

Simultaneous Recognition of Horizontal and Vertical Text in Natural Images

Chankyu Choi, Youngmin Yoon, Junsu Lee, Junseok Kim

NAVER Corporation

{chankyu.choi,youngmin.yoon,junsu.lee,jun.seok}@navercorp.com

Abstract. Recent state-of-the-art scene text recognition methods have primarily focused on horizontal text in images. However, in several Asian countries, including China, large amounts of text in signs, books, and TV commercials are vertically directed. Because the horizontal and vertical texts exhibit different characteristics, developing an algorithm that can simultaneously recognize both types of text in real environments is necessary. To address this problem, we adopted the direction encoding mask (DEM) and selective attention network (SAN) methods based on supervised learning. DEM contains directional information to compensate in cases that lack text direction; therefore, our network is trained using this information to handle the vertical text. The SAN method is designed to work individually for both types of text. To train the network to recognize both types of text and to evaluate the effectiveness of the designed model, we prepared a new synthetic vertical text dataset and collected an actual vertical text dataset (VTD142) from the Web. Using these datasets, we proved that our proposed model can accurately recognize both vertical and horizontal text and can achieve state-of-the-art results in experiments using benchmark datasets, including the street view test (SVT), IIIT-5k, and ICDAR. Although our model is relatively simple as compared to its predecessors, it maintains the accuracy and is trained in an end-to-end manner.

Keywords— Directional Encoding Mask; Selective Attention Network; Supervised Learning; Horizontal and Vertical Text Recognition

1 Introduction

Optical character recognition (OCR) in the wild is a challenging problem in computer-vision research, and several approaches have been studied to perform OCR on images acquired from various environments such as documents or natural scenes. OCR is usually segmented into separate modules such as scene text detection and scene text recognition. The scene text detection module identifies the regions of interest in images that contain text; various algorithms handle this task. Subsequently, the scene text recognition module translates these regions into labeled text. This study presents a novel method to perform scene text recognition.

Majority of the prior scene text recognition algorithms only attempted to cover the horizontal text. However, in several Asian countries, including China, large amounts



Fig. 1: Majority of the prior scene text recognition algorithms only attempted to cover the horizontal text. However, in several Asian countries, including China, large amounts of text in signs, books, and TV commercials are vertically directed.

of text in signs, books, and TV commercials are vertically oriented. Horizontal and vertical texts exhibit different characteristics that should be simultaneously solved for. They occupied different aspect ratio text region in image. Therefore, a clever algorithm must be developed to deal with both types of text. Fig.1 depicts the reason due to which the OCR algorithm is able to handle the vertically directed text in real environments. A recent study used arbitrarily oriented text recognition (AON)[1] encoded image input in four directions (top-to-bottom, bottom-to-top, left-to-right, and right-to-left) to detect a freely rotated word. AON can, therefore, adapt in an unsupervised manner to tackle the rotated text using four orientations of directional information.

In this study, inspired by the AON method, we simultaneously recognize the horizontal and vertical texts. Our method uses carefully designed information formulated via the directional encoding mask (DEM) and selective attention network (SAN) methods. Both the modules help the network to learn quickly and accurately. DEM is designed to extract the direction of a text region before it is passed to the text recognition model. Therefore, it should be created in advance during the preparation time of the ground truth data. It encapsulates the directional information for the text region. An advantage of DEM is that it helps the network to quickly and accurately learn to distinguish between vertical and horizontal text in a supervised manner. Our contributions are as follows.

- 1) We created two alternative methods, namely, DEM and SAN that can simultaneously deal with both horizontal and vertical characters.
- 2) To train for vertical text, we generated a synthetic vertical text dataset (SVTD). In addition, to evaluate the performance with respect to the vertical text, we collected an actual vertical text dataset (VTD142), including building signs and titles. All the real data were collected from the webpages.
- 3) Our model achieves state-of-the-art results with respect to various benchmark datasets.

2 Related Work

Previously, the development of OCR algorithms mostly used handcrafted features, which are difficult to design and exhibit limited performance in specific areas; however,

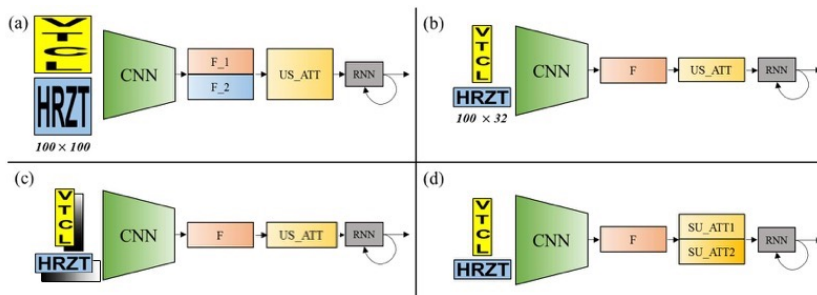


Fig. 2: Various models for recognizing the horizontal and vertical text : (a) AON style[1] with 2-directional clue network), (b) Baseline, (c) DEM, (d) SAN. F means features, SU_ATT means supervised attention and US_ATT means unsupervised attention mechanism respectively.

currently, neural networks offer a promising alternative for powerful OCR algorithms[2–5]. Convolutional recurrent neural network (CRNN)[6] encode the input images using convolutional neural network (CNN) and further decode the encoded features using recurrent neural network (RNN) by connectionist temporal classification (CTC). Recurrent and recursive convolution network[7, 8] improve the representational power and efficiency of the encoders. Attention OCR[9] improves the overall performance using positional encoding and attention mechanisms. Gated recurrent convolutional neural network (GRCN)[10] adopt gated recurrent CNNs to further improve the performance. Currently, even though there are several deep neural network (DNN)[1, 11, 12] in use, only a few studies that focus on the vertical and horizontal texts have been conducted. FOTS[13] handles oriented text, which is essentially different from recognizing characters the in vertical direction. The results also exhibited lower performance as compared to the attention-based decoder because they have experimented with a CTC-based decoder. AON[1] is an unsupervised learning method that learns a characters direction by network oneself. Conversely, we examine the orientation of a word during the training stage and add this feature to the network so that the network can be learned in a supervised manner. We extended the AON to a supervised learning method to recognize the horizontal and vertical texts.

3 Proposed Method

We begin by explaining the manner in which a unified network can be designed to simultaneously handle the horizontal and vertical texts using an images aspect ratio. We examine the direction of a word during the training stage and add this feature to the network so that the network can be learned in a supervised manner. To explain it more clearly, our methods determine the direction of the text after checking the width-to-height ratio of the input image. Further, the vertical text is rotated by 90 and input to the network. In this study, we introduce the two methods that are required to achieve this.

4 C. Choi et al.

3.1 Directional Encoding Mask

DEM combines the input image and the directional information of the characters, which are further encoded by the CNN encoder. Because DEM contains the directional information of the text, the network uses this information while training. In this study, we encode the directional information in cosine for the horizontal text and in sine for the vertical text. The equation below describes the DEM.

$$\text{DEM} = \begin{cases} \{y \mid y = \sin(0.5 * \text{normalized_width} * \pi)\}, & \text{if width} > \text{height} \\ \{y \mid y = \cos(0.5 * \text{normalized_width} * \pi)\}, & \text{otherwise} \end{cases} \quad (1)$$

In the equation(1), the normalized width is the x coordinate in normalized space [0,1]. As can be observed from the equation, we encode DEM for the horizontal text using a cosine function. The DEM for vertical text is encoded using a sine function. The kernel function of DEM can be replaced by similar functions; therefore, we chose it experimentally.

$$X = \text{concat}(I, \text{DEM}) \quad (2)$$

DEM is not a single value, but a two-dimensional image of the same size as the input image. DEM and the input image are concatenated and entered into the CNN. X is used as the input of the CNN.

3.2 Selective Attention Network

SAN is another supervised learning method. It has two attention masks, each of which reacts according to the corresponding text direction; therefore, they work separately. SAN can turn its own attention weight on and off depending on the direction of the text. This differs from AONs clue network because it is based on supervised learning. The *character feature vector* for the RNN is defined as :

$$x_t^c = W_c c_{t-1}^{\text{OneHot}} \quad (3)$$

where c_{t-1} is one hot encoding of the previous letter (ground truth during training and predicted during test time).

Let s_t be the hidden state of the LSTM at time t . Further, we compute the output and the subsequent state of the LSTM, as shown

$$(y_t, s_t) = \text{LSTM}_{\text{step}}(x_t, s_{t-1}) \quad (4)$$

Let us denote the cnn features as $f = f_{i,j,c}$ where i, j are the index locations in the feature map and c indexes the channels. We verified the aspect ratios of the image inputs and created two attention modules. These attention masks are denoted as $\alpha_t^{(h,v)}$, where h and v indicate the horizontal and vertical texts, respectively. We name this as *selective attention network* because it responds to the text direction.

$$\alpha_t^{(h,v)} = \text{softmax}(V_a^T \circ \tanh(W_s^{(h,v)} s_t + W_f f_{i,j})) \quad (5)$$

SAN comprises only one set of attention weights for each horizontal and vertical text. In the equation(5), \circ is the Hadamard product and \tanh is applied element-wise to its vector argument. Further, the *visual feature vector* (also known as the *context vector*), x_t^v is computed as the weighted image features based on the attention at time step t as shown :

$$x_t^v = \sum_{i,j} \alpha_{t,i,j}^{(h,v)} f_{i,j,c} \quad (6)$$



Fig. 3: SVTD is a new synthetic vertical text dataset in a similar fashion as [1]. VTD142 is real vertical text dataset, which was collected from web.

We can define x_t as the input of LSTM, as depicted bellow.

$$x_t = x_t^c + x_t^v \quad (7)$$

The final predicted distribution over the letters at time t is described as shown.

$$\tilde{y}_t = W_{\tilde{o}} y_t + b_{\tilde{o}} \quad (8)$$

$$\tilde{Y} = \{\tilde{y}_0, \dots, \tilde{y}_t, \dots, \tilde{y}_{T-1}\} \quad (9)$$

In the equations, T is the sequence length of the RNN and \tilde{Y} is the result of our LSTM decoder.

4 Experiments

4.1 Dataset

A range of datasets of images of words are used to train and test the model.

SVT[14] includes cropped images, including 257 training images and 647 test images, captured from Google Street View

IIT5k[15] contains 2000 cropped training images and 3000 testing images downloaded from the Google image search engine

ICDAR 2003[16] (IC03 in short) contains 258 full-scene and 936 cropped images for training, and 251 full-scene and 860 cropped images for testing.

ICDAR 2013[17] (IC13 in short) contains 1015 cropped text images, but no lexicon is associated.

ICDAR 2015[18] (IC15 in short) contains 2077 cropped text images without lexicon.

SVTD is a new synthetic vertical text dataset in a similar fashion as [1]. It includes 1 million training word images and 3,000 testing word images.

VTD142 is real vertical text dataset including 142 word images. We collect it from web.

Large amounts of datasets are needed to create a good-performance, deep-learning model. However, producing annotated training data requires considerable time and effort. Compared to other tasks in computer vision, generating high-quality synthetic data using various font files in optical character recognition (OCR) is easy. In previous studies, millions of synthetic images were created for recognition[19] and they have been

6 C. Choi et al.

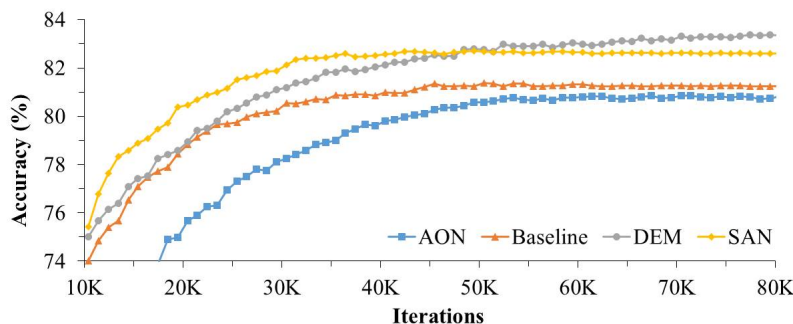


Fig. 4: Comparison of the performance of the baseline and the proposed method. The x-axis represents the number of iterations and the y-axis represents the average accuracy of the test DB.

further developed for use in detectors[20]. However, those studies only dealt with horizontal characters, not vertical ones, which could not be directly used. We applied the method described in [20] to find the area in which the letters were rendered. The depth and segmentation maps were estimated from the background image to be synthesized, and the vertical characters were rendered in the homogeneous regions. We also applied augmentation techniques, such as affine transformations, contrast changes, and certain noises, to generate various images. To our knowledge, there are no publicly available vertical-text datasets that can be used for evaluation; hence, we had to collect new data to properly evaluate our proposed method. To this end, we chose several appropriate keywords (e.g., vertical sign, buildings sign, and vertical text) and used them in the web. We usually gathered vertical signs on the buildings because that is the way they are mostly seen in real environments, such as streets, shops, and buildings.

We follow the evaluation protocol of [14], which is set to recognize words, that include only alphanumeric characters (0-9 and A-Z) and are at least three characters long. Results are not case-sensitive.

4.2 Exploring various models

To recognize the vertical and horizontal texts, we consider various models in Fig.2. (a) depicts a network based on AON. It is an unsupervised method that learns the text direction without using any relevant information. The input image has a square ratio and is encoded with respect to two directions (top-to-bottom and left-to-right). The clue network (CN) identifies the text direction and focuses on that direction. (b) depicts an unsupervised method, which is our baseline network. It differs from AON because the input image is observed to be rectangular. We hypothesized that a rectangle would be more efficient than a square because we only need to consider two directions, i.e., horizontal and vertical. In fact, the vertical text images are transposed. This hypothesis was proved via experiments that depicted an approximate performance improvement of 2–3%p. (c) depicts that DEM (S), our proposed supervised module, provides the text direction to the network. (d) depicts that SAN, our second proposed supervised network, works selectively depending on the text direction.

Method	IIIT5k		SVT		IC03		IC13	IC15	SVTD	VTD
	50	None	50	None	50	None	None	None	None	None
ABBY[14]	24.3	-	35.0	-	56.0	-	-	-	-	-
Wang <i>et al.</i> [14]	-	-	57.0	-	76.0	-	-	-	-	-
Mishra <i>et al.</i> [15]	64.1	-	73.2	-	81.8	-	-	-	-	-
Novikova <i>et al.</i> [21]	64.1	-	72.9	82.8	-	-	-	-	-	-
Wang <i>et al.</i> [22]	-	-	70.0	-	90.0	-	-	-	-	-
Bissacco <i>et al.</i> [23]	-	-	90.4	-	-	-	87.6	-	-	-
Goel <i>et al.</i> [24]	-	-	77.3	-	89.7	-	-	-	-	-
Alsharif <i>et al.</i> [25]	-	-	74.3	-	93.1	-	-	-	-	-
Almazan <i>et al.</i> [26]	91.2	-	89.2	-	-	-	-	-	-	-
Yao <i>et al.</i> [27]	80.2	-	75.9	-	88.5	-	-	-	-	-
Su and Lu <i>et al.</i> [28]	-	-	83.0	-	92.0	-	-	-	-	-
Jaderberg <i>et al.</i> [29]	95.5	-	93.2	71.7	97.8	89.6	81.8	-	-	-
Rodriguez <i>et al.</i> [30]	76.1	-	-	70.0	-	-	-	-	-	-
Gordo <i>et al.</i> [31]	93.3	-	91.8	-	-	-	-	-	-	-
Lee <i>et al.</i> [8]	96.8	78.4	96.3	80.7	97.9	88.7	90.0	-	-	-
Shi <i>et al.</i> [6]	97.6	78.2	96.4	80.8	98.7	93.1	90.8	-	-	-
Shi <i>et al.</i> [12]	96.2	81.9	95.5	81.9	98.3	90.1	88.6	-	-	-
Baoguang <i>et al.</i> [12]	96.2	81.9	95.5	81.9	98.3	90.1	88.6	-	-	-
Chen-Yu <i>et al.</i> [8]	96.8	78.4	96.3	80.7	97.9	88.7	90.0	-	-	-
Cheng <i>et al.</i> [1]	99.6	87.0	96.0	82.8	98.5	91.5	92.3	68.2	-	-
baseline	96.5	82.3	93.3	83.4	96.5	93.1	91.5	67.9	87.0	62.0
DEM	97.5	83.3	93.8	83.6	98.1	93.5	92.7	69.7	87.9	65.2
SAN	97.1	83.0	93.7	85.5	98.5	93.3	92.5	68.0	87.6	64.8
ResNet+DEM	98.6	85.8	94.1	86.7	98.3	94.0	93.7	73.1	90.5	70.9
ResNet+SAN	98.0	86.0	94.6	87.4	98.5	94.1	94.2	73.5	89.6	66.9

Table 1: Results of recognition accuracy on general benchmarks. “50” are lexicon sizes, “None” means lexicon-free.

4.3 Implementation Details

The input images in our experiment are resized to 100×32 for VGG or 256×32 for ResNet[11] and have one channel (gray-scale). The pixel values are normalized to the range (-0.5, 0.5) before they are provided as input to the network. All the LSTM cells used in our experiments have 512 memory units, and the dropouts are not used. The RMSProp method is used to conduct training with a batch size of 512. The initial learning rate is 0.001 and decreases by a factor of 0.9 as the validation errors stop decreasing for 10 iterations. Gradient clipping is applied at a magnitude of 5, and training is terminated after 100,000 iterations. To train our network, 8 million synthetic data released by[19] and 4 million synthetic instances (excluding images that contain non-alphanumeric characters) cropped from 90,000 images[20] were used. In addition, we modified the open source by[20] to generate 1 million synthetic vertical data for training and 3000 data for testing. We did not apply any data augmentation processes to the training dataset. Because we implemented our method using TensorFlow v1.6, CUDA9.0, and cuDNN v7.0, the proposed model can be GPU-accelerated.



Fig. 5: Results of attention and recognition for horizontal and vertical characters.

The experiments were performed on a workstation with one Intel Xeon(R) E5-2650 2.20GHz CPU and one NVIDIA Tesla M40 GPU.

5 Results

Fig.4 compares the performance of the baseline with that of the proposed method. As exhibited by the results, our baseline performed better than the AON-based method in terms of speed and accuracy. This is because an input image of 100×32 is effective while recognizing the horizontal and vertical characters that maintain an image input of 100×100 . Additionally, it was better at recognizing characters while using the DEM and SAN methods of supervised learning than while using the baseline of unsupervised learning. In most of the benchmark datasets, DEM performed slightly better than SAN.

Table1 compares the performance of the state-of-the-art method and the proposed method. The combination of a ResNet-based CNN backbone exhibits the same effect as that exhibited by DEM and SAN. We achieved state-of-the-art performance for SVT, IC13, IC15, SVDT, and VDT142 based on a lexicon-free evaluation protocol. This exhibits that our algorithm works better in real world environments when there is no lexicon. Fig.5 depicts the results of attention and recognition for horizontal and vertical characters, both of which confirm that the method works well. The main purpose of our study is to compare the performance between a supervised proposed method and an unsupervised baseline method for text direction. We would like to emphasize that our model not only recognizes vertical text but also recognizes horizontal text as compared with existing methods.

6 Conclusions

We presented a unified network to simultaneously deal with the horizontal and vertical texts. Our proposed DEM and SAN modules are supervised methods and exhibit better perceived performances than those exhibited by the previous methods. We achieved state-of-the-art results with respect to popular benchmark datasets and verified that our model works for a real vertical text dataset (VTD142).

References

1. Cheng, Z., Liu, X., Bai, F., Niu, Y., Pu, S., Zhou, S.: Arbitrarily-oriented text recognition. CoRR **abs/1711.04226** (2017)

2. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in neural information processing systems. (2012) 1097–1105
3. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2015) 1–9
4. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2016) 770–778
5. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: European conference on computer vision, Springer (2016) 630–645
6. 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 transactions on pattern analysis and machine intelligence* **39** (2017) 2298–2304
7. Liang, M., Hu, X.: Recurrent convolutional neural network for object recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2015) 3367–3375
8. Lee, C.Y., Osindero, S.: Recursive recurrent nets with attention modeling for ocr in the wild. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2016) 2231–2239
9. Wojna, Z., Gorban, A., Lee, D.S., Murphy, K., Yu, Q., Li, Y., Ibarz, J.: Attention-based extraction of structured information from street view imagery. arXiv preprint arXiv:1704.03549 (2017)
10. Wang, J., Hu, X.: Gated recurrent convolution neural network for ocr. In: Advances in Neural Information Processing Systems. (2017) 335–344
11. Cheng, Z., Bai, F., Xu, Y., Zheng, G., Pu, S., Zhou, S.: Focusing attention: Towards accurate text recognition in natural images. *CoRR* **abs/1709.02054** (2017)
12. Shi, B., Wang, X., Lv, P., Yao, C., Bai, X.: Robust scene text recognition with automatic rectification. *CoRR* **abs/1603.03915** (2016)
13. Liu, X., Liang, D., Yan, S., Chen, D., Qiao, Y., Yan, J.: Fots: Fast oriented text spotting with a unified network. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2018) 5676–5685
14. Wang, K., Babenko, B., Belongie, S.: End-to-end scene text recognition. In: Computer Vision (ICCV), 2011 IEEE International Conference on, IEEE (2011) 1457–1464
15. Mishra, A., Alahari, K., Jawahar, C.: Scene text recognition using higher order language priors. In: BMVC-British Machine Vision Conference, BMVA (2012)
16. Lucas, S. M., P.A.S.L.T.A.W.S..Y.R.: Icdar 2003 robust reading competitions. *IEEE* (2003) 682
17. Karatzas, D., S.F.U.S.I.M.i.B.L.G.M.S.R.D.L.H.L.P.: Icdar 2013 robust reading competition. In *Document Analysis and Recognition (ICDAR)* (2013) pp. 1484–1493
18. Karatzas, D., Gomez-Bigorda, L., Nicolaou, A., Ghosh, S., Bagdanov, A., Iwamura, M., Matas, J., Neumann, L., Chandrasekhar, V.R., Lu, S., et al.: Icdar 2015 competition on robust reading. In: *Document Analysis and Recognition (ICDAR), 2015 13th International Conference on, IEEE* (2015) 1156–1160
19. Jaderberg, M., Simonyan, K., Vedaldi, A., Zisserman, A.: Synthetic data and artificial neural networks for natural scene text recognition. *CoRR* **abs/1406.2227** (2014)

10 C. Choi et al.

20. Gupta, A., Vedaldi, A., Zisserman, A.: Synthetic data for text localisation in natural images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2016) 2315–2324
21. Novikova, T., Barinova, O., Kohli, P., Lempitsky, V.: Large-lexicon attribute-consistent text recognition in natural images. In: European Conference on Computer Vision, Springer (2012) 752–765
22. Wang, T., Wu, D.J., Coates, A., Ng, A.Y.: End-to-end text recognition with convolutional neural networks. In: Pattern Recognition (ICPR), 2012 21st International Conference on, IEEE (2012) 3304–3308
23. Bissacco, A., Cummins, M., Netzer, Y., Neven, H.: Photoocr: Reading text in uncontrolled conditions. In: Proceedings of the IEEE International Conference on Computer Vision. (2013) 785–792
24. Goel, V., Mishra, A., Alahari, K., Jawahar, C.: Whole is greater than sum of parts: Recognizing scene text words. In: Document Analysis and Recognition (ICDAR), 2013 12th International Conference on, IEEE (2013) 398–402
25. Alsharif, O., Pineau, J.: End-to-end text recognition with hybrid hmm maxout models. arXiv preprint arXiv:1310.1811 (2013)
26. Almazán, J., Gordo, A., Fornés, A., Valveny, E.: Word spotting and recognition with embedded attributes. IEEE transactions on pattern analysis and machine intelligence **36** (2014) 2552–2566
27. Yao, C., Bai, X., Shi, B., Liu, W.: Strokelets: A learned multi-scale representation for scene text recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2014) 4042–4049
28. Su, B., Lu, S.: Accurate scene text recognition based on recurrent neural network. In: Asian Conference on Computer Vision, Springer (2014) 35–48
29. Jaderberg, M., Simonyan, K., Vedaldi, A., Zisserman, A.: Deep structured output learning for unconstrained text recognition. arXiv preprint arXiv:1412.5903 (2014)
30. Rodriguez-Serrano, J.A., Gordo, A., Perronnin, F.: Label embedding: A frugal baseline for text recognition. International Journal of Computer Vision **113** (2015) 193–207
31. Gordo, A.: Supervised mid-level features for word image representation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. (2015) 2956–2964