Unambiguous Scene Text Segmentation with Referring Expression Comprehension

Xuejian Rong, *Student Member, IEEE,* Chucai Yi, *Member, IEEE,* and Yingli Tian†, *Fellow, IEEE*

Abstract—Text instance provides valuable information for the understanding and interpretation of natural scenes. The rich, precise high level semantics embodied in text could be beneficial for understanding the world around us, and empower a wide range of real-world applications. While most recent visual phrase grounding approaches focus on general objects, this paper explores extracting designated texts and predicting unambiguous scene text segmentation mask, i.e. scene text segmentation from natural language descriptions (referring expressions) like *orange text on a little boy in black swinging a bat*. The solution of this novel problem enables accurate segmentation of scene text instances from the complex background. In our proposed framework, a unified deep network jointly models visual and linguistic information by encoding both region-level and pixel-level visual features of natural scene images into spatial feature maps, and then decode them into saliency response map of text instances. To conduct quantitative evaluations, we establish a new scene text referring expression segmentation dataset: *COCO-CharRef*. Experimental results demonstrate the effectiveness of the proposed framework on the text instance segmentation task. By combining image-based visual features with language-based textual explanations, our framework outperforms baselines that are derived from state-of-the-art text localization and natural language object retrieval methods on *COCO-CharRef* dataset.

Index Terms—Natural Language Description, Text Detection, Text Retrieval, Text Recognition, Deep Neural Network, Referring Expression

1 Introduction

As a category of self-described object entities, text instances such as characters, words, and sentences in a scene image provide one of the most concise and accurate cues to help people understand and interpret natural scenes. Reading text information from scene natural images, namely scene text extraction, could play an important role in image search [1], [2], instant translation [3], robot navigation and self-driving autonomous system [4], industrial automation [5], augmented reality [6], and many other vision-language applications.

† *Corresponding author.*
X. Rong and Y. Tian are with the Department of Electrical Engineering, The City College, City University of New York, New York, NY, 10031.
E-mail: {xrong, ytian}@ccny.cuny.edu.
C. Yi received his Ph.D. degree in computer science in 2014 from The Graduate Center, City University of New York. He is now working at Google on augmented reality, Mountain View, CA, 94043.
E-mail: gschucai@gmail.com

This material is based upon work supported by the National Science Foundation under award numbers IIS-1400802 and EFRI-1137172. Manuscript received May 15, 2019.

Characterness, a.k.a. scene text segmentation mask, as a measure of scene text saliency in natural images first introduced in [7], has been widely used similar to *objectness* for object saliency detection. One main motivation for predicting characterness, a.k.a. scene text segmentation mask, is the fact that *text attracts human attention*, even when amongst a cluttered background [7]. This has been shown by a range of authors including Judd et al. [8] and Cerf et al. [9] who verified that humans tend to focus on text in natural scenes. Also, as indicated in Bylinskii et al. [10] on the research of saliency prediction and human eye fixations, text instances are widely existing in the daily life and always attract much human visual attention, and quantitatively accounts 29% out of all under-predicted outdoor labels by DeepFix [11] model in the CAT2000 dataset [12].

The *referring expression*-based text saliency and segmentation has several significant benefits for a variety of scene understanding tasks. Compared with a set of unordered text bounding boxes predictions traditionally generated by scene text detection approaches, a precise pixel-wise segmentation of target text instances is more valuable because: 1) First, the text bounding boxes of annotated text regions were usually contaminated by background objects, especially for the text instances in irregular shapes. 2) Second, a predicted text bounding box did not completely match its text instance in most cases, and many methods would add post processing to adjust their positions and sizes. These defective text regions would drastically affect the following text recognition, transcription, and other associated tasks depending on the accuracy of text regions. In contrast, segmentation-based pixel-wise text extraction methods could generate more accurate and useful text annotations. 3) Accurate pixel-level text segmentation mask is also necessary for many applications such as scene text image inpainting/completion and the automatic removal of image annotations or video subtitles.

However, even though scene text segmentation could provide a more comprehensive extraction output with respect to scene text detection, not all text instances are equivalently important in a scene image. The informativeness of a text instance in the background context, and the allure of the text instance on its own, can affect how long individual observers fixate on it, and what proportion observers gaze at [10]. Most previous approaches of scene text extraction merely regarded text instances as one generic category of objects, and attempted to exhaustively generate their lo-

Fig. 1: In this work we approach the novel problem of *unambiguous scene text segmentation*. Unlike traditional semantic image segmentation and scene text detection/segmentation, the proposed framework jointly models visual features and context descriptions to predict specific scene text segmentation mask, which results in unambiguous pixel-wise scene text segmentation from natural language expressions. From left to right: 1) original image; 2) semantic image segmentation (on category *human*); 3) region-level scene text detection represented with bounding boxes; 4) scene text segmentation mask from referring expression query *"orange text on a little boy in black swinging a bat"*.

cation (*spatial*) and character sequence (*literal*) information while ignoring higher level *semantic* and *contextual* information. It means that text instances were treated the same as other objects for scene understanding and description, even though they are probably straightforward explanations of surrounding context environment and semantically self-described.

As we know, text information could help people understand surrounding environments only if the relationship between text instances and the context environment were perceived. It means that user-specific text instances could be more valuable than a set of unorganized text instances from scene images. Therefore, this paper addresses the following problem: *from an image containing scene text instances, and a referring expression that describes specific text instances with their context, how to predict the text regions corresponding to the referring expression and further segment text instances in pixel level* (i.e., *unambiguous scene text segmentation / characterness prediction)*. For example, as shown in Figure 1, for the query {*orange text on a little boy in black swinging a bat*} our framework would predict a segmentation mask that covers the orange word exactly on the target entity, but not the others. This problem is related to but different from recent work on *natural language object retrieval/segmentation* [13], [14], *referring expression comprehension* [15], and *phrase grounding* [16], [17], since the highly variant appearance, scale, and density, and the style of self-description tell the significant difference between text instances and generic objects. Besides, the recent semantic segmentation approaches were not specifically designed for text instances which are usually inconsistent, fragmental, and lack of well-defined shape/boundary and amorphous regions, so they hardly work well on text segmentation. Our proposed model outperforms the state-of-the-art approach on the newly collected benchmark dataset, named COCO-CharRef, for this novel task.

High-quality and user-specific scene text segmentation from referring expressions can underpin many vision-language applications which rely on natural language interfaces, such as communicating with a grocery shopping aid for blind or visually impaired users (e.g., *Alexa, please read me the price of non-fat CHOBANI Greek yogurt on the top shelf*), or interacting with video editing software (e.g., *Premiere, please transcribe/mosaic all the identity information if photo IDs or credit cards appear in this video*). In addition, it is a good testbed for research in the area of vision and language systems on scene text images. The proposed framework could also help dramatically boost the efficiency of the whole scene text reading system while avoiding the exhaustive search and recognition of all text instances.

Contributions. The contributions of our work are three-fold.

- We propose a new framework which accurately segments pixel-level text instances from context description by the prediction of unambiguous scene text segmentation mask.
- The relationship between text instances and their context are explored and modeled in this paper as {*text-predicate-context*} triplets, where the *predicate* can be *spatial, preposition, comparative, visual attributes* or their combinations. While the number of such triplet relationship phrases could grow combinatorially, the proposed model tends to handle a large number of relations by sharing parameters among them. For instance, a single "on" classifier can be learned to recognize both {*text on bags*} and {*text on dogs*}, even when {*text on dogs*} has never been seen in training.
- A new large-scale benchmark dataset is constructed to evaluate the performance of the new task and other text-based scene understanding approaches. And the proposed approach achieves remarkable performance improvement compared with several strong baseline methods.

The remainder of this paper is organized as follows: Sec. 2 reviews related work recently published. In Sec. 3, the baseline methods and proposed model are described in detail. The construction of new benchmark dataset, experimental results, and analytic discussions are presented in Sec. 4. We conclude our paper and describe future work in Sec. 5.

2 Related Work

Text detection and segmentation, word recognition, word image retrieval, image captioning and description, visual question answering, generation and comprehension of referring expressions, image-language alignment, phrase grounding, and visual relation reasoning/modeling could be considered as diverse subtasks of the supertask of systematic visual and linguistic interaction, which jointly models the natural language and scene image information. We

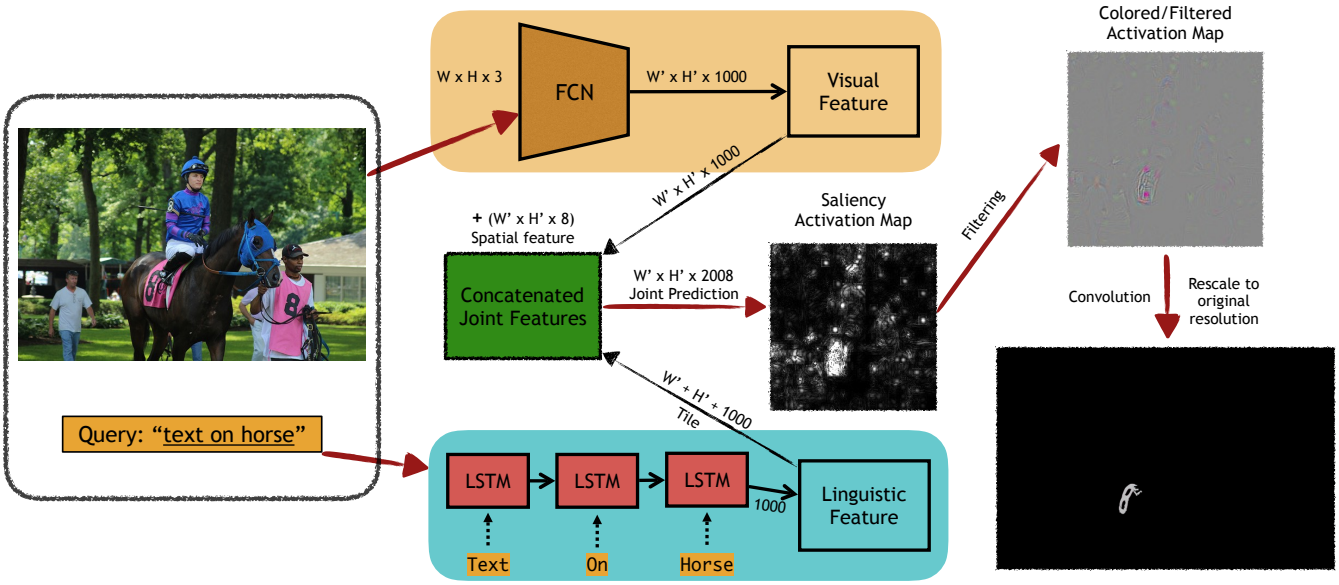

Fig. 2: The Architecture of Our Proposed Model. Given an input image containing text instances, the model encodes the image-based data along with the language-based query of referring expression into convolutional and recurrent features, then decodes them into the saliency response map through a fully convolution classification network. After upscaling, the final text segmentation masks are generated from the saliency response map.

discuss the related work with the proposed new task as follows.

Text extraction in the wild. Text extraction from the natural scene has been increasingly popular in academic research and practical applications. Most existing text extraction methods localize scene text instances with word-level bounding box annotations, although they depend on multiple schemes, including sliding-window based [18], [19], [20], connected components-based [21], [22], [23], or deep neural network-based [24], [25], [26], [27], [28], [29], [30], [31], [32], [33], [34], [35], [36], [37]. However, bounding box-based text regions generated by traditional text detection approaches are usually not tight enough and contain many background outliers (especially when text instances are in irregular or curved shapes), which makes the following scene text recognition difficult. In contrast, segmentation-based pixel-wise text extraction methods could generate more accurate and useful text annotations.

In our proposed framework, by modeling both visual and linguistic information in context, text instances are able to be segmented from the scene image with pixel-wise labeling. Most conventional solutions of text segmentation employed a bottom-up pipeline with heuristic features [23], [38], [39] which were not robust and reliable. It results in the high dependency of low-level image filtering. Even though deep neural network could substantially improve generic semantic object segmentation [40], [41], [42], text segmentation from scene images with complex background is still a problem to be addressed. The main challenge is that text instances always reveal high variants of appearances, scales, and structures that are difficult to model. Some recent approaches [43], [44], [45], [46], [30], [34] based on deep networks presented well-performed text segmentation of zoomed-in views, but could not satisfyingly handle cluttered environments with large-scale and a wide variety of context information from natural scenes. In our proposed framework, text instances and their context are jointly modeled. The relationship modeling of a text instance with its surrounding objects further benefits the characterness prediction and pixel-wise segmentation.

Fully convolutional networks for segmentation. Fully Convolutional Networks (FCNs) are deep neural networks consisting of only convolutional (and pooling) layers. FCNs pioneered the use of deep learning for semantic segmentation, and are still one of the most commonly adopted backbone of the state-of-the-art methods for semantic segmentation over a pre-defined set of semantic categories [41], [47], [48]. One advantage of FCNs is that spatial information is consistently preserved, which makes this kind of networks suitable for segmentation tasks that require spatial grid output. In our framework, FCNs are used to extract features and generate segmentation labels, with the purpose of handling image-based text instances at both region and pixel levels.

Alignment of images with language. Learning correspondences between sentence structure and image regions has been explored with the visual-semantic alignment. This architecture has been used for applications in image retrieval and caption generation [49], [50]. With new datasets proposed which provide bounding box-level natural language annotations [15], recent work has also investigated region-wise image captioning and description for the tasks of natural language object retrieval [13], dense captioning [51], scene graph parsing [52], and visual common sense reasoning [53]. Our proposed framework has a similar idea that aligns a language triplet with regions of pixels in the image. Typically, existing approaches do not explicitly represent relations between noun phrases in a sentence to improve visual-semantic alignment. We believe that understanding these relations will lead to better scene understanding including phrase grounding and comprehension, and scene

graph generation and reasoning.

Visual relation modeling. Triplet learning has been addressed in various tasks such as mining typical relations (knowledge extraction) [54], reasoning [55], object detection [56], or image retrieval [57]. In this work, we address the task of relationship modeling in scene text segmentation from language based explanations. Early work on human-object interactions [58] models the triplet in the form $(person, action, object)$. Recently, the work in [59] tried to generalize the similar setting to non-human subjects by developing a language model sharing knowledge among visual detections related with each other. Inspired by the idea but different from these approaches, we restrict the $subject$ to be a text instance and cover a broader class of predicates that include prepositions and comparatives. In our work this combinatorial challenge can be addressed by developing a new visual representation with better generalization into unseen triplets $\{$*text-predicate-object*$\}$ *and* without depending on a strong language model.

Grounding visual explanations. Our proposed framework is an innovative combination of the recent work on object localization and segmentation from natural language descriptions, i.e., referring expression comprehension. In those work, the task is to localize/segment a target object in a scene based on its natural language referring expression (by drawing a bounding box over it, or pixel wisely assigning the foreground label to it). The methods of [13] and [15] are built upon image captioning frameworks such as LRCN [60] or mRNN [61], and localize objects by selecting the bounding box where the expression has the highest probability. The authors of [62] firstly proposed a natural language based scene text extraction methods, but the framework is not trained end-to-end and cannot output pixel-wise text annotations. In [16], the authors proposed a model to localize a textual phrase by attending to a region on which the phrase can be best reconstructed. In [63], a joint embedding space of visual features and words is learned to localize target object by searching the closest region in the joint embedding space. [14] proposes an end-to-end training method for generating object segmentation mask from natural language descriptions. The proposed model encodes the given expression into a real-valued vector using LSTM networks [64], and extracts a spatial feature map from the image using a Convolutional Network. Then it performs pixel-wise classification based on the encoded referring expression and feature map to output an image mask covering the visual entity described by the expression. Liu et al. [65] further propose to learn the word-to-image interaction instead of modeling image and sentence features independently. The proposed method achieves top results on general object segmentation with language explanations, and also shows that the combination of visual and linguistic features for scene text segmentation is worth exploring.

To the best of our knowledge, all previous methods of natural language based detection and retrieval can only return a bounding box or segmentation mask of the generic objects, and no prior work has learned to directly segment text instances given a natural language description as a query. The most related approach with our work is the recent unambiguous text localization and retrieval model (CRTR) proposed in [62], which pioneered the task of natural language based scene text detection and retrieval. However, CRTR is only capable of generating bounding box level prediction, and relies on preceding results from a sequential text detection model. Also, the framework proposed in [62] are not end-to-end trainable (though being end-to-end evaluated). In comparison, we construct strong baselines with the foreground segmentation based on the bounding boxes obtained by combining the state-of-the-art text detection method [27] and natural language retrieval method [13]. The CRTR model proposed in [62] is also regarded as one start-of-the-art to compare with (see details in Sec. 4).

3 Proposed Model

In this section, we first describe our proposed framework in detail from Sec. 3.1 to Sec.3.3, including text image visual feature encoding, referring expression linguistic feature encoding, fusion feature decoding, and saliency response map upsampling. Then several effective baseline methods are presented, which are derived from previous state-of-the-art approaches in Sec. 3.4.

3.1 Spatial feature map extraction

As shown in Figure 3, given an image with scene text instances, our proposed framework first computes an effective feature representation that is able to encode visual appearance of text characters and their relations with surrounding objects. This feature representation should preserve the spatial information to enable the correct spatial prediction of a segmentation mask. This is accomplished by a fully convolutional network model similar to FCN-32s [41], where the image is fed through a series of convolutional (and pooling) layers obtaining a feature map containing encoded spatial information fused from different levels. In our framework, the network is further modified to encode both region-wise and pixel-wise information, resulting in a more effective feature representation compatible with varieties of text instances.

Given an input scene text image of size $W \times H$, we apply a convolutional network to the image and obtain a $W' \times H'$ spatial feature map, with each position on the feature map containing D_{im} channels (D_{im} dimensional local descriptors). At each position of the spatial feature map, the D_{im} dimensional local descriptor is further L2-normalized in order to obtain a more robust feature representation with respect to degradations. In this way, we can extract a $W' \times H' \times D_{im}$ spatial feature map as the representation of each image.

To allow the model to reason about spatial relationships, two extra channels are inserted into the feature maps: the x and y coordinates of each spatial location. The top left corner and the bottom right corner of a feature map are represented as $(-1, -1)$ and $(+1, +1)$, respectively. In this way, we obtain a $w \times h \times (D_{im} + 2)$ feature representation containing both local image descriptors and spatial coordinates.

In our implementation, the VGG-16 architecture [66] is adopted as a fully convolutional network by transforming fully-connected layers `fc6`, `fc7` and `fc8` to convolutional

layers, which outputs $D_{im} = 1000$ dimensional local descriptors. The resulting feature map size is $W' = W/s$ and $H' = H/s$, where $s = 32$ is the pixel stride on fc8 layer output. It means that a unit on the spatial feature map has a large enough receptive field of 384 pixels, which aggregates the information of context concepts for text instances from neighboring regions. This design can help reason about interaction between visual text entities and context concepts.

3.2 Encode referring expressions with LSTM

For the input natural language expression that describes a scene text region, we model the query sequence into a vector because it is more efficient to process fixed-length vectors than variable-length sequences. In our encoder, for the natural language expression in a sequence-to-sequence manner [67], each word is first embedded into a vector through a word embedding matrix, and then a recurrent LSTM [64] network with D_{text} dimensional hidden state is used to scan through the embedded word sequence. For a text sequence $S = (w_1, ..., w_T)$ with T words (where w_t is the vector embedding for the t-th word), at each iteration t, the LSTM network takes as input the embedded word vector w_t from the word embedding matrix. At the final iteration $t = T$ when the LSTM network has seen the whole text sequence, the hidden state h_T in LSTM network is regarded as the encoded vector representation of the expression. Similar to the encoding of spatial features in Sec. 3.1, we also L2-normalize the D_{text} dimensions in h_T and set $D_{text} = 1000$ in our implementation.

Specifically, every word w_t is one-hot encoded and mapped to a word embedding $\mathbf{w}_t$. The entire sentence is then encoded with an LSTM into a vector $\mathbf{h}_T$ of size D_{text}, where $\mathbf{h}_t$ represents the hidden state of LSTM at time step t:

$$\text{LSTM} : (\mathbf{w}_t, \mathbf{h}_{t-1}, \mathbf{c}_{t-1}) \rightarrow (\mathbf{h}_t, \mathbf{c}_t), \tag{1}$$

$$\mathbf{c}_t = \mathbf{f} \odot \mathbf{c}_{t-1} + \mathbf{i} \odot \mathbf{g}, \tag{2}$$

$$\mathbf{h}_t = \mathbf{o} \odot \tanh(\mathbf{c}_t), \tag{3}$$

where n is the size of the LSTM cell. $\mathbf{c}_t$ is the memory states at time step t. The vector $\mathbf{h}_T$ is then concatenated with the image features and spatial coordinates at all locations to produce a $W' \times H' \times (D_{im} + D_{text} + 8)$ tensor.

3.3 Spatial classification and upsampling

After extracting the spatial feature map from a scene image in Sec. 3.1 and the encoded referring expression h_T in Sec. 3.2, we need to determine whether a spatial location on the feature map belongs to the foreground, which denotes the text instances described by the natural language expression. In our framework, this is accomplished by a fully convolutional classifier over the local image descriptor and the encoded referring expression. We first tile and concatenate h_T to the local descriptor at each spatial location in the spatial grid to obtain a $W' \times H' \times D^*$ (where $D^* = D_{im} + D_{text} + 2$) spatial map containing both visual and linguistic features. Then, a two-layer classification network is trained with a D_{cls} dimensional hidden layer, which takes as input the D^* dimensional representation and outputs a score to indicate the probability that a spatial location belongs to the target image region. We set $D_{cls} = 500$ in our implementation.

This classification network contains two 1×1 convolutional layers (with ReLU none linearity between them) and is applied to the underlying $W' \times H'$ feature map. The fully convolutional classification network outputs a $W' \times H'$ coarse *low-resolution scene text response map* containing classification scores, which can be regarded as a low-resolution segmentation of the text instances described by the referring expression.

In order to obtain a segmentation mask with higher resolution, upsampling is further performed by deconvolution (swapping the forward and backward pass of convolution operation) [41], [48]. In our framework a $2s \times 2s$ deconvolution filter is used with stride $s = 32$ for the VGG-16 network architecture, similar to the FCN-32s [41]. The deconvolution filter produces a $W \times H$ *high resolution saliency response map* that has the same size as the input image, and the values on this map represent the confidence of determining if a pixel belongs to target object. We use the pixel-wise classification results (i.e., whether or not a value on the response map is greater than 0) as the final segmentation prediction.

In training phase, a training sample is a tuple (I, T, M), where I is an scene text image, T is a natural language expression describing specific text region(s) within that image, and M is a binary segmentation mask of the corresponding text region(s). The training loss function is defined to be average pixel-wise loss, as presented in the following equation.

$$Loss = \frac{1}{WH} \sum_{i=1}^{W} \sum_{j=1}^{H} L(v_{ij}, M_{ij}) \tag{4}$$

where W and H are image width and height, v_{ij} is the response value (score) on the high resolution saliency response map and M_{ij} is the binary ground-truth label at pixel (i, j). L is the per-pixel weighted logistic regression loss as follows

$$L(v_{ij}, M_{ij}) = \begin{cases} \alpha_f \log(1 + \exp(-v_{ij})) & \text{if } M_{ij} = 1 \\ \alpha_b \log(1 + \exp(v_{ij})) & \text{if } M_{ij} = 0 \end{cases} \tag{5}$$

where α_f and α_b are loss weights of foreground and background pixels respectively. In experiments, we observe that training converges faster in the presence of higher loss weights for foreground pixels, and we set $\alpha_f = 3$ and $\alpha_b = 1$ in $L(v_{ij}, M_{ij})$.

All the parameters of the network are initially derived from a VGG-16 network [66], which is pretrained on the 1000-class ILSVRC classification task [68], and fine-tuned on newly collected *COCO-CharRef* dataset. The deconvolution filter for upsampling is initialized from bilinear interpolation. All the other parameters in our proposed framework, including the word embedding matrix, the LSTM parameters and the classifier parameters, are randomly initialized. The whole network is trained with standard back-propagation using SGD with momentum.

3.4 Baseline methods

As far as we know, no previous work was able to directly predict the pixel-wise segmentation of text instances based on referring expression in natural language. To evaluate

our method, we construct several baseline methods of text segmentation, and compare their performance with our proposed text segmentation framework.

Classification over *Characterness* proposals. To re-implement this baseline method, we first extract a set of text segmentation proposals using the original Characterness method [7], and then train a binary classifier to determine whether or not a text segmentation proposal matches the expression. First, the encoded query is concatenated with visual features that are extracted from each proposal using a VGG-16 [66] network pretrained on [68]. Then, the concatenated features are adopted to train a classifier to predict a segmentation proposal to be foreground or background. In contrast, our framework employs fully convolutional network to perform pixel-wise classification in an end-to-end way, without relying on the generation of candidate text segmentations.

Foreground segmentation from text region proposals. To re-implement this baseline method, a recent text localization method (TextBoxes [27]) is employed to exhaustively obtain all potential text bounding boxes, which are then scored and ranked by the Spatial Context Recurrent ConvNet (SCRC) [13] based on the given referring expression. The state-of-the-art scene text detector "TextBoxes" [27] is derived from the Single Shot Detector [69]. It achieves highly competitive results on standard text detection benchmarks, for example, f-measure 0.85 on ICDAR 2011/13 dataset. We select the top 50 text region proposals from [27] and input them into SCRC. Next, the foreground segmentation is extracted from the top text bounding boxes using GrabCut [70]. SCRC localizes a referring expression by scoring/ranking candidate bounding boxes based on an image captioning model. Our framework takes from a testing image the top-ranked 50 candidate bounding boxes that best match the associated testing referring expression, and performs two types of evaluations, either using the globally highest scoring bounding box as the segmentation, or the foreground segmentation resulted from GrabCut [70].

4 Experiments

4.1 Dataset Construction

Although there are many datasets for the evaluations of scene text segmentation, semantic object segmentation, and image captioning respectively, no benchmark dataset is available with both pixel-level scene text annotations and corresponding natural language descriptions. The ReferIt dataset [71] has been widely used in image captioning and natural language object retrieval. However, it does not provide any image-based annotations or language-based referring expressions for the scene text instances. The ICDAR dataset (Task 2 for Robust Reading Challenges 2013 and 2015) [72] contains real world images of text on sign boards, books, posters and other objects with pixel-level foreground/background annotations. However, most text bounding boxes from these datasets are in extremely focused view and rarely contain useful context concept entities. Therefore, to evaluate the performance of the baseline methods (as described above in Sec. 3) and our proposed framework in the way of unambiguous scene text segmentation, we establish a new *Characterness Referring Expression* dataset named as *COCO-CharRef*. Specifically, it is built on the basis of several existing externally annotated datasets [1] [73], [74], [75], all of which were further originated from MS COCO [76]

MSCOCO dataset. MSCOCO [76] is an image recognition, segmentation, and captioning dataset. It contains more than $300,000$ images and 2 million instances across 80 object categories. It serves as a fundamental data source for the three datasets introduced below.

COCO-Text dataset. COCO-Text is a large-scale dataset for text detection and recognition in natural images, based on MSCOCO dataset. It contains more than $63,000$ images and $173,000$ text instances. However, it does not provide pixel-level text segmentation annotations and natural language descriptions. Also, there exist many annotated text instances that are illegible for transcription and recognition, resulting from too small text height or too much occlusion. Our dataset ignores these kinds of illegible text samples.

RefCOCO dataset. RefCOCO [75] integrates four earlier referring expression datasets (RefCOCOg, RefCOCO, RefCOCO+, and RefClef), and provides instance level natural language referring expressions. Due to the compensation between COCO-Text and RefCOCO datasets, we select the overlapped part of them to establish half of the new *COCO-CharRef* dataset. The ground truth pixel-level annotations are generated from the ground truth word-level bounding box annotations through GrabCut [70], and further filtered with human assessment. The complete referring expressions for text instances are generated by combining the detected relationship between text instances and context concepts, with the existing referring expressions for those context concepts. The illegible text instances are treated as "Don't Care" regions, and cannot affect the final evaluation performance [72]. In total we get $2,427$ legible text images for final dataset construction.

COCO-Stuff dataset. COCO-Stuff [74] is a dataset which augments the MSCOCO dataset with pixel-level *stuff* annotations. Besides the originally annotated 80 object classes, COCO-Stuff provides $10,000$ complex images which are further annotated with 91 stuff classes including both indoor-stuff (e.g., floor, wall, ceiling) and outdoor-stuff (e.g., building, ground, sky). To enable the accurate pixel-level evaluation of text segmentation, we follow the text rendering methods in [77] and generate synthetic text images. The complete referring expressions are obtained from the intersection of synthetic COCO-Stuff text images and RefCOCO datasets, which are similar to those in COCO-Text. Since scene text mostly appears at urban artificial environments, in the generation of text images, we discard some categories of outdoor objects that should not belong to those environments. In total $2,573$ text images are generated in the established dataset.

New Created COCO-CharRef dataset. In summary, our constructed COCO-CharRef benchmark dataset is built on the COCO-Text, RefCOCO, and COCO-Stuff datasets. Specifically, it contains $5,000$ images ($4,000$ for training and $1,000$ for evaluation), and $84,112$ referring expressions annotated on $19,241$ text regions. To date, the COCO-

1. http://mscoco.org/external/

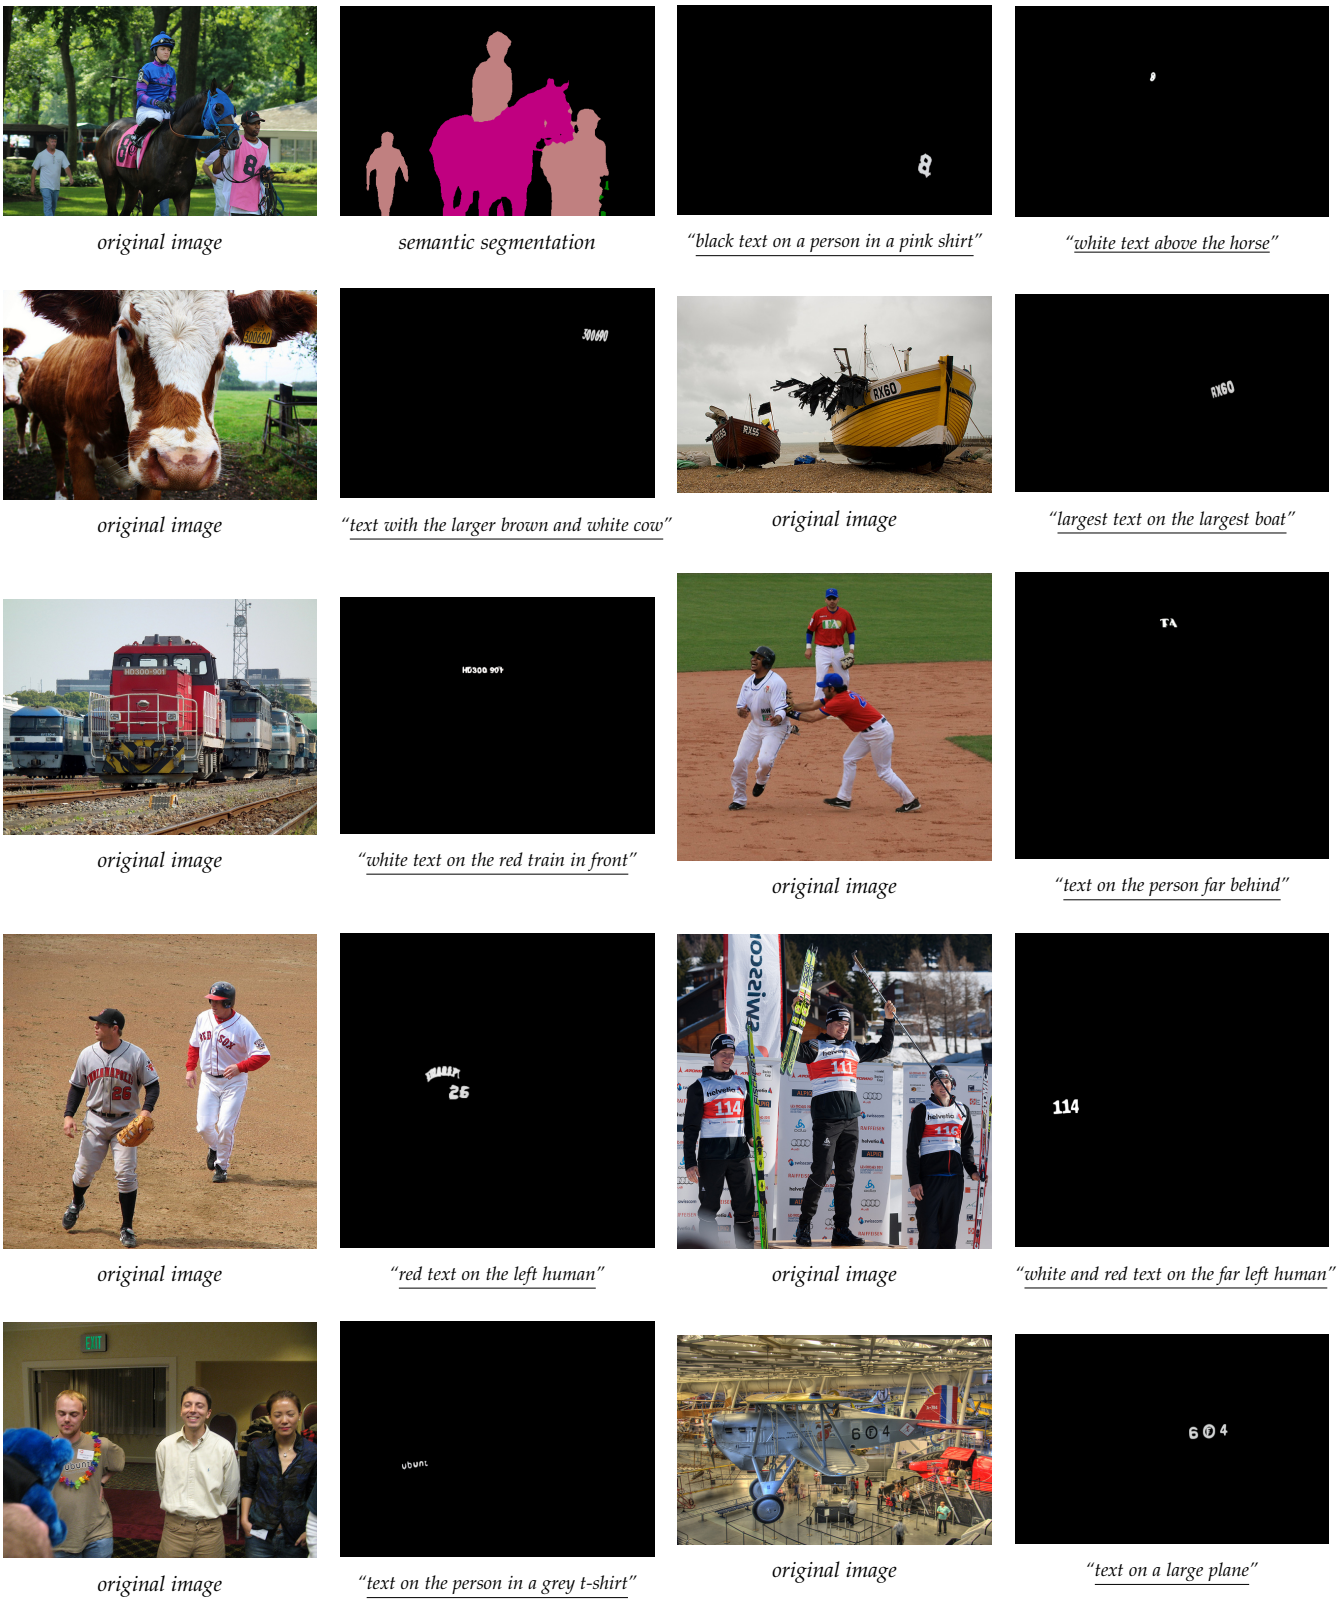

Fig. 3: Examples of unambiguous segmentation of text regions from the *COCO-CharRef* dataset. Text instances in not-high resolutions are often embedded in complex background context and cluttered, but are still successfully segmented accordingly.

CharRef dataset is the biggest available dataset that contains natural language expressions annotated on segmented scene text regions as far as we know. The dataset and related evaluation code will be made available.

Method	@0.5	@0.6	@0.7	@0.8	@0.9	overall IoU
Whole image	0.0%	0.0%	0.0%	0.0%	0.0%	0.7%
Characterness classification [7]	3.5%	1.5%	0.4%	0.3%	0.0%	9.4%
SCRC [13] + TextBoxes [27]	5.3%	2.1%	0.7%	0.0%	0.0%	14.9%
SCRC [13] + TextBoxes [27] + GrabCut	7.6%	3.7%	1.9%	0.4%	0.4%	18.2%
TextSeg [14] (retrained)	11.4%	5.1%	2.6%	0.7%	0.3%	20.5%
CRTR [62] + GrabCut	10.7%	5.8%	3.3%	1.8%	0.6%	25.4%
Proposed: low resolution	9.5%	4.2%	2.7%	1.3%	0.4%	24.7%
Proposed: high resolution	**12.3%**	**6.9%**	**4.6%**	**2.1%**	**0.9%**	**28.1%**

TABLE 1: The performance of our framework and baseline methods on the COCO-CharRef benchmark dataset under precision metric and overall IoU metric. Numbers in bold indicate the best performance.

4.2 Evaluation on COCO-CharRef dataset

Metrics. On the newly-collected COCO-CharRef dataset, we evaluate the performance of our framework and the baseline methods. Two metrics are used for evaluation: the *overall intersection-over-union* (overall IoU) metric and the *precision* metric. In the experiments, the area of a region is represented by the number of pixels at that region. The overall IoU is the total intersection area divided by the total union area (calculated pixel-wisely), where both intersection area and union area are accumulated over all testing samples (notice that a testing sample is a pair of scene image and referring expression as well). Although the overall IoU metric is the standard metric used in PASCAL VOC segmentation [13], our evaluation metric has to measure the accuracy of segmenting foreground text instances specified by the input referring expression from the background. The precision metric goes along with 5 different IoU thresholds from low to high: Precision@X ($X \in \{0.5, 0.6, 0.7, 0.8, 0.9\}$) where Precision@X means the percentage of images with IoU higher than X. Precisely, the precision metric is the percentage of test samples where the IoU between prediction and ground-truth passes the threshold. For example, precision@0.5 is the percentage of expressions where the predicted segmentation overlaps with the ground-truth text region by at least 50% IoU.

Results. The evaluation results of our experiments are summarized in Table 1. By simply returning the whole image, the baseline method only achieves 0.7% overall IoU. This is partially caused by the fact that COCO-CharRef dataset contains some large context regions and the overall IoU metric put more weights on large regions such as walls and buildings. Moreover, a reasonable overall IoU can be obtained through per-word segmentation (Per-word). The prediction of text bounding box as a whole from SCRC [13] (`SCRC + TextBoxes`) obtains better performance than the prediction of the referring expression-based segmentation proposal [7] ("Characterness classification"). Also, it shows that GrabCut [70] obtains better performance of foreground segmentation from SCRC-ranked TextBoxes (`SCRC + TextBoxes + GrabCut`) than just (`SCRC + TextBoxes`).

As to the method of end-to-end object segmentation from language descriptions proposed in [14], we retrain the segmentation network on scene text datasets [72] and conduct the evaluation on the COCO-CharRef dataset. The state-of-the-art referring scene text localization method is also combined with GrabCut with the performance presented. In summary, [14] achieves decent precision on lower settings, but does not perform well on higher settings (P@0.8 and P@0.9). The combined {CRTR + GrabCut} model achieves decent and consistent performance, but still inferior compared with the high-resolution version of the proposed model.

We believe that the Precision metric is more compatible with the task of text instance segmentation from natural language descriptions. In practical applications, user would prefer the specified segmentation from their referring expression query, rather than the detection or segmentation accuracy of all text instances.

As the evaluation results, our proposed framework outperforms all the competing methods by a large margin under both precision metric and overall IoU metric. In Table 1, the second last row ("low resolution") denotes the process of bi-linear upsampling over the coarse response map obtained from the coarse-level model in low-resolution space, and the last row ("high resolution") denotes the performance of the full model in high-resolution space, and the effectiveness of spatial upsampling at low-resolution space. In addition, it demonstrates that the final high-resolution model achieves both higher precision and higher overall IoU, compared with the baseline methods. Some examples of text instance segmentation are illustrated in Figure 3, where the top row depicts different segmentation results from different referring expressions on the same image. Images and other rows present more unambiguous text segmentation results.

Figure 4 illustrates some failure cases on COCO-CharRef dataset, where the IoU between prediction and ground-truth segmentation is less than 50%. In some examples, our framework fails to accurately segment the strokes of text instances, but still produces reasonable text saliency response maps covering the target regions specified by natural language referring expressions. Our framework could be continuously fine-tuned by integrating more training examples associated with the failure cases.

5 Conclusion

In this paper, we have addressed the challenging problem of unambiguous scene text segmentation, i.e. segmenting

Fig. 4: Some failure cases where IoU $< 50\%$ between prediction and ground-truth, due to incorrect region-level prediction or inaccurate pixel-level estimation of the corresponding models.

specific scene text instances described by the referring expressions from natural scenes. In our proposed framework, spatial information and context descriptions of scene text instances benefit from each other. Scene text instances could provide pivotal and precise information for context descriptions of the entire or a region of the scene image, while context description could provide a more user-friendly way to incorporate the extracted text information and its context into practical applications. Experimental results on the newly collected benchmark dataset demonstrate that our framework outperforms baseline methods by a large margin. Our future work will focus on more comprehensively modeling of the relationships between scene text instances and their contexts.

REFERENCES

[1] Georg Schroth, Sebastian Hilsenbeck, Robert Huitl, Florian Schweiger, and Eckehard Steinbach. Exploiting text-related features for content-based image retrieval. In *2011 IEEE International Symposium on Multimedia*, pages 77–84. IEEE, 2011.

[2] Sam S Tsai, Huizhong Chen, David Chen, Georg Schroth, Radek Grzeszczuk, and Bernd Girod. Mobile visual search on printed documents using text and low bit-rate features. In *2011 18th IEEE International Conference on Image Processing*, pages 2601–2604. IEEE, 2011.

[3] Christopher Parkinson, Jeffrey Jacobsen, David Ferguson, and Stephen Pombo. Instant translation system. 2016.

[4] Ruth Schulz, Ben Talbot, Obadiah Lam, Feras Dayoub, Peter Corke, Ben Upcroft, and Gordon Wyeth. Robot navigation using human cues: A robot navigation system for symbolic goal-directed exploration. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1100–1105. IEEE, 2015.

[5] Zhiwei He, Jilin Liu, Hongqing Ma, and Peihong Li. A new automatic extraction method of container identity codes. *IEEE Transactions on intelligent transportation systems*, 6(1):72–78, 2005.

[6] Marc Petter, Victor Fragoso, Matthew Turk, and Charles Baur. Automatic text detection for mobile augmented reality translation. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 48–55. IEEE, 2011.

[7] Yao Li, Wenjing Jia, Chunhua Shen, and Anton van den Hengel. Characterness: An indicator of text in the wild. *IEEE transactions on image processing*, 23(4):1666–1677, 2014.

[8] Tilke Judd, Krista Ehinger, Frédo Durand, and Antonio Torralba. Learning to predict where humans look. In *Computer Vision, 2009 IEEE 12th international conference on*, pages 2106–2113. IEEE, 2009.

[9] Moran Cerf, E Paxon Frady, and Christof Koch. Faces and text attract gaze independent of the task: Experimental data and computer model. *Journal of vision*, 9(12):10–10, 2009.

[10] Zoya Bylinskii, Adrià Recasens, Ali Borji, Aude Oliva, Antonio Torralba, and Frédo Durand. Where should saliency models look next? In *European Conference on Computer Vision*, pages 809–824. Springer, 2016.

[11] Srinivas SS Kruthiventi, Kumar Ayush, and R Venkatesh Babu. Deepfix: A fully convolutional neural network for predicting human eye fixations. *IEEE Transactions on Image Processing*, 26(9):4446–4456, 2017.

[12] Ali Borji and Laurent Itti. Cat2000: A large scale fixation dataset for boosting saliency research. *CVPR workshop*, 2015.

[13] Ronghang Hu, Huazhe Xu, Marcus Rohrbach, Jiashi Feng, Kate Saenko, and Trevor Darrell. Natural language object retrieval. In *Computer Vision and Pattern Recognition (CVPR), 2016 IEEE Conference on*, pages 4555–4564. IEEE, 2016.

[14] Ronghang Hu, Marcus Rohrbach, and Trevor Darrell. Segmentation from natural language expressions. In *European Conference on Computer Vision*, pages 108–124. Springer, 2016.

[15] Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 11–20, 2016.

[16] Anna Rohrbach, Marcus Rohrbach, Ronghang Hu, Trevor Darrell, and Bernt Schiele. Grounding of textual phrases in images by reconstruction. In *European Conference on Computer Vision*, pages 817–834. Springer, 2016.

[17] Mingzhe Wang, Mahmoud Azab, Noriyuki Kojima, Rada Mihalcea, and Jia Deng. Structured matching for phrase localization. In *European Conference on Computer Vision*, pages 696–711. Springer, 2016.

[18] Xu-Cheng Yin, Xuwang Yin, Kaizhu Huang, and Hong-Wei Hao. Robust text detection in natural scene images. *IEEE transactions on pattern analysis and machine intelligence*, 36(5):970–983, 2014.

[19] Qixiang Ye and David Doermann. Text detection and recognition in imagery: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 37(7):1480–1500, 2015.

[20] Lukáš Neumann and Jiří Matas. Efficient scene text localization and recognition with local character refinement. In *Document Analysis and Recognition (ICDAR), 2015 13th International Conference on*, pages 746–750. IEEE, 2015.

[21] Chucai Yi and YingLi Tian. Text string detection from natural

scenes by structure-based partition and grouping. *IEEE Transactions on Image Processing*, 20(9):2594–2605, 2011.

[22] Xu-Cheng Yin, Wei-Yi Pei, Jun Zhang, and Hong-Wei Hao. Multi-orientation scene text detection with adaptive clustering. *IEEE transactions on pattern analysis and machine intelligence*, 37(9):1930–1937, 2015.

[23] Lukáš Neumann and Jiří Matas. Real-time lexicon-free scene text localization and recognition. *IEEE transactions on pattern analysis and machine intelligence*, 38(9):1872–1885, 2016.

[24] Zheng Zhang, Chengquan Zhang, Wei Shen, Cong Yao, Wenyu Liu, and Xiang Bai. Multi-oriented text detection with fully convolutional networks. *CVPR*, 2016.

[25] Z. Tian, W. Huang, T. He, Pan. He, and Y. Qiao. Detecting text in natural image with connectionist text proposal network. *ECCV*, 2016.

[26] Tong He, Weilin Huang, Yu Qiao, and Jian Yao. Accurate text localization in natural image with cascaded convolutional text network. *CoRR abs/1507.05717*, 2016.

[27] Minghui Liao, Baoguang Shi, Xiang Bai, Xinggang Wang, and Wenyu Liu. Textboxes: A fast text detector with a single deep neural network. In *AAAI*, pages 4161–4167, 2017.

[28] Pan He, Weilin Huang, Tong He, Qile Zhu, Yu Qiao, and Xiaolin Li. Single shot text detector with regional attention. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 3047–3055, 2017.

[29] Yuliang Liu and Lianwen Jin. Deep matching prior network: Toward tighter multi-oriented text detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1962–1969, 2017.

[30] Baoguang Shi, Xiang Bai, and Serge Belongie. Detecting Oriented Text in Natural Images by Linking Segments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2550–2558, 2017.

[31] Xinyu Zhou, Cong Yao, He Wen, Yuzhi Wang, Shuchang Zhou, Weiran He, and Jiajun Liang. East: an efficient and accurate scene text detector. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 5551–5560, 2017.

[32] Yuliang Liu, Lianwen Jin, Shuaitao Zhang, and Sheng Zhang. Detecting curve text in the wild: New dataset and new solution. *arXiv preprint arXiv:1712.02170*, 2017.

[33] Yingying Jiang, Xiangyu Zhu, Xiaobing Wang, Shuli Yang, Wei Li, Hua Wang, Pei Fu, and Zhenbo Luo. R2cnn: Rotational region cnn for orientation robust scene text detection. *arXiv preprint arXiv:1706.09579*, 2017.

[34] Dan Deng, Haifeng Liu, Xuelong Li, and Deng Cai. Pixellink: Detecting scene text via instance segmentation. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

[35] Tong He, Zhi Tian, Weilin Huang, Chunhua Shen, Yu Qiao, and Changming Sun. An end-to-end textspotter with explicit alignment and attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5020–5029, 2018.

[36] Xuebo Liu, Ding Liang, Shi Yan, Dagui Chen, Yu Qiao, and Junjie Yan. Fots: Fast oriented text spotting with a unified network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5676–5685, 2018.

[37] Pengyuan Lyu, Minghui Liao, Cong Yao, Wenhao Wu, and Xiang Bai. Mask textspotter: An end-to-end trainable neural network for spotting text with arbitrary shapes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 67–83, 2018.

[38] Lukáš Neumann and Jiří Matas. Real-time scene text localization and recognition. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 3538–3545. IEEE, 2012.

[39] Shijian Lu, Tao Chen, Shangxuan Tian, Joo-Hwee Lim, and Chew-Lim Tan. Scene text extraction based on edges and support vector regression. *International Journal on Document Analysis and Recognition (IJDAR)*, 18(2):125–135, 2015.

[40] Evan Shelhamer, Jonathan Long, and Trevor Darrell. Fully convolutional networks for semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(4):640–651, 2017.

[41] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.

[42] Panqu Wang, Pengfei Chen, Ye Yuan, Ding Liu, Zehua Huang, Xiaodi Hou, and Garrison Cottrell. Understanding convolution for semantic segmentation. *arXiv preprint arXiv:1702.08502*, 2017.

[43] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Deep features for text spotting. In *European conference on computer vision*, pages 512–528. Springer, 2014.

[44] 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.

[45] Youbao Tang and Xiangqian Wu. Scene text detection and segmentation based on cascaded convolution neural networks. *IEEE Transactions on Image Processing*, 26(3):1509–1520, 2017.

[46] Jonathan Fabrizio, Myriam Robert-Seidowsky, Séverine Dubuisson, Stefania Calarasanu, and Raphaël Boissel. Textcatcher: a method to detect curved and challenging text in natural scenes. *International Journal on Document Analysis and Recognition (IJDAR)*, 19(2):99–117, 2016.

[47] Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip HS Torr. Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1529–1537, 2015.

[48] Hyeonwoo Noh, Seunghoon Hong, and Bohyung Han. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1520–1528, 2015.

[49] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3128–3137, 2015.

[50] Andrej Karpathy, Armand Joulin, and Li F Fei-Fei. Deep fragment embeddings for bidirectional image sentence mapping. In *Advances in neural information processing systems*, pages 1889–1897, 2014.

[51] Justin Johnson, Andrej Karpathy, and Li Fei-Fei. Densecap: Fully convolutional localization networks for dense captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4565–4574, 2016.

[52] Rowan Zellers, Mark Yatskar, Sam Thomson, and Yejin Choi. Neural motifs: Scene graph parsing with global context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5831–5840, 2018.

[53] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. *arXiv preprint arXiv:1811.10830*, 2018.

[54] Mark Yatskar, Vicente Ordonez, and Ali Farhadi. Stating the obvious: Extracting visual common sense knowledge. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 193–198, 2016.

[55] Richard Socher, Danqi Chen, Christopher D Manning, and Andrew Ng. Reasoning with neural tensor networks for knowledge base completion. In *Advances in neural information processing systems*, pages 926–934, 2013.

[56] Mohammad Amin Sadeghi and Ali Farhadi. Recognition using visual phrases. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1745–1752. IEEE, 2011.

[57] Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David Shamma, Michael Bernstein, and Li Fei-Fei. Image retrieval using scene graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3668–3678, 2015.

[58] V. Ramanathan et al. Learning semantic relationships for better action retrieval in images. *CVPR*, 2015.

[59] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In *European Conference on Computer Vision*, pages 852–869. Springer, 2016.

[60] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. *CVPR*, 2015.

[61] Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, Zhiheng Huang, and Alan Yuille. Deep captioning with multimodal recurrent neural networks (m-rnn). *arXiv preprint arXiv:1412.6632*, 2014.

[62] Xuejian Rong, Chucai Yi, and Yingli Tian. Unambiguous text localization and retrieval for cluttered scenes. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3279–3287. IEEE, 2017.

[63] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k

entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Computer Vision (ICCV), 2015 IEEE International Conference on*, pages 2641–2649. IEEE, 2015.
[64] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
[65] Chenxi Liu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, and Alan Yuille. Recurrent multimodal interaction for referring image segmentation. *arXiv preprint arXiv:1703.07939*, 2017.
[66] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *ICLR*, 2015.
[67] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.
[68] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
[69] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. SSD: Single Shot Multibox Detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
[70] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. Grabcut: Interactive foreground extraction using iterated graph cuts. In *ACM transactions on graphics (TOG)*, volume 23, pages 309–314. ACM, 2004.
[71] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 787–798, 2014.
[72] Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluis Gomez i Bigorda, Sergi Robles Mestre, Joan Mas, David Fernandez Mota, Jon Almazan Almazan, and Lluis Pere De Las Heras. Icdar 2013 robust reading competition. In *Document Analysis and Recognition (ICDAR), 2013 12th International Conference on*, pages 1484–1493. IEEE, 2013.
[73] Andreas Veit, Tomas Matera, Lukas Neumann, Jiri Matas, and Serge Belongie. COCO-Text: Dataset and Benchmark for Text Detection and Recognition in Natural Images. *arXiv preprint*, 2016.
[74] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. COCO-Stuff: Thing and Stuff Classes in Context. *arXiv preprint arXiv:1612.03716*, 2017.
[75] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *European Conference on Computer Vision*, pages 69–85. Springer, 2016.
[76] T. Lin et al. Microsoft COCO: Common Objects in Context. *ECCV*, 2014.
[77] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. Synthetic data for text localisation in natural images. *CVPR*, 2016.

Xuejian Rong is a researcher at Media Lab of The City College, City University of New York, where he works in the intersection of Deep Learning, Computer Vision, and Computational Photography. He is finishing a PhD advised by Prof. Yingli Tian. His PhD research broadly covers the scene understanding area, specifically on scene text extraction (including scene text detection, retrieval, and recognition), i.e., how to localize and transcribe useful text instances (e.g., characters, words, sentences) widely existing in our surrounding environments. The related methods can be used in many applications, including image retrieval, instant translation, robots navigation, and PhotoOCR. He also occasionally works on image restoration and other inverse problems in computer vision. He received the B.E. degree from Nanjing University of Aeronautics and Astronautics with honors thesis in 2013.

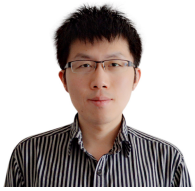

Chucai Yi is currently working at Google corporate on computer vision research and development. Before joining Google, he worked at HERE Map Technologies and Amazon corporates. He received his Ph.D. degree in Computer Science in 2014 from The Graduate Center of City University of New York. As an academic researcher, his research work includes object/signage detection and recognition, human involved event detection, 3D point cloud processing, and camera calibration. As a professional software developer, he has been designing and implementing end-to-end applications and services that apply machine learning and computer vision techniques to solve specific problems on large-scale data.

Yingli Tian (M'99 – SM'01 – F'18) received the B.S. and M.S. degrees from Tianjin University, China, in 1987 and 1990, and the Ph.D. degree from Chinese University of Hong Kong, Hong Kong, in 1996. After holding a faculty position at National Laboratory of Pattern Recognition, Chinese Academy of Sciences, Beijing, she joined Carnegie Mellon University in 1998, where she was a postdoctoral fellow at the Robotics Institute. She then worked as a research staff member in IBM T. J. Watson Research Center from 2001 to 2008. She is one of the inventors of the IBM Smart Surveillance Solutions. She is a professor in the Department of Electrical Engineering at the City College and the department of Computer Science at the Graduate Center, City University of New York since 2008. Her current research focuses on a wide range of computer vision problems from scene understanding, human behavior analysis, facial expression recognition to assistive technology and medical imaging analysis. She is a fellow of IEEE.