experiments on both regular and irregular scene text datasets to evaluate the performance of our proposed method.

Datasets

Synthetic Datasets There are three public synthetic datasets, namely Synth90K (Jaderberg et al. 2015), SynthText (Gupta, Vedaldi, and Zisserman 2016) and SynthAdd (Li et al. 2019). Synth90K is randomly generated based on the 90K common English words, which contains 9-million image instances. In SynthText, texts are randomly blended on full images and text samples are cropped. There are in total 8-million images in SynthText. SynthAdd is also a synthetic dataset with only text line annotations to compensate the lack of special characters. There are 1.6-million images in SynthAdd.

Regular datasets mainly contains text images with horizontal layout of characters and equal spacing between characters. These images can be simply recognized by reading from left to right.

- **IIIT5K-Words (IIIT5K)** (Mishra, Alahari, and Jawahar 2012) is collected from Google image searches, which contains 2,000 images for training and 3,000 images for evaluation.
- **Street View Text (SVT)** (Wang, Babenko, and Belongie 2011) is collected from Google Street Image, with 257 training images and 647 testing images. Some of these outdoor street images are of low-resolution.
- **ICDAR 2003 (IC03)** (Lucas et al. 2003) is a regular text dataset cropped from real scene text images. It contains 1,156 training images and 1,110 testing images. Instead of filtering the samples which contain non-alphanumeric characters or have fewer than three characters, we use the whole dataset for testing.
- **ICDAR 2013 (IC13)** (Karatzas et al. 2013) has 848 cropped text instances for training and 1095 for testing. It inherits most of IC03’s images.

Irregular datasets contain many hard cases of scene text images. Many of these are curved, rotated and distorted text images.

- **ICDAR 2015 (IC15)** (Karatzas et al. 2015) contains 4,468 images for training and 2,077 images for evaluation. These images are cropped from natural images captured by Google Glasses. Thus, many images are blurry, curved and rotated.
- **SVT Perspective (SVT-P)** (Quy Phan et al. 2013) consists of 645 cropped images, which are from side view images and contain perspective distortions.
- **CUTE80** (Risnumawan et al. 2014) contains 288 text patches cropped from natural scene images. Many of these are curved text images of high resolution.
- **COCO-Text (COCO)** (Veit et al. 2016) contains 42618 real text images for training and 9837 images for testing.

Implementation Details

We implement our proposed network structure with PyTorch and conduct all experiments on NVIDIA Tesla V100 GPUs

Table 2: Text line recognition accuracy (in percentages) on public benchmarks, including both regular and irregular datasets. All experiments are compared in a lexicon-free basis. In each column, the state-of-the-art result is shown in **bold**, and the second best result is shown in underline. The methods marked with “*” are carefully evaluated with rotation strategy. The methods marked with “+” are trained with both word-level and character-level annotations. Our best model outperforms all the compared methods in the overall recognition rate, and achieves the state-of-the-art on most scene text datasets.

Method		Regular Text				Irregular Text			Infer-Time ms/image
		IIT5K	IC03	IC13	SVT	IC15	SVT-P	CUTE80	
CTC	(Liu et al. 2016)	83.3	89.9	89.1	83.6	-	73.5	-	-
	(Wang and Hu 2017)	80.8	91.2	-	81.5	-	-	-	-
	(Shi, Bai, and Yao 2016)	81.2	89.9	89.6	82.7	-	66.8	54.9	2.7
Attention	(Lee and Osindero 2016)	78.4	88.7	90.0	80.7	-	-	-	-
	(Shi et al. 2016)	81.9	90.1	88.6	81.9	-	71.8	59.2	-
	(Yang et al. 2017)+	-	-	-	-	-	75.8	69.3	-
	(Cheng et al. 2017)+	87.4	94.2	93.3	85.9	70.6	-	-	-
	(Liu et al. 2018)+	87.0	93.1	92.9	-	-	-	-	-
	(Liu, Chen, and Wong 2018)+	92.0	92.0	91.1	85.5	74.2	78.9	-	-
	(Bai et al. 2018)+	88.3	94.6	94.4	87.5	73.9	-	-	-
	(Zhan and Lu 2019)	93.3	-	91.3	90.2	76.9	79.6	83.3	-
	(Shi et al. 2018)	93.4	94.5	91.8	93.6	76.1	78.5	79.5	73.1
	(Luo, Jin, and Sun 2019)	91.2	95.0	92.4	88.3	68.8	76.1	77.4	87.3
	(Liao et al. 2019)+	91.9	-	91.5	86.4	-	-	79.9	-
	(Li et al. 2019)*	95.0	-	<u>94.0</u>	91.2	78.8	86.4	89.6	-
	(Wang et al. 2019)	93.3	-	91.3	88.1	77.9	80.2	85.1	-
Ours	CTC Baseline	95.4	93.6	91.8	89.2	76.4	80.1	85.7	9.5
	GTC (1D)	96.0	95.8	93.2	91.8	79.5	85.6	91.3	10.6
	GTC (2D)	95.0	94.6	92.6	91.2	79.3	83.4	90.6	10.6
	CTC + GCN	95.2	93.9	92.4	90.6	76.6	81.7	88.2	11.0
	GTC (1D) + GCN	95.5	95.2	94.3	<u>93.2</u>	80.4	85.5	92.0	11.0
	GTC (2D) + GCN	<u>95.8</u>	<u>95.5</u>	94.4	92.9	79.5	85.7	<u>92.2</u>	11.0
	GTC (2D) + GCN *	95.5	95.2	94.3	92.9	82.5	<u>86.2</u>	92.3	-

with 16GB memory. We use a batch size of 32 on each GPU, with 32 GPUs in total. ADAM optimizer is chosen for training, with the initial learning rate set to 10^{-3} and a decay rate of 0.1 every 30000 iterations. We directly train our network using synthetic data (Synth90k, SynthText and SynthAdd) and the training images provided from public benchmarks (IIT5K, SVT, IC03, IC13, IC15, COCO), which is the same as what is described in (Li et al. 2019). We randomly sample 5.6-million images from those training images for training. The input images are resized to have a fixed height of 64 pixels and a varying length, but not longer than 160 pixels.

During the evaluation, the attentional guidance is abandoned. We directly evaluate the input images by using CTC decoder without any rotation or prediction strategies. Greedy decoding is adopted.

Experimental Results

All experiments were evaluated in a lexicon-free condition. We directly evaluated the test images of 7 public datasets and the results are shown in Table 2. ‘1D’ denotes the attentional guidance described in the Methodology. We also evaluated GTC by using another attentional guidance from (Li et al. 2019), which shows the robustness of GTC. The related experimental results are denoted by ‘2D’. (Li et al. 2019) rotated testing images by 90 degrees clockwise and anti-clockwise respectively, and recognized them together with

the original image. For fair comparison, our experiment in the last row of Table 2 also used the rotation strategy. Our other experiments directly evaluated the test images without rotation. Therefore, we claim that we achieved the state-of-

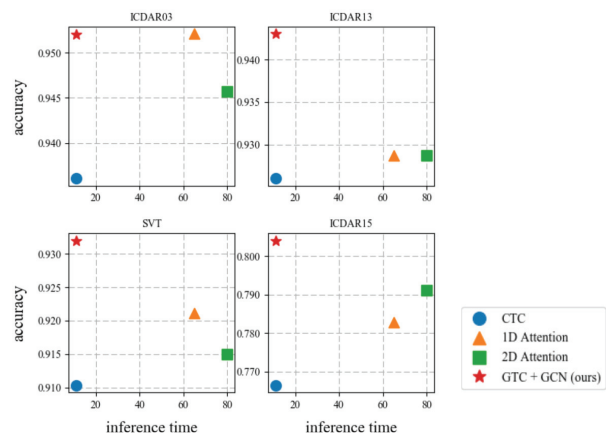


Figure 3: Accuracy versus inference time (ms/image) trade-off tables of different approaches. GTC outperforms CTC and attention-based methods in all datasets.

Table 3: Comparison between GTC and guided training of attention. The performance was measured by using overall accuracy on all public benchmarks.

Experiment	overall accuracy
Use CTC to Guide Attention	87.87%
Use Attention to Guide CTC (GTC)	90.06%

the-art by using only word annotations and public datasets.

Our GTC method outperforms CTC-based methods a lot while it maintains a fast inference. The GCN module has also been verified to model a better feature sequence for irregular texts. Note that all GTC results are based on the predictions of CTC decoder.

To evaluate the efficiency of our approach, we also conduct experiments to analyze the inference time of different methods. We fix the batch size as 1 and run all test experiments on the same device. The inference time is measured on a single NVIDIA Titan X GPU with 12 GB memory. The comparison of our method and other methods is shown in Figure 3. The results show that our GTC method achieves 6 times faster in the inference than attention-based methods and maintains the highest recognition rate. The GCN+CTC decoder is even 14+ times faster than the attention decoder. The inference time is 54 ms/image for attention decoder and 3.7 ms/image for the GCN+CTC decoder.

Ablation Study

In this experiment, CTC is used to guide the attention decoder instead. The encoder is trained with the CTC loss and the attention decoder is used for evaluation. The result in Table 3 shows that CTC is not an effective guidance compared with Attention. The result also indicates that the CTC loss harms the training process and produces poor feature representations. Therefore, this experiment suggests that a powerful guidance is necessary in the guided training in order to improve the overall performance.

Besides, we find that GTC also has better transformation results compared with STN+CTC framework (see Figure 4). As attention mechanism is more sensitive to spatial information, the transformation module can be more responsive.

To clearly describe how GCN works, we present visualizations in Figure 5. With distance matrix and similarity matrix multiplied point-wisely, we get the adjacency matrix of GCN, which focuses on local similar features. Thus, it establishes the local correlations for better sequence modelling.




Figure 4: Examples of original images and rectified images from different approaches.

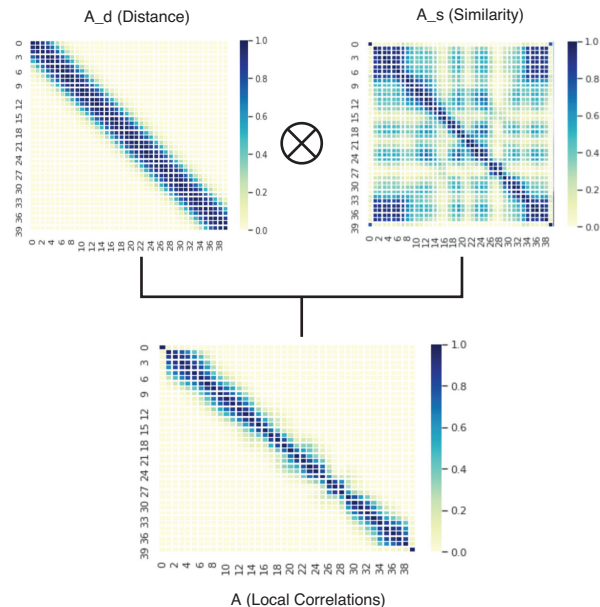


Figure 5: Visualization of the similarity matrix, the distance matrix and their point-wise product.

Conclusion

In this paper, we propose an effective and efficient GTC method to significantly improve the robustness and performance on scene text recognition, which overcomes the limitations of CTC. This work is the first attempt to use GCN to learn the local correlations of feature sequences, which further improves CTC performance. We conduct experiments to evaluate the effectiveness of our method on 7 public benchmarks and the method achieves a new state-of-the-art on most datasets.

References

Bai, F.; Cheng, Z.; Niu, Y.; Pu, S.; and Zhou, S. 2018. Edit probability for scene text recognition. In *Proceedings of the IEEE Conf. Comp. Vis. Patt. Recogn.*, 1508–1516.

Bissacco, A.; Cummins, M.; Netzer, Y.; and Neven, H. 2013. Photoocr: Reading text in uncontrolled conditions. In *Proceedings of the IEEE International Conf. Comp. Vis.*, 785–792.

Cheng, Z.; Bai, F.; Xu, Y.; Zheng, G.; Pu, S.; and Zhou, S. 2017. Focusing attention: Towards accurate text recognition in natural images. In *Proceedings of the IEEE International Conf. Comp. Vis.*, 5076–5084.

Cheng, Z.; Xu, Y.; Bai, F.; Niu, Y.; Pu, S.; and Zhou, S. 2018. Aon: Towards arbitrarily-oriented text recognition. In *Proceedings of the IEEE Conf. Comp. Vis. Patt. Recogn.*, 5571–5579.

Gupta, A.; Vedaldi, A.; and Zisserman, A. 2016. Synthetic data for text localisation in natural images. In *Proceedings of the IEEE Conf. Comp. Vis. Patt. Recogn.*, 2315–2324.

- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conf. Comp. Vis. Patt. Recogn*, 770–778.
- Jaderberg, M.; Simonyan, K.; Zisserman, A.; et al. 2015. Spatial transformer networks. In *Advances in neural information processing systems*, 2017–2025.
- Jaderberg, M.; Simonyan, K.; Vedaldi, A.; and Zisserman, A. 2016. Reading text in the wild with convolutional neural networks. *International Journal of Computer Vision* 116(1):1–20.
- Karatzas, D.; Shafait, F.; Uchida, S.; Iwamura, M.; Bigorda, L. G.; Mestre, S. R.; Mas, J.; Mota, D. F.; Almazan, J. A.; and De Las Heras, L. P. 2013. Icdar 2013 robust reading competition. In *2013 12th International Conference on Document Analysis and Recognition*, 1484–1493. IEEE.
- Karatzas, D.; Gomez-Bigorda, L.; Nicolaou, A.; Ghosh, S.; Bagdanov, A.; Iwamura, M.; Matas, J.; Neumann, L.; Chandrasekhar, V. R.; Lu, S.; et al. 2015. Icdar 2015 competition on robust reading. In *2015 13th International Conference on Document Analysis and Recognition*, 1156–1160. IEEE.
- Kim, S.; Hori, T.; and Watanabe, S. 2017. Joint ctc-attention based end-to-end speech recognition using multi-task learning. In *2017 IEEE ICASSP*, 4835–4839. IEEE.
- Kipf, T. N., and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Lee, C.-Y., and Osindero, S. 2016. Recursive recurrent nets with attention modeling for ocr in the wild. In *Proceedings of the IEEE Conf. Comp. Vis. Patt. Recogn*, 2231–2239.
- Li, H.; Wang, P.; Shen, C.; and Zhang, G. 2019. Show, attend and read: A simple and strong baseline for irregular text recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 8610–8617.
- Liao, M.; Zhang, J.; Wan, Z.; Xie, F.; Liang, J.; Lyu, P.; Yao, C.; and Bai, X. 2019. Scene text recognition from two-dimensional perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 8714–8721.
- Liu, W.; Chen, C.; Wong, K.-Y. K.; Su, Z.; and Han, J. 2016. Star-net: a spatial attention residue network for scene text recognition. In *BMVC*, volume 2, 7.
- Liu, Z.; Li, Y.; Ren, F.; Goh, W. L.; and Yu, H. 2018. Squeezedtext: A real-time scene text recognition by binary convolutional encoder-decoder network. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Liu, W.; Chen, C.; and Wong, K.-Y. K. 2018. Char-net: A character-aware neural network for distorted scene text recognition. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Lucas, S. M.; Panaretos, A.; Sosa, L.; Tang, A.; Wong, S.; and Young, R. 2003. Icdar 2003 robust reading competitions. In *Seventh International Conference on Document Analysis and Recognition, 2003. Proceedings.*, 682–687. Citeseer.
- Luo, C.; Jin, L.; and Sun, Z. 2019. Moran: A multi-object rectified attention network for scene text recognition. *Pattern Recognition* 90:109–118.
- Luong, M.-T.; Pham, H.; and Manning, C. D. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*.
- Mishra, A.; Alahari, K.; and Jawahar, C. 2012. Scene text recognition using higher order language priors. In *Proc. British Mach. Vis. Conf.*, 1–11.
- Quy Phan, T.; Shivakumara, P.; Tian, S.; and Lim Tan, C. 2013. Recognizing text with perspective distortion in natural scenes. In *Proceedings of the IEEE International Conf. Comp. Vis.*, 569–576.
- Risnumawan, A.; Shivakumara, P.; Chan, C. S.; and Tan, C. L. 2014. A robust arbitrary text detection system for natural scene images. *Expert Systems with Applications* 41(18):8027–8048.
- Shi, B.; Bai, X.; and Yao, C. 2016. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *IEEE transactions on pattern analysis and machine intelligence* 39(11):2298–2304.
- Shi, B.; Wang, X.; Lyu, P.; Yao, C.; and Bai, X. 2016. Robust scene text recognition with automatic rectification. In *Proceedings of the IEEE Conf. Comp. Vis. Patt. Recogn*, 4168–4176.
- Shi, B.; Yang, M.; Wang, X.; Lyu, P.; Yao, C.; and Bai, X. 2018. Aster: An attentional scene text recognizer with flexible rectification. *IEEE transactions on pattern analysis and machine intelligence*.
- Veit, A.; Matera, T.; Neumann, L.; Matas, J.; and Belongie, S. 2016. Coco-text: Dataset and benchmark for text detection and recognition in natural images. *arXiv preprint arXiv:1601.07140*.
- Wang, J., and Hu, X. 2017. Gated recurrent convolution neural network for ocr. In *Advances in Neural Information Processing Systems*, 335–344.
- Wang, K.; Babenko, B.; and Belongie, S. 2011. End-to-end scene text recognition. In *2011 International Conf. Comp. Vis.*, 1457–1464. IEEE.
- Wang, P.; Yang, L.; Li, H.; Deng, Y.; Shen, C.; and Zhang, Y. 2019. A simple and robust convolutional-attention network for irregular text recognition. *arXiv preprint arXiv:1904.01375*.
- Yang, X.; He, D.; Zhou, Z.; Kifer, D.; and Giles, C. L. 2017. Learning to read irregular text with attention mechanisms. In *IJCAI*, volume 1, 3.
- Zhan, F., and Lu, S. 2019. Esir: End-to-end scene text recognition via iterative image rectification. In *Proceedings of the IEEE Conf. Comp. Vis. Patt. Recogn*, 2059–2068.
- Zhang, Y.; Nie, S.; Liu, W.; Xu, X.; Zhang, D.; and Shen, H. T. 2019. Sequence-to-sequence domain adaptation network for robust text image recognition. In *Proceedings of the IEEE Conf. Comp. Vis. Patt. Recogn*, 2740–2749.