

Improving Text Proposals for Scene Images with Fully Convolutional Networks

Dena Bazazian^{*†}, Raúl Gómez^{*‡}, Angelos Nicolaou^{*†}, Lluís Gómez^{*†},
Dimosthenis Karatzas^{*†}, Andrew D. Bagdanov[§]

^{*} Universitat Autònoma de Barcelona (UAB), Barcelona, Spain

[†] Computer Vision Centre (CVC), Barcelona, Spain

[‡] Eurecat - Catalunya Technology Center, unit of Multimedia Technologies, Barcelona, Spain

[§] Media Integration and Communication Center (MICC), University of Florence, Florence, Italy

Email: {dbazazian, angelos, lgomez, dimos}@cvc.uab.es,

raul.gomez@ce.eurecat.org, andrew.bagdanov@unifi.it

Abstract—Text Proposals have emerged as a class-dependent version of object proposals – efficient approaches to reduce the search space of possible text object locations in an image. Combined with strong word classifiers, text proposals currently yield top state of the art results in end-to-end scene text recognition. In this paper we propose an improvement over the original Text Proposals algorithm of [1], combining it with Fully Convolutional Networks to improve the ranking of proposals. Results on the ICDAR RRC and the COCO-text datasets show superior performance over current state-of-the-art.

I. INTRODUCTION

Text understanding in unconstrained scenarios, such as text in scene images or videos, has attracted a lot of attention from the computer vision community. Though extensively studied in recent years, localizing text in complex images remains quite challenging.

Object proposal techniques have emerged as an efficient approach to reducing the search space of possible object locations in an image by generating candidate class-independent object locations and extents. Such generic object proposal methods are typically designed to detect single-body objects, and are not appropriate for text detection which aims to detect groups of disjoint, atomic objects (characters or text strokes). The Text Proposals technique was recently suggested by Gomez et al. [1] as an alternative, class-specific object proposal method that takes into account the particular characteristics of text. Text Proposals combined with a strong word classifier produces state-of-the-art results in end-to-end scene text recognition [1].

In this paper we look at how the power of Fully Convolutional Networks (FCN) can be leveraged to improve Text Proposals. The robustness of FCN stems from the use of convolutional layers instead of fully-connected layers as in conventional networks, which results in spatial information being preserved in the final discriminative layer of the network. FCNs have recently been applied to text detection with notable results [27]. The output of such a network is a coarse heatmap of locations of interest, which must be further processed to produce a proper localization result.

The Text Proposals algorithm, on the other hand, produces a short list of bounding boxes of probable text locations. The method follows a strategy based on an over-segmentation of the image and the efficient construction of meaningful groupings of regions, resulting in very accurate localization. Yet, like all object proposals approaches, it is optimized for high recall and produces a considerable number of false positives.

In this paper we train an FCN for text detection and use resulting heatmap to rerank the set of possible text locations obtained by the Text Proposals algorithm. Different ranking strategies are explored that yield state-of-the-art text proposals results over several datasets.

The main contribution of this paper is the development and evaluation of different ranking strategies for combining the Text Proposals approach with the FCN obtained text probability scores, yielding significantly better recall rates while considering an order of magnitude fewer proposals.

The remainder of the article is organized as follows. Section II presents related work, followed by a description of our approach and architecture in Section III. Section IV reports on the experimental results of our approach, and conclusions are drawn in section V.

II. RELATED WORK

Text detection and recognition in scene images are gaining increasing attention from the computer vision community. In spite of the immense effort which has been devoted to improving the performance of text detection in unconstrained environments, it is still quite challenging due to the diversity of text appearance and geometry in the scene and the highly complicated backgrounds.

Text detection and localization algorithms are divided in two main approaches. First, connected component-based methods perform image segmentation and then classify and group the detected components into text candidates [2], [3], [4], [5], [6]. Following this idea, but in the other way around, is [7] where the interested components are detected by employing

grouping. Second, sliding-window methods perform region classification at multiple scales and follow standard post-processing to derive precise text locations [8], [9], [10], [11], [12]. Most top-ranking approaches [13], [5], [6], [12], [14] in the latest editions of the ICDAR Robust Reading Competition [15] are connected component-based methods based on variants of the Maximally Stable Extremal Regions (MSERs) algorithm [16]. The MSER algorithm is particularly suited to text detection as it efficiently leverages local contrast patterns. Different approaches in this category include the Stroke Width Transform (SWT) [17] and recent variants like the Stroke Feature Transform (SFT) [18]. Connected component-based approaches permit early fusion with the language model at the time of filtering components and forming words. A recent example [19] proposes a model which combines bottom-up cues from individual character detection and top-down cues from a lexicon.

CNN classifiers are increasingly incorporated in text detection approaches. Wang et al. [12] employed CNN models for text/non-text classification. He et al. [20] and Huang et al. [21] exploit a CNN model to filter out non-character MSER components, while CNN based region classification has been repeatedly employed for sliding-window approaches [8], [12]. Tian et al. proposed TextFlow [9] for character classification, which simplifies multiple post-processing steps by utilizing a minimum cost flow network [12]. These approaches all designed features for detecting character candidates, and they apply a pre-trained CNN model to assign a score to each candidate. On a different line, Zhang et al. [11] trained two classifiers that work at character level and text region level, respectively.

Despite the immense success of CNN models for tasks such as character classification or word-spotting once text regions are localized, the problem of text localization still poses significant challenges. To this end, the use of generic object proposal techniques for scene text understanding has been exploited by Jaderberg et al. [22]. Their end-to-end pipeline combines the EdgeBoxes [23] object proposals algorithm and a trained aggregate channel features detector [24] with a powerful deep Convolutional Neural Network for holistic word recognition. Their method uses a CNN-based bounding box regression module on top of region proposals in order to improve their quality before word recognition is performed.

The use of generic object proposals is not optimal when text is to be detected, as demonstrated in [25]. Gomez et al. [1] propose instead a text-specific object proposals method that is based on generating an hierarchy of word hypotheses based on a similarity region grouping algorithm. Contrary to generic object proposal algorithms, this approach takes into account the specific characteristics of text regions that are fundamentally different from the typical notion of single-body object as normally used in for object detection. Replacing the overly complicated pipeline used for text localization in Jaderberg [22] with the single step localization offered by [1] is currently topping the state of the art for end-to-end methods on the demanding Challenge 4 (incidental text) of the Robust

Reading Competition.

On the other hand, Fully Convolutional Networks (FCNs) [26] have recently attracted considerable attention from the robust reading community [27], [28], [29], [30]. FCN-based methods replace fully-connected layers with convolutional layers which allows them to preserve coarse spatial information which is essential for text localization tasks. Zhang et al. [27] integrated semantic labeling by FCN with MSER to provide a natural solution for handling text at arbitrary orientations. In a parallel work Zhang et al. [28] designed a character proposal network based on an FCN which simultaneously predicts characteriness scores and refines the corresponding locations. The characteriness score is used for proposal ranking.

He et al. [29] proposed a Cascaded Convolutional Text Network (CCTN) that combines two customized convolutional networks for coarse-to-fine text localization. The CCTN detects text regions roughly from a low-resolution image, and then accurately localizes text lines from each enlarged region. In addition, to precise text detection, Gupta et al. [30] presented an extreme variant of Hough voting, inspired by Fully-Convolutional Networks [26] and the YOLO technique [31]. The algorithm of YOLO follows the idea of Hough voting, but avoids voting and uses instead individual predictions directly. Therefore, it allows the framework to pick up contextual information which is useful in detecting small text occurrences in natural scene images.

In this paper we take into account the respective advantages of FCN models and Text Proposals in order to propose an improved text proposals algorithm. We employ text probability scores obtained through FCN over individual proposals for ranking, efficiently suppressing non-text proposals. Contrary to Zhang et al. [27] who propose a coarse-to-fine method to obtain bounding boxes from the FCN heatmap with the purpose of maximizing localization accuracy, we employ the FCN heatmap to produce text proposals in order to maximize recall by producing the necessary number of proposals in an efficient way. We choose to use the FCN result to re-rank text proposals obtained by [1], and thereby improving the performance of the text proposals algorithm.

III. THE PROPOSED METHOD

In this section we introduce our text proposal framework. The pipeline, shown in Figure 1, comprises three steps. First, a **text proposal stage** decomposes an image into regions and creates the text group hypotheses based on [1] (see Section III-A). In parallel, a **pixel-wise prediction by FCN** stage estimates the pixel-level text probability for the whole image (see Section III-B). Finally, **Hypotheses Ranking** leverages the FCN predictions to re-rank the text hypotheses (see Section III-C). The remainder of this section discusses these three components in detail.

A. Text Proposals

In our framework we use the Text Proposals algorithm of Gomez et al. [1], which is divided in three main steps: (1)




Fig. 1. Text proposals produce a list of hypotheses for text locations as well as their associated grouping score, while a text detection FCN is simultaneously applied to the image to obtain per pixel text-probability. Text hypotheses are then re-ranked based on the FCN information to obtain the final hypotheses ranking.




Fig. 2. Architecture of the Fully Convolutional Network (FCN) used in our pipeline based on [26].

an initial over-segmentation of the input image from which we obtain a set of connected components; (2) the creation of text hypotheses through several bottom-up agglomeration processes; and (3) a ranking strategy that prioritizes the best text proposals.

Region decomposition, based on the Maximally Stable Extremal Regions (MSER) algorithm, aims to detect the atomic parts that will give rise to subsequent text groupings. The subsequent grouping process builds an hierarchy of groupings of the initial set of MSER regions according to similarity cues such as intensity and color mean and stroke width.

The groupings are first classified as text/non-text based on features efficiently computed in an incremental way along the hierarchy structure. Once non-text groupings have been filtered out, remaining groupings are scored based on a text structure quality score that stems from perceptual organization principles and the notion of perceptual meaningfulness of the particular grouping, and form the resulting text proposals. For further information on the specifics of the Text Proposal algorithm, please see [1].

B. FCN-based Text Prediction

Fully convolutional networks (FCNs) are designed to provide pixel-level predictions [26]. Every layer in an FCN computes a local operation on relative spatial coordinates. Since there is no fully-connected layer, it is possible to use FCN on variable dimension images and produce an output of the corresponding input dimension as well as preserving coarse spatial information of the image which is essential for the text detection tasks. Hence, in this research we use an FCN to perform per-pixel prediction and estimate a *text heatmap* for the input image. The architecture of our network is based on [26] and is shown in Figure 2.

First, we transform the pre-trained VGG network [32] into a fully convolutional form following Long et al. [26]. The original FCN network was designed for the semantic segmentation of images into the twenty classes of the PASCAL VOC dataset. Thus, we customized the network to our purpose



Fig. 3. An example image from our training set, the corresponding ground truth text mask, and an FCN heatmap superimposed on the original image. The heatmap indicates the degree of “textness” estimated by the FCN.

of performing text/none-text segmentation. We apply softmax normalization on the FCN output in order to employ it as text probability for the subsequent hypotheses ranking step. In Figure 3 we show an example image from our training dataset, along with the ground truth text annotations and text heatmap output by our FCN.

C. Hypotheses Ranking

Once text proposals and text heatmaps are computed then we rank the set of proposals according to three different ranking strategies as explained next.

The **baseline** (BAS) strategy ranks regions using the grouping quality score provided by the Text Proposals algorithm [1] as explained in section III-A (does not re-rank text proposals using the FCN heatmap).

The **mean text probability** (MTP) strategy ranks regions using the mean text probability, obtained from the FCN heatmap averaging over the proposed text region. FCN heatmaps provide only coarse information about text localization. Directly using this information to rank bounding boxes can be tricky as smaller regions corresponding to the inner parts of a text block might appear to be better candidates than the full text block bounding box that covers the extents of the text including lower probability pixels at the borders. Indeed, using the average FCN score over text hypotheses as a ranking mechanism has the adverse effect that small patches within text regions are ranked better than full regions corresponding to words or lines. On the other hand, the grouping quality score prioritizes regions with high probability to be words or text lines, but it produces many false positives.

Therefore, we introduce a **suppression** (SUP) strategy, which optimally combines these two complementary behaviors. The idea is to suppress the text proposals which have low mean text probability, in order to filter false positives. Subsequently, the rank of the remaining regions is established according to the grouping quality score which prioritizes well-structured text blocks. The smaller regions corresponding to the inner parts of text blocks are not suppressed by their FCN probability, however, they are generally ranked low based on their grouping quality score as they only cover a small set of atomic parts. The suppression strategy allows us to discard a significant number of false positives, resulting in superior detection with a smaller number of proposals while preserving the high recall rate.

IV. EXPERIMENTAL RESULTS

In this section we describe experiments we performed to evaluate the proposed method. First, we detail the datasets

we used. Afterwards, we quantify the effect of training FCN model over different datasets. Furthermore, we compare between the baseline and proposed ranking strategies, and finally evaluate the effect of different suppression thresholds.

A. Datasets and Experimental Setup

1) *Datasets*: The experiments of our proposed methods were performed on three current text detection datasets covering a range of conditions: ICDAR-Challenge4, SynthText in the Wild, and COCO-Text.

- **ICDAR-Challenge4**: ICDAR Challenge 4 focuses on incidental scene text [15]. Incidental scene text refers to text that appears in the scene without the user having taken any specific prior action to cause its appearance or to improve its positioning/quality in the frame. While focused scene text (explicitly photographed by the user) is the expected input for applications such as translation on demand, incidental scene text represents another wide range of applications linked to wearable cameras or massive urban captures where the acquisition process is difficult or undesirable to control. This challenge has 1000 images with publicly available ground truth annotations. We randomly selected and set aside 200 of these images to use for testing and used the rest for training.
- **SynthText in the Wild**: SynthText in the Wild is a synthetic large scale dataset that provides detailed ground-truth annotations and scalable alternatives to annotating images manually [30]. It comprises 858,750 images, which are divided into 200 different themes. Each theme consists of about 4000 images of text synthetically generated onto almost 40 different backgrounds. For training we considered 15000 random images of the SynthText Dataset with different backgrounds from various themes. Since this dataset is synthetic and it was not use for the validation by the original authors. Therefore, we just consider it for the training stage and not for testing.
- **COCO-Text**: The COCO-Text dataset [33] is based on the MS COCO dataset, which contains images of complex everyday scenes. The COCO-Text dataset contains non-text images, legible text images and illegible text images. A training set and a validation set are defined, from which we consider only those images containing at least one instance of legible text. Notice that some of that images may also contain illegible text instances. That is a total of 22184 images for training and 7026 images for validation. We evaluate on the validation set, since the test set of COCO-Text is withheld.

2) *Training Process*: In this work we trained the FCN network for the task of text detection on the three datasets described in section IV-A1. The training for each of the train sets for these datasets is done by fine-tuning a VGG16 network pre-trained on Imagenet [32].

The FCN networks are all trained using Stochastic Gradient Descent with mini-batches of 20 images, a momentum of 0.9, L2 weight-decay of 5^{-4} , and a fixed learning rate of 10^{-4} . We



Fig. 4. Examples of top proposals by strategy. Top: red regions are those ranked high by Baseline strategy (BAS). Bottom: blue boxes are those ranked high by the FCN.

apply dropout after the Conv6 and Conv7 layers with a rate of 0.5 as shown in Figure 2.

B. Performance Evaluation

The time consumed to generate the FCN heatmaps without taking into account reading the image from disk and writing the heatmap to disk is 0.15 seconds per image. All experiments were conducted on a system with a GPU GeForce GTX TITAN and an Intel® Core™2 Quad CPU Q9300@ 2.50GHz processor.

The quantitative results are shown in Figure 5 and Table I. Each method is coded as [Ranking Strategy]_[FCN Training Set] while for the Suppression Strategy, we also indicate the suppression threshold used.

1) *Effect of Training on different datasets*: We trained the FCN using all three different training sets, and we evaluate each trained model on both ICDAR-Challenge4 and COCO-Text. We report results using the mean text probability (MTP) ranking strategy (MTP_COCO, MTP_ICDAR and MTP_SYNT). Examining the corresponding plots it can be seen that adapting on different training sets has minor influence on the ranking results and the quality of heatmaps produced while there is no observed difference between training on real (COCO, ICDAR) and synthetic (SYNT) data.

At the same time, it can be observed that all curves based on the mean text probability ranking strategy (using only the FCN heatmap information to rank the regions returned by the Text Proposals method) perform notably worse than the baseline.

Looking at the top-N proposals for both the Baseline(BAS) and the MTP strategies in Figure 4 it can be appreciated that by using solely the FCN information, the method selects more relevant regions, however promotes small parts of text blocks. While the Baseline(BAS), making use of the grouping quality score, then returns regions that better match the extent of the text block, but also produces a lot of false positives.

2) *Suppression Strategy*: The suppression strategy has a superior performance over the FCN information by combining it with the grouping quality score.

As the training set produces little difference in the quality of the model, here we report results using the FCN model trained

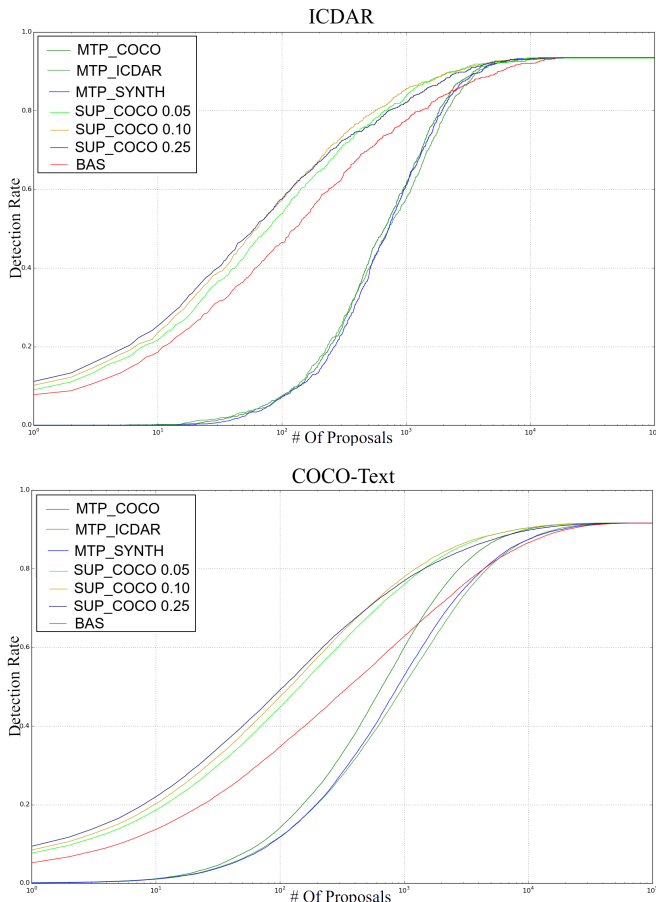


Fig. 5. Evaluating different detection rate of different strategies on ICDAR-Challenge4 (top) and COCO-Text (bottom) datasets.

TABLE I
DETECTION RATE OF DIFFERENT STRATEGIES USING 10, 100 AND 1000 PROPOSALS ON ICDAR-CHALLENGE4 AND COCO-TEXT DATASETS.

# of Proposals	ICDAR			COCO-Text		
	10	100	1000	10	100	1000
SUP_COCO 0.10	0.23	0.57	0.85	0.20	0.47	0.78
SUP_COCO 0.05	0.21	0.53	0.83	0.18	0.44	0.76
SUP_COCO 0.25	0.25	0.57	0.82	0.21	0.49	0.77
BAS	0.18	0.46	0.77	0.13	0.34	0.63
MTP_COCO	0.00	0.07	0.61	0.01	0.15	0.60
MTP_SYNTH	0.00	0.07	0.61	0.01	0.11	0.52
MTP_ICDAR	0.00	0.07	0.57	0.01	0.11	0.50

on the COCO-Text dataset only and focus on the selection of the suppression threshold (using models trained on ICDAR or SYNTH produce very similar results).

Figure 5 shows results employing different thresholds. A threshold of 0.1 yielded the best results in our experiments, since according to Figure 5 and Table I, it demonstrates that by considering smaller threshold the detection rate gets less.

Moreover, in Table I, we present indicative quantitative results for detection rate at 10, 100 and 1000 proposals, which represents a reasonable trade-off for real-life applications.

C. Qualitative Results

Figure 6 illustrates the bounding boxes of the top 500 ranked regions using the baseline and suppression (threshold 0.10) ranking strategies on selected images from the COCO-Text dataset. Similarly, in Figure 6 we give examples from ICDAR-Challenge4. The proposed ranking strategy is capable of detecting varied types of text, including different fonts, handwritten text, various orientations or deformations, diverse text length, and also different languages.

V. CONCLUSIONS

In this paper we proposed an approach to improving the quality of a text-specific object proposals algorithm by exploiting local predictions from a Fully Convolutional Network trained to segment text from non-text regions. Our method re-scores text region proposals by integrating the pixel-wise text probability provided by the FCN model. This allows us to achieve high text region recall while considering significantly fewer candidate regions. Hence, it is capable of accelerating text detection. Experimental results show that our proposed approach achieves a performance superior to the baseline algorithm on COCO-Text and ICDAR2015-Challenge4 benchmarks. In the future, we plan to extend the proposed method to perform as an end-to-end text recognition system.

ACKNOWLEDGMENT

This work was supported in part by the project TIN2014-52072-P and by the Eurecat Catalan Technology Center.

REFERENCES

- [1] Gomez-Bigorda, L and Karatzas, D, TextProposals: a Text-specific Selective Search Algorithm for Word Spotting in the Wild, submitted to Elsevier, arxiv:1604.02619, 2016.
- [2] Kang, L and Li, Y and Doermann, D, Orientation robust text line detection in natural images, IEEE Computer Vision and Pattern Recognition (CVPR), 2014.
- [3] Li, Y and Jia and Shen, C and van den Hengel, A, Characterness: An indicator of text in the wild, IEEE Trans. Image Processing (TIP), 2014.
- [4] Yao, C and Bai, X and Liu, W and Ma, Y and Tu, Z, Detecting texts of arbitrary orientations in natural images, In IEEE Computer Vision and Pattern Recognition (CVPR), 2012.
- [5] Yin, X and Pei, W and Zhang, J and Hao, H, Multi-orientation scene text detection with adaptive clustering, IEEE Trans. Pattern Analysis and Machine Intelligence (TPAMI), 2015.
- [6] Yin, X and Yin, X and Huang, K and Hao, H, Robust text detection in natural scene images, IEEE Trans. Pattern Analysis and Machine Intelligence (TPAMI), 2014.
- [7] Gomez, L and Karatzas, D, Multi-script text extraction from natural scenes, 12th International Conference on Document Analysis and Recognition (ICDAR), 2013.
- [8] Jaderberg, M and Vedaldi, A and Zisserman, A, Deep features for text spotting, In European Conference on Computer Vision, European Conference on Computer Vision (ECCV), 2014.
- [9] Tian, S and Pan, Y and Huang, C and Lu, S and Yu, K and Tan, C.L, Text flow: A unified text detection system in natural scene images, IEEE International Conference on Computer Vision (ICCV), 2015.
- [10] Wang, K and Babenko, B and Belongie, S, End-to-end scene text recognition, IEEE International Conference on Computer Vision (ICCV), 2011.
- [11] Zhang, Z and Shen, W and Yao, C and Bai, X, Symmetry-based text line detection in natural scenes, IEEE Computer Vision and Pattern Recognition (CVPR), 2015.
- [12] Wang, T and Wu, D and Coates, A and Ng, Y, End-to-end text recognition with convolutional neural networks, International Conference on Pattern Recognition (ICPR), 2012.




Fig. 6. ICDAR-Challenge4 and COCO-Text ranking examples. In red, regions ranked by Baseline strategy (BAS), and in green those ranked by the suppression strategy (0.10). The top 500 ranked regions by each strategy are drawn.



Fig. 7. Top ranked regions from COCO-Text and ICDAR-Challenge4 using suppression strategy with a threshold of 0.10.



Fig. 8. Multi-language examples. Top ranked regions from COCO-Text and ICDAR-Challenge4 using suppression strategy with a threshold of 0.10 are shown.

- [13] Koo, H and Kim, D, Scene text detection via connected component clustering and nontext filtering, Image Processing, IEEE Transactions on, vol. 22, no. 6, pp.2296–2305,2013.
- [14] Neumann, L and Matas, J, On combining multiple segmentations in scene text recognition, 12th International Conference on Document Analysis and Recognition (ICDAR), pp. 523–527, 2013.
- [15] Karatzas, D and Gomez-Bigorda, L and Nicolaou, A and Ghosh, Suman and Bagdanov, Andrew and Iwamura, Masakazu and Matas, Jiri and Neumann, Lukas Chandrasekhar, Vijay and Ramaseshan and Lu, Shijian and Shafait, Faisal and Uchida, Seiichi and Valveny, Ernest, ICDAR robust reading competition, 2015.
- [16] Matas, J and Chum, O and Urban, M and Pajdla, T, Robust wide-baseline stereo from maximally stable extremal regions, Image and vision computing (IVC), 2004.
- [17] Epshtein, B and Ofek, E and Wexler, Y, Detecting text in natural scenes with stroke width transform, Computer Vision and Pattern Recognition (CVPR), 2010.
- [18] Huang, W and Lin, Z and Yang, J and Wang, J, Text localization in natural images using stroke feature transform and text covariance descriptors, International Conference on Computer Vision (ICCV), 2013.
- [19] Mishra, A and Alahari, K and Jawahar, C. V. , Enhancing Energy Minimization Framework for Scene Text Recognition with Top-Down Cues, Computer vision and pattern recognition, arxiv:1601.03128, 2016.
- [20] He, T and Huang, W and Qiao, Y and Yao, J, Text-attentional convolutional neural networks for scene text detection, Computer vision and pattern recognition, arxiv:1510.03283, 2015.
- [21] Huang, W and Qiao, Y and Tang, X, Robust scene text detection with convolutional neural networks induced MSER trees, European Conference on Computer Vision (ECCV), 2014.
- [22] Jaderberg, M and Simonyan, K and Vedaldi, A and Zisserman, A, Reading text in the wild with convolutional neural networks, International Journal of Computer Vision, 2016.
- [23] Zitnick, L and Dollar, P, Edge boxes: Locating object proposals from edges, European Conference on Computer Vision (ECCV), 2014.
- [24] Dollar, P and Appel, R and Belongie, S and Perona, P, Fast feature pyramids for object detection, Pattern Analysis and Machine Intelligence, IEEE Transactions on 36 (8)15321545, 2014.
- [25] Gomez, L and Karatzas, D, Object proposals for text extraction in the wild, Document Analysis and Recognition (ICDAR), 13th International Conference on, pages 206–210, 2015.
- [26] Long, J Shelhamer, E and Darrell, T, Fully convolutional networks for semantic segmentation, Computer Vision and Pattern Recognition (CVPR), 2015.
- [27] Zhang, Z and Zhang, Ch and Shen, W and Yao, C and Liu, W and Bai, X, Multi-Oriented Text Detection with Fully Convolutional Networks, Computer vision and pattern recognition, arxiv:1604.04018, 2016.
- [28] Zhang, S and ZLin, M and Chen, T and Jin, L and Lin, L, Character Proposal Network for Robust Text Extraction, Computer vision and pattern recognition, arxiv:1602.04348, 2016.
- [29] He, T and Huang Yu Qiao, W and Yao, J, Accurate Text Localization in Natural Image with Cascaded Convolutional Text Network, Computer vision and pattern recognition, arxiv:1603.09423, 2016.
- [30] Gupta, A and Vedaldi, A and Zisserman, A, Synthetic Data for Text Localisation in Natural Images, Computer vision and pattern recognition, arxiv:1604.06646, 2016.
- [31] Redmon, J and Divvala, S and Girshick, R and Farhadi, A, You Only Look Once: Unified, Real-Time Object Detection, Computer vision and pattern recognition (CVPR), 2015.
- [32] Simonyan, K and Zisserman, A, Very deep convolutional networks for large-scale image recognition, International Conference on Learning Representations (ICLR), 2015.
- [33] Veit, A and Madera, A and Neumann, L and Matas, J and Belongie, S, COCO-Text: Dataset and Benchmark for Text Detection and Recognition in Natural Images, Computer Vision and Pattern Recognition, arxiv:1601.07140, 2016.