

“© 2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

MorphText: Deep Morphology Regularized Accurate Arbitrary-shape Scene Text Detection

Chengpei Xu, Wenjing Jia, Ruomei Wang, Xiaonan Luo, and Xiangjian He, *Senior member, IEEE*

Abstract—Bottom-up text detection methods play an important role in arbitrary-shape scene text detection but there are two restrictions preventing them from achieving their great potential, *i.e.*, 1) the accumulation of false text segment detections, which affects subsequent processing, and 2) the difficulty of building reliable connections between text segments. Targeting these two problems, we propose a novel approach, named “MorphText”, to capture the regularity of texts by embedding deep morphology for arbitrary-shape text detection. Towards this end, two deep morphological modules are designed to regularize text segments and determine the linkage between them. First, a Deep Morphological Opening (DMOP) module is constructed to remove false text segment detections generated in the feature extraction process. Then, a Deep Morphological Closing (DMCL) module is proposed to allow text instances of various shapes to stretch their morphology along their most significant orientation while deriving their connections. Extensive experiments conducted on four challenging benchmark datasets (CTW1500, Total-Text, MSRA-TD500 and ICDAR2017) demonstrate that our proposed MorphText outperforms both top-down and bottom-up state-of-the-art arbitrary-shape scene text detection approaches.

Index Terms—Arbitrary-shape Scene Text Detection, Deep Morphology, Bottom-up Methods, Regularized Text Segments.

I. INTRODUCTION

ARBITRARY-shape scene text detection has gained widespread attention in the age of deep learning. The mainstream approaches for arbitrary-shape scene text detection can be grouped into two types: top-down and bottom-up. The top-down approaches usually model the text areas globally and produce the final text areas in a single pass, while bottom-up methods usually model text areas as a connected-component linkage problem. The top-down methods align with the current design strategy of end-to-end deep learning models, so they often require fewer post-processing steps compared with the bottom-up methods. Moreover, some successful modules designed for object detection and semantic/instance segmentation can be directly embedded into top-down methods to enhance detection performance. However, texts differ significantly from general objects in terms of scale, aspect ratio, etc. Some direct

C. Xu, W. Jia, and X. He are with Faculty of Engineering and IT, University of Technology Sydney, Australia (Email: Chengpei.Xu@student.uts.edu.au, Wenjing.Jia@uts.edu.au, Xiangjian.He@uts.edu.au). X. He will be with the Department of Computer Science, University of Nottingham Ningbo China from May 2022 (Xiangjian.He@gmail.com). Corresponding author: X. He

R. Wang is with School of Data and Computer Science, National Engineering Research Center of Digital Life, Sun Yat-sen University, China (Email: isswrm@mail.sysu.edu.cn).

X. Luo is with the School of Computer Science and Information Security, Guilin University of Electronic Technology, China (Email: luoxn@guet.edu.cn).

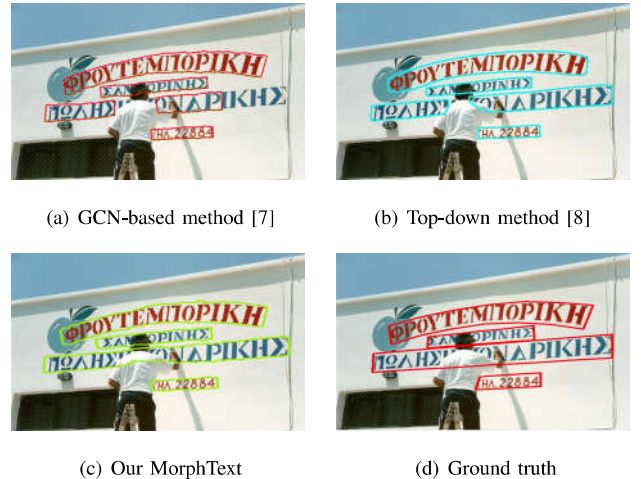


Fig. 1. Both the GCN-based method [7] and the top-down method [8] have failed (as shown in (a) and (b)) when the text instance is separated due to heavy occlusion. With our morphology regularization, such separated text segments can still be connected into a single text instance (as shown in (c)).

embedding operations have not considered the characteristics of the texts themselves. As a consequence, top-down methods sometimes fail to detect long texts due to the limited reception field of CNNs [1], [2]. This has triggered some reception-field expansion methods, such as the non-local [3] and deformable convolution [4] modules in [1], [5], [6]. However, their improvement to arbitrary-shape text detection is limited.

Bottom-up methods are naturally advantageous for processing long or arbitrary-shape texts, since the rules for forming words and sentences in the real world intrinsically follow a bottom-up manner. Bottom-up methods can be more robust in terms of dealing with long texts when the connections between text segments are handled properly. Nevertheless, there are two issues that prevent them from achieving their great potential.

First, the bottom-up methods are more likely to accumulate false positive detections. This is because bottom-up methods usually need to combine feature representations generated from different modalities, *e.g.*, the visual representation of individual text segments as well as the relational representation between text segments. The processes of learning the visual representation of individual text segments and learning the relational representation between different text segments are usually independent or partially independent of each other and are therefore difficult to be optimized simultaneously. For example, some state-of-the-art bottom-up methods [1], [9], [7] have explored embedding Graph Convolutional Neural

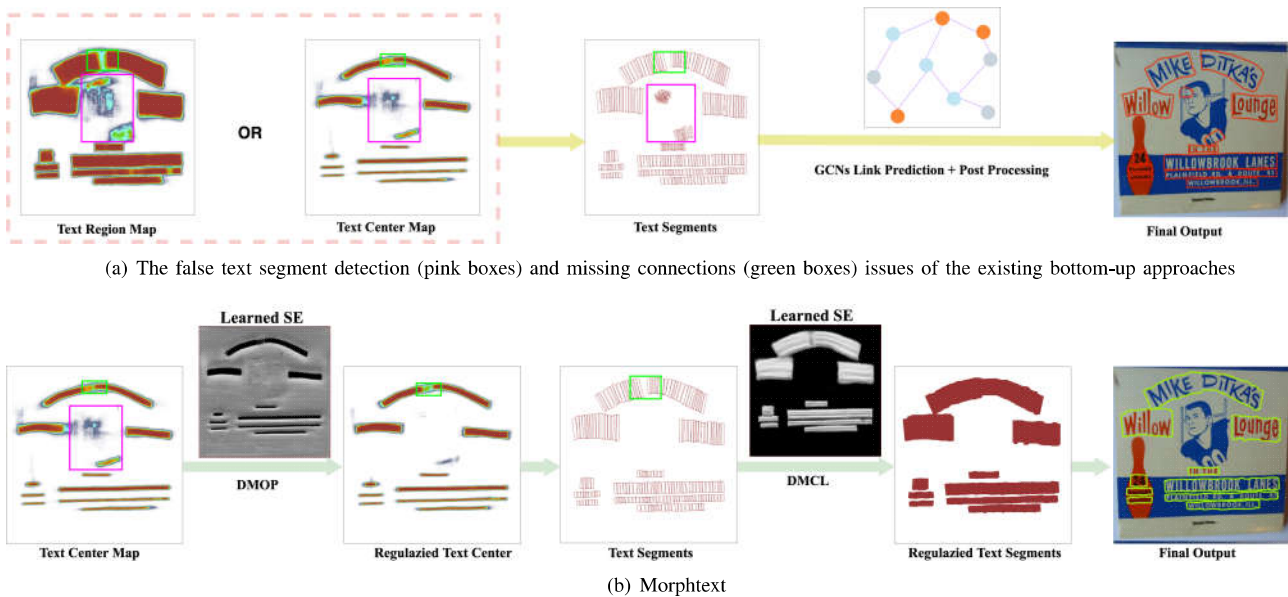


Fig. 2. Our proposed MorphText approach effectively addresses two key issues that restrain the performance of the bottom-up methods. The pink boxes indicate the false detection areas accumulated from the earlier processing and the green boxes indicate the disconnected areas.

Networks (GCNs) for learning the relational information. However, the training of the visual representation and GCN modules are partially independent, so their gradient update may also be partially independent from each other. Also, the cross-modality training sometimes increases the difficulty of training and can be more complicated than the purely CNN-based top-down methods [8], [10], [11], [12], [13]. Moreover, the GCN module only makes inferences on the relationship between text segments and makes no selection or rectification on the text segment candidates, so it does not discriminate whether a text segment is a true positive or false positive detection. Undiscriminated false text segments can lead to error accumulation, which affects the performance of the subsequent processing.

The second constraint of bottom-up methods is the difficulty of reasoning the connections between text segments. GCNs seem to be an ideal solution for relational reasoning compared to the previous sequence-based, rule-based, RNN-based or geometric-location-based methods in [14], [15], [16], [17], [18], since GCNs are able to capture the complex topological structure between text segments. However, GCNs also require large computational resources for building the topological structure of all text segments. GCNs adopted by the bottom-up methods are expected to address the linkage problem of spatially separated text segments, but the performance of GCN link prediction is dependent on the quality of the visual representation of text segments.

Before generating text segments, the network first needs to predict the entire candidate text area so that text segments can be generated by slicing or regression. However, the segmentation of text may not strictly align with the spatial interval of the text instance, unless the dataset provides character-level annotation. Therefore, there are simple cases where the text segments produced by the visual representation stage are

clearly spatially adjacent, as well as difficult cases where text segments are spatially separated.

Thus, on one hand, spatially adjacent text segments (simple cases) are already connected in the original text map. But GCNs unnecessarily compute their connectivity again, adding unnecessary computational burden. On the other hand, for spatially separated text segments (difficult cases), it is also challenging for GCNs to determine whether the separation is due to the accumulation of errors from the previous CNN stage or the text instance is indeed separated at this point (see Fig. 1(a)). More linkage failure cases from GCN-based methods [7] are shown in Fig. 8. This is one of the reasons that the performance of GCN-based bottom-up methods [1], [9], [7] for arbitrary shape text detection is lower than those of top-down methods [8], [10], [19], [6], since GCNs in these methods only perform link prediction while having no ability to prevent error accumulation. Therefore, with additional mechanisms re-checking false text segments, GCNs can be an ideal framework for addressing the linkage problem.

Is there a method that can remove false detections while simultaneously addressing the linkage problem of bottom-up approaches?

Before the dominance of the deep learning techniques, morphological operations had experienced great prosperity in the areas of image processing and text detection. Morphological operations, including erosion, dilation, opening and closing, are designed to analyze the shape and form of an object in the image, so as to filter target patterns and connect distinct patterns in the image. They can be directly applied to gray scale images with a predefined Structure Element (SE). Compared to CNNs, the efficiency of filtering a specific pattern by a morphological operation is higher as a non-linear function, because CNNs need to stack several layers to accomplish the same task, whereas morphological operations

only need to select a proper SE. However, determining the shape, size and flatness of an SE is a process of trial and error and does not always bring performance gains. Handcrafted SEs have significant limitations and cannot adapt to different patterns. This is why traditional morphological operations have gradually lost their attraction for image processing tasks.

Currently, traditional morphological operations have also evolved into Deep Morphological Networks (DMNs) with learnable SEs to capture different patterns in images [20], [21]. These methods have also attempted to prove that DMNs can substitute for CNNs as a universal feature extractor or classifier. However, existing experiments have shown that this is not an area in which morphology performs well, and have often led to incremental improvement only. In our work, we propose utilizing deep morphology in a supplementary role with CNNs to take advantage of their strength in terms of preserving the regular patterns of targets and tackle the linkage problem that linear filters suffer.

Building on the strength of deep morphological operations in preserving patterns of regularity, in this paper we propose a novel arbitrary-shape scene text detection approach to regularize text segments and alleviate the false detection and missing linkage problems of the existing bottom-up methods. Fig. 2 shows such an example of handling the existing problems.

In particular, for removing false detection patterns, we design a Deep Morphological Opening (DMOP) module and embed it into the generation process of text segments and their center-line maps so as to suppress the noisy text-line interfering patterns. Our proposed DMOP is different from the opening operations in [20], [21], as its purpose is to regularize the text segment detection results, rather than being used to extract features. Our proposed embedding of the DMOP module utilizes the core strength of the traditional morphological opening operation combined with trainable kernels to filter those text-like noisy patterns. The figures in Fig. 2 illustrate this. Thus, our DMOP is designed to regularize the results of CNNs under the guidance of our carefully designed loss functions. Moreover, we propose to adopt the residual connection from [22] to ensure our deep morphological opening retains its partial differentiable properties and does not over-process the detection results. Thanks to the trainability of the DMOP module, the error-prone post-processing steps used in existing bottom-up methods are unnecessary.

After applying the DMOP regularization, to link different text segments we design a Deep Morphological Closing (DMCL) module, which operates directly on the morphology of each text segment to fill the gaps between them and build the link. Here, even if the partial text segments are missing due to insufficient feature representation (*e.g.*, Figs. 1(a) and 2(b)), the DMCL module can still have a chance to determine their morphology. The DMCL relaxes the restriction of existing bottom-up methods on text segments having fixed geometric forms, allowing text instances of various shapes to stretch in the most significant direction.

The major contributions of this paper are threefold:

1) Towards accurate detection of arbitrary-shape text, we propose a novel MorphText approach by effectively embedding deep morphology for regularizing text segments.

2) Our method enables the overall network to be trained in an end-to-end manner and replaces the error-prone post-processing steps in bottom-up methods with two trainable deep morphological modules.

3) Our method outperforms the state-of-the-art arbitrary-shape text detection methods on several competitive benchmark datasets. To the best of our knowledge, this is the first time that deep morphology is introduced into this area.

II. RELATED WORK

A. Modeling Methods of Arbitrary-shape Scene Text Detection

The top-down methods typically consider arbitrary-shape scene text as an instance segmentation problem. Therefore, the classic framework Mask-RCNN [23] and its variations have been widely used, playing an important role in top-down methods, *e.g.* ContourNet [10], Xiao *et al.* [11], TextRPN in Wang *et al.* [24], TextFuseNet [8]. There are also some segmentation-based methods [5], [2], [25], which can be considered as top-down methods since, similar to Mask-RCNN-based methods, these methods also generate mask-like text segmentation maps as their final results. In the methods described in [5], [2], [25], the text segmentation map is constructed from a text center line and its offset to the boundary. This step is similar to the direct regression of text boundaries and does not require further post-processing steps. Thus, top-down methods generally obtain text detection results from the segmentation maps of the region proposal networks.

The bottom-up methods often need to connect text segments or generate text areas progressively from their centers to the border areas. These methods include TextSnake [17], CRAFT [15], Scribble Lines [26], etc. Recently, GCNs have been introduced to ensure the linkage between text segments. PuzzleNet [9], RelaText [1] and DRRG [7] are similar works that have utilized GCNs to link text segments. They share the same idea, *i.e.*, first using a text segment proposal network to generate text segments, and then connecting text segments using GCNs. However, their methods suffer from the error accumulation problem: if a text-like object is proposed by the text segment proposal network, the GCNs will indiscriminately link this text-like object with other text segments, resulting in error accumulation. Furthermore, the bottom-up methods [9], [1], [7], [26], [17] often require additional post-processing to determine the visiting order of the contour points. Compared to the top-down methods, they cannot directly generate text regions from the segmentation maps. Although others such as [27], [28] generate text areas from the inside out with a dilation process, they also require rule-based post-processing to refine the final text contour.

In our approach, we propose to embed deep morphology into arbitrary-shape text detection, aiming to retain the regularity of texts so as to regularize false text segment detection and also address the problem of missing linkages between them. Towards this end, two deep morphological modules are designed, both trainable to adapt to different situations. This approach increases the overall integrity and robustness and avoids the tedious post-processing commonly seen in bottom-up methods.

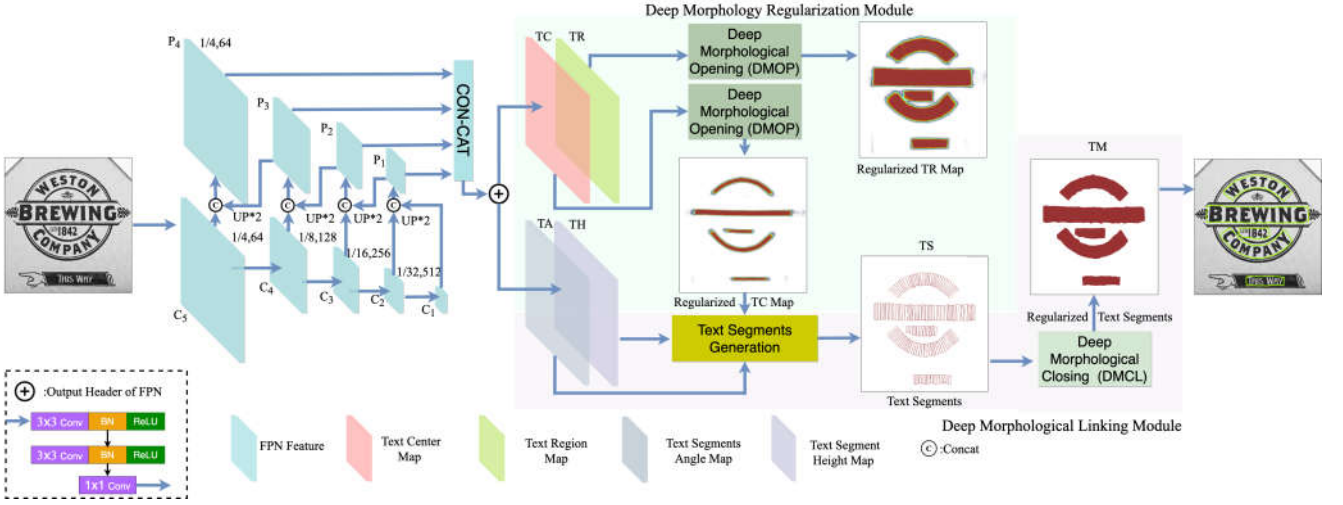


Fig. 3. The overall structure of our network, where “1/4,64”, “1/8,128”,... and “1/32,512” indicate the resize ratio and the channel number.

B. Deep Morphological Networks

The selection of structure elements (SEs) and the enumerative combination of the two basic morphological operations, *i.e.*, erosion and dilation [29] in traditional morphological operations, are the key steps for extracting morphological features and succeeding on the elaborate tasks. Selecting SEs and the combination of morphological operations manually requires expert knowledge and yet the robustness of these handcrafted steps cannot be guaranteed. Thus, it is logical to consider the idea of learning the SEs automatically and this has become feasible during the current deep learning era [30]. Zamora *et al.* [30] trained morphological neurons with a one dimensional SE using a fully connected neural network. Recently, deep morphological methods [20], [31], [32], [21] extended the representation of SEs to two dimensions, and enabled 2D images to be used as input for deep morphological networks. Among them, Nogueira *et al.* [20] used a flat SE to build morphological neurons for extracting features as a replacement for linear convolutions. However, flat SEs only define the shape of the filtering process and do not affect pixel values, so they may lead to insufficient feature representation on images. Later, trainable non-flat SEs were introduced in [32] where a morphology-based pooling method was developed to replace the max pooling process in CNNs. Mondal *et al.* [31], [21] also adopted trainable non-flat SEs and proved that an alternative stacking of deep morphological erosion and dilation could help with image restoration without needing CNNs.

Although the methods in [20], [31], [21] showed that deep morphological networks contain far fewer parameters than CNNs, it is difficult for deep morphological networks to compete with CNNs in terms of feature extraction and classification capabilities. This conclusion can be drawn if we look at the results of image dehazing, image classification and de-raining tasks in [31], [21], [20]. The min/max algebra in the morphological operation can only be piece-wise differentiable during training, resulting in the inability to precisely and progressively update the gradient at each pixel when compared

with the back-propagation process in CNNs. Moreover, the min/max algebra is similar to the max pooling layer in CNNs and it inevitably results in the loss of information.

III. METHODOLOGY

The architecture of our method, illustrated in Fig. 3, contains three parts, namely: a text segment proposal module; a deep morphology-based regularization module; and a morphological linking module. Text images are first fed into the text segment proposal module, which uses ResNet50 [22] as the backbone and FPN [33] as the neck to extract multi-scale features of texts. Then, the multi-scale features are concatenated and processed by the output head. The output head generates a text probability map and its center probability map indicating the center points of text segments, and two probability maps indicating the height and rotation angle of text segments. The DMOP module regularizes the noisy text-like patterns in the text and text center maps with learned SEs and removes possible false text segments. The regularized text center map is then used to guide the fusion of the text segment height and angle maps and produce candidate text segments. Finally, the DMCL module processes the resultant text segments, determining the connections between them based on their morphology and producing the final result.

In the following sub-sections, we describe the two deep morphology modules and then give details of the text segment proposal module and the objective function.

A. Deep Morphological Operations

The deep morphological operations in our approach perform similar functions to the traditional image morphological operations but with non-flat, trainable structure elements [29], [31], [34]. Moreover, the input of our deep morphology modules is multiple-channel image features generated from the text segment proposal module and FPN, rather than gray scale images, as indicated in Fig. 3.

Let us denote the multi-channel image feature of size $C \times W \times H$ by \mathcal{I} , where C , W and H are the channel, width and

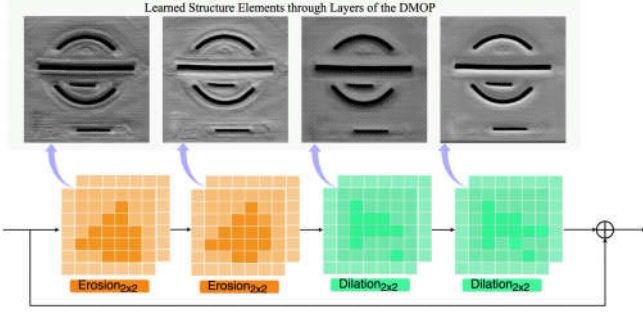


Fig. 4. The visualization of the learned structure elements through DMOP.

height of the input image features, respectively. Similarly, let us denote the size of the structure element S by $C \times M \times N$, where C , M and N are the channel, width and height of the structure element, respectively.

The two basic image morphological operations, *i.e.*, dilation (\oplus) and erosion (\ominus), can then be expressed by the operation of the structure element S on the input I as:

$$(\mathcal{I} \oplus S)[c, w, h] = \max_{(i,j) \in D_s} (\mathcal{I}(c, w+i, h+j) + S(c, i, j)), \quad (1)$$

and

$$(\mathcal{I} \ominus S)[c, w, h] = \min_{(i,j) \in D_s} (\mathcal{I}(c, w+i, h+j) - S(c, i, j)), \quad (2)$$

where, $c \in [1, C]$, $w \in [1, W]$, $h \in [1, H]$, and D_s is the definition field of the structure element S , which can be determined by M and N . In Eqs. (1) and (2), $\mathcal{I}(c, w+i, h+j)$ means that the structure element S slides on the c -th channel of the image feature \mathcal{I} , in a similar manner to the convolution operation in CNNs. In our experiments, the structure elements for dilation and erosion are initialized as zero matrices, which are then updated through optimizing the objective function.

B. Deep Morphology based Text Segment Regularization

False text segments affect the accuracy of both bottom-up and top-down methods. False detection means that text-like interfering patterns are also included in the resultant text segments.

The traditional morphological opening operation [29] is an erosion operation followed by a dilation operation, and it can be used for removing small, irregular objects from detection. The basic principle of our design is also a series of deep erosion operations followed by a series of deep dilation operations. Fig. 4 shows the structure of our Deep Morphological Opening (DMOP) module, which is composed of two erosion blocks followed by two dilation blocks. As can be seen in Fig. 4, the learned SEs gradually suppress the noise patterns with the deep morphological operations. Here, the learned SEs are obtained by splicing the structure elements according to their positions in the image feature.

Considering an input text center map TC , the text center map after applying the DMOP module, denoted by TC' , can be written as:

$$TC' = (TC \ominus S_1^{2 \times 2 \times 2} \ominus S_2^{2 \times 2 \times 2}) \oplus S_3^{2 \times 2 \times 2} \oplus S_4^{2 \times 2 \times 2}, \quad (3)$$

where $S_1^{2 \times 2 \times 2}$ to $S_4^{2 \times 2 \times 2}$ are trainable structure elements with size 2×2 of the same channels as the text center maps. The text region map is processed by DMOP in a similar way.

Note that, a major difference between our DMOP and the deep morphological methods in [31], [21], [20], [32] is that we introduce the residual connection. Since the min/max operation in erosion and dilation are piece-wise differentiable, the gradient descent of the min/max values will be updated in SEs during the training process. For instance, if we consider L as the loss of the structure element S for Eq. (1), the gradient update of S can be represented by $\frac{\partial L}{\partial S}$, which can be piece-wise calculated according to the chain rule:

$$\frac{\partial L}{\partial S} = \sum_{(i,j) \in D_s} \frac{\partial (\mathcal{I} \oplus S[c, w, h])}{\partial S(c, i, j)} \cdot \frac{\partial L}{\partial (\mathcal{I} \oplus S[c, w, h])}. \quad (4)$$

According to the max/min algebra,

$$\frac{\partial (\mathcal{I} \oplus S[c, w, h])}{\partial S(c, i, j)} = 0 \quad \text{or} \quad 1. \quad (5)$$

That is to say that Eq. (5) equates one if and only if $\mathcal{I}(c, w+i, h+j) + S(c, i, j)$ reaches the maximum. Thus, when the expression in Eq. (5) is equal to zero it causes the gradient update of the SE to ignore the non-min/max values and not be constrained. Therefore, when the SE removes true positive text areas and causes over-correction, the back-propagation from non-min/max values should also be preserved to compensate for the gradient update.

In our approach, we propose introducing a residual connection into the deep morphology regularization module, where the regularized text center map, denoted by TC° , after applying DMOP regularization, is computed as:

$$TC^\circ = TC + TC', \quad (6)$$

where TC' is computed by Eq. (3).

Furthermore, different from the methods in [31], [21], [20], [32], which tried to concatenate a sequence of erosion and dilation and their dual sequence and used them as a classifier or feature extractor, our method focuses on using the deep morphological operation as a regularization module. Thus, we select small SEs and shallow deep morphological layers to constrain the influence of the DMOP module onto regularizing false detection only. Section IV-C shows ablation studies validating the effectiveness of the selection of the SE size and the effect of different sizes and ways of combination.

C. Deep Morphology based Relational Reasoning

The center location of text segments TS can now be positioned using TC° . The other geometric features of text segments are described in Sect. III-D. After the text segments are obtained, some of the existing bottom-up methods try to firstly determine the linkage relationship between them and then the visiting order of their contour points. However, if we look at the text segments in Fig. 3, their shapes and contours are already very close to the ground truth text areas. The only problem that we need to solve is if the intervals between text segments should be retained and if the holes in each text segment should be filled up.

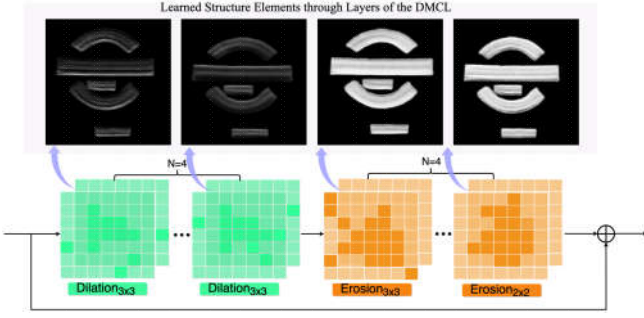


Fig. 5. Visualization of the learned structure elements through DMCL.

The traditional morphological closing operation [29], a dilation operation followed by an erosion operation, has been used to fill small holes and connect some patterns in images. Similar to the DMOP, we propose a Deep Morphological Closing (DMCL) module, which uses four deep morphological dilation operations followed by four erosion operations to build the link between text segments and fill the holes in them (see Fig. 5). Similarly, to avoid the over-correction problem, the residual connection is also adopted in our DMCL module.

Thus, the regularized text segment map, denoted by TS^* , after applying the DMCL module can be represented by:

$$TS^* = TS + TS'', \quad (7)$$

where

$$TS'' = TS' \ominus S_5^{1 \times 3 \times 3} \ominus S_6^{1 \times 3 \times 3} \ominus S_7^{1 \times 3 \times 3} \ominus S_8^{1 \times 3 \times 3}. \quad (8)$$

In this equation,

$$TS' = TS \oplus S_1^{1 \times 3 \times 3} \oplus S_2^{1 \times 3 \times 3} \oplus S_3^{1 \times 3 \times 3} \oplus S_4^{1 \times 3 \times 3}. \quad (9)$$

Here, the size of the SEs is set to 3×3 empirically. The stacking of four deep morphological layers is also determined empirically to fill small holes in the text segment maps. The learned SEs in Fig. 5 show the process of text segments being reshaped and filled to a smooth text instance.

After applying the DMCL module, the text detection results can be directly found in the text segment maps with the hole and connection problem fixed. Thus, our DMCL introduces trainable SEs to build the link and fill the holes, and to replace the tedious post-processing steps, such as those in methods [1], [9], [7], [10]. Our method does not need to locate the contour points or identify their visiting order. The contour drawing process now is the same as in top-down methods. Thus, post-processing steps such as segment grouping and false detection removal are packed into an end-to-end manner.

D. Text Segment Proposal Module

The proposed text segment proposals in our work are rectangles with rotation angles. Each text segment can be represented by $[x, y, h, w, \theta]$, namely, the x and y coordinates of the center points, the height, width and rotation angle of the text segment. Here, we concatenate the FPN features $P1 - P4$ and use an output head (see Fig. 3) to obtain the 6-channel geometric features GF of the text segments, where

$GF \in \mathbb{R}^{6 \times W \times H}$. The GF contains four text probability maps, namely, two channels for the Text Region (TR) map, two channels for the Text Center (TC) map, one channel for the Text Height (TH) map and one channel for Text Angle (TA) map, i.e., $TR \in \mathbb{R}^{2 \times W \times H}$, $TC \in \mathbb{R}^{2 \times W \times H}$, $TH \in \mathbb{R}^{1 \times W \times H}$ and $TA \in \mathbb{R}^{1 \times W \times H}$. Among them, the TR map is not directly involved in the calculation of text segments but is used to guide the training of the FPN feature extractor. The ground truth of TC maps is generated by shrinking the polygon using the clipping algorithm proposed in [35], where the shrinking ratio is set to 0.2.

To generate the ground truth of TH and TA maps, we first use the method in [36] to find the bottom long side of the polygon representation of each text region. Then, the bottom long side is sectioned to n line segments with $n - 1$ points. Here, for every four pixels we sample a point. To generate the TH map, for each point in the TC map, we find the distance from this point to its closest line segment and two times the distance is used to approximate the height of this point. Then, we calculate the angle of the distance line of each point in the TC map with a horizontal line to obtain the TA map.

Note that GF does not produce the feature map for the width of the text segments. This is because our DMOP module has trainable SEs, so that it is not necessary to predict the exact width of the text segments, as each text segment will go through the dilatation operation in the DMOP module. Thus, in our approach, we simply calculate the Text Width (TW) map based on the resultant TH map. Moreover, we design the text segments as slender rectangles, mimicking the general shape of characters. Thus, the TW map is calculated by:

$$TW = Clip(2, \frac{1}{4}TH, 8), \quad (10)$$

where $Clip$ is the clipping function that limits the values in the TW map to the range of 2 – 8 pixels.

Then, with TC, TH, TW and TA, the TS can be generated for the DMOP module. For every point in the TC map, a rectangle text segment is generated where text segments with over 50% overlapping are removed by non-maximum suppression to reduce unnecessary computation.

E. Objective Function

The overall objective function consists of five parts:

$$L = \lambda_1 L_{TR} + \lambda_2 L_{TC} + \lambda_3 L_{TH} + \lambda_4 L_{TA} + \lambda_5 L_{TM}, \quad (11)$$

where L_{TR} and L_{TC} are the losses from the regularized TR and TC maps, and L_{TM} is the loss from text segments after the morphological operation. They are all class-balanced cross-entropy losses for guiding the training of text areas and their center line areas. L_{TM} can be calculated using the same loss function as TR. Online hard example mining [37] is adopted for training the L_{TR} and L_{TM} functions, keeping the ratio of positive sample number to negative sample number to 1 : 3. L_{TC} is also a class-balanced cross-entropy loss with the balancing factor between positive samples and negative samples being set to 0.75 empirically to address the unbalanced-sample problem between the text pixels and background pixels. L_{TH} and L_{TA} are Smoothed-L1 losses for the height and angle of



Fig. 6. Examples of text detection results obtained with the proposed MorphText approach on benchmark datasets.

each text segment. In our paper, λ_1 , λ_2 , λ_3 , λ_4 and λ_5 are set to 1, 2, 1, 1 and 1, respectively, for simplicity.

IV. EXPERIMENTS

A. Datasets and Evaluation Protocol

SynthText: The SynthText dataset [38] consists of 800,000 synthetic images, created by mixing rendered words on natural images. In our work, we follow [39], [15], [5], [7], [1] and use SynthText as a pre-train dataset.

CTW1500: The CTW1500 dataset [40] consists of arbitrary-shape English and Chinese text instances and is a challenging dataset for long curved texts, with 1,500 training and 1,000 testing images.

Total-Text: The Total-Text dataset [41], consisting of 1,255 training images and 300 testing images, was created for detecting arbitrary-shape texts and consists of multi-oriented and curved text instances with polygon annotations.

MSRA-TD500: The MSRA-TD500 dataset [42] is designed for detecting multi-oriented and multi-lingual long texts. It contains 300 training images and 200 testing images with line-level annotation.

ICDAR2017: The ICDAR2017 dataset [43] is designed for detecting multi-oriented and multi-lingual texts of nine languages with 7,200 training images and 1,800 validation images and 9,000 testing images. Following [19], [27], [44], [7], [11], we use ICDAR2017 as part of our pre-train data.

B. Implementation Details

Our MorphText is implemented using the PyTorch 1.7 framework. The backbone network ResNet50 is pre-trained on ImageNet, with FPN being adopted for multi-scale feature extraction. The channels of the up-sampling pipeline are set to 512, 256, 128 and 64 respectively.

The training of MorphText can be divided into two steps. First, we pre-trained our model on SynthText for five epochs using the Adam optimizer with a learning rate of 0.001. The input image size for pre-training was set to 640×640 pixels. Then, before we fine-tuned our model on specific datasets, ICDAR2017 was used as an additional pre-train dataset for

TABLE I
THE EFFECTIVENESS OF OUR PROPOSED DMOP AND DMCL ON CTW1500 AND TOTAL-TEXT.

Datasets	OP	CL	DMOP	DMCL	P (%)	R (%)	F (%)
CTW1500	×	✓	×	×	85.7	82.7	84.2
	✓	✓	×	×	86.5	82.9	84.7
	×	✓	✓	×	87.2	83.0	85.0
	×	×	×	✓	87.6	83.3	85.4
	✓	×	×	✓	88.0	83.2	85.5
	×	×	✓	✓	89.0	83.2	86.0
Total-Text	×	✓	×	×	86.5	84.8	85.6
	✓	✓	×	×	87.1	85.2	86.1
	×	✓	✓	×	87.9	84.8	86.3
	×	×	×	✓	87.4	85.5	86.4
	✓	×	×	✓	87.7	85.5	86.6
	×	×	✓	✓	88.4	85.5	86.9

fine-tuning for another 100 epochs. We adopted the SGD optimizer and an initial learning rate of 0.01, which was decayed by a factor 0.1 for every 100 epochs. The momentum and weight decay of SGD were set to 0.9 and 0.0001, respectively. Following the existing practice, the input size was set to 640×640 for CTW1500 and Total-Text and 960×960 for MSRA-TD500 and ICDAR2017. The numbers of fine-tuning epochs for CTW1500, Total-Text, MSRA-TD500 and ICDAR2017 were set to 200, 200, 250, and 400, respectively. Data augmentation techniques including random cropping, resizing, color variations, adding random noise, flipping and rotation were adopted.

During the testing, following the existing practice, the input images were resized to 640×640 , 800×800 , 1024×1024 and 1024×1024 for CTW1500, Total-Text, MSRA-TD500 and ICDAR2017, respectively. The training was conducted on a single NVIDIA RTX 3090 GPU and the testing was conducted on a single NVIDIA Quadro P6000 GPU with a 3.60GHz Intel Xeon Gold 5122 CPU. Our proposed approach runs at a decent speed, with an average speed of 2.0 fps, 1.8 fps, 1.2 fps and 1.1 fps on the above tested datasets, respectively. Moreover, the processing speed is mainly limited by the non-maximum suppression process, which contributes over 60% of the total processing time irrespective of the input sizes (see the Supplementary for more details).

TABLE II
THE IMPACT OF THE KERNEL SIZE IN DMOP AND DMCL ON F-MEASURE.

Kernel Size	CTW1500				Total-Text			
	Min F (%)	Max F (%)	Mean F (%)	Std dev F	Min F (%)	Max F (%)	Mean F(%)	Std dev F
$2 \times 2, 2 \times 2$	82.4	83.1	82.6	0.24	82.0	82.2	82.1	0.09
$2 \times 2, 3 \times 3$	85.6	86.0	85.9	0.16	86.6	86.9	86.7	0.10
$2 \times 2, 4 \times 4$	79.5	80.2	79.7	0.26	82.0	82.7	82.4	0.26
$2 \times 2, 5 \times 5$	75.6	75.8	75.7	0.08	75.6	76.0	75.7	0.14
$3 \times 3, 2 \times 2$	82.0	82.5	82.3	0.17	82.0	82.3	82.1	0.12
$3 \times 3, 3 \times 3$	85.7	85.9	85.8	0.09	86.0	86.4	86.2	0.16
$3 \times 3, 4 \times 4$	79.5	79.8	79.6	0.12	81.5	81.8	81.6	0.11
$3 \times 3, 5 \times 5$	75.2	75.4	75.3	0.08	75.3	75.6	75.5	0.10
$4 \times 4, 2 \times 2$	81.8	82.2	82.0	0.16	81.2	81.8	81.5	0.20
$4 \times 4, 3 \times 3$	83.1	83.5	83.3	0.15	82.4	82.7	82.5	0.13
$4 \times 4, 4 \times 4$	79.0	79.5	79.3	0.17	81.4	81.7	81.5	0.12
$4 \times 4, 5 \times 5$	75.0	75.4	75.3	0.15	75.3	75.7	75.4	0.14
$5 \times 5, 2 \times 2$	78.2	78.8	78.5	0.20	77.2	77.8	77.6	0.21
$5 \times 5, 3 \times 3$	82.1	82.4	82.3	0.11	82.0	82.4	82.2	0.14
$5 \times 5, 4 \times 4$	77.1	77.7	77.4	0.19	79.7	80.5	80.2	0.27
$5 \times 5, 5 \times 5$	74.1	74.7	74.4	0.23	74.7	75.0	74.9	0.11

TABLE III
THE IMPACT OF THE RESIDUAL CONNECTION ON THE DETECTION.

Residual	CTW1500			Total-Text		
	P (%)	R (%)	F (%)	P (%)	R (%)	F (%)
×	88.2	81.6	84.8	87.2	82.7	84.9
✓	89.0[+0.8]	83.2[+1.6]	86.0[+1.2]	88.4[+1.2]	85.5[+2.8]	86.9[+2.0]

C. Ablation Studies

1) *The Effectiveness of DMOP and DMCL*: To validate the effectiveness of our proposed DMOP and DMCL modules, we conducted ablation studies on the two most representative datasets of arbitrary-shape text detection, *i.e.*, CTW1500 and Total-Text. Table I shows the comparison results. In this table, ‘P’, ‘R’ and ‘F’ represent Precision, Recall, and F-measure, respectively, ‘DMOP’ and ‘DMCL’ denote our DMOP false detection removal branch and DMCL relational reasoning branch, respectively, and ‘OP’ and ‘CL’ denote the traditional morphological opening and closing, where the kernel size of OP and CL are 2×2 and 3×3 , respectively, the same as those in our DMOP and DMCL. Here, the baseline is the model where our proposed DMCL is substituted with the traditional closing operation but without the DMOP noise suppression. Since our method follows a bottom-up design, a closing module, either CL or DMCL, is required to ensure the connectivity of the text segments.

From Table I, the results show that our idea of using deep morphological operations to remove false detection and connect separated text segments is effective. Although the SEs and iteration selection of OP and CL are based on trial and error, the OP and CL adhere to the concept of our bottom-up design and can bring 0.5% and 0.5% gains of F-measure on CTW1500 and Total-Text, respectively. When OP and CL are replaced by our proposed DMOP and DMCL, more significant improvements have been achieved on both datasets: 1.8% and 1.3%, respectively. We attribute this performance gain to the guidance of the loss functions, which allows the DMOP and DMCL to be flexible and robust. To be specific, when CL is replaced by DMCL, the design idea of stretching text segments to fill the gap becomes trainable, and it brings 1.2% and 0.8%

gains on CTW1500 and Total-Text, respectively.

The purpose of OP and DMOP is to remove false detections (spurious text-like objects), which clears the way for the subsequent relational reasoning. The effectiveness of this can be seen from the result of combining OP and DMCL, where even the traditional OP can remove some of the small text-like objects. The selected kernel of OP is a flat all-ones matrix with a size of 2×2 . There is a performance gain of 1.3% and 1.0% on the two datasets. When OP is replaced by DMOP, the combination of DMOP and CL can also bring 0.8% and 0.7% gains on the datasets. The selected SE for CL here is a flat all-ones matrix with a size of 3×3 . The DMOP contains trainable kernels and can decide whether to filter according to different situations of text instances. From the results shown in Table I, we can conclude that our DMOP and DMCL modules are complementary. Through the gradient update process of MorphText, DMOP and DMCL complement each other to boost the final performance.

2) *The Impact of the SE Size*: In traditional morphological operations, the selection of SEs can greatly affect the final results. Here, we mainly test the impact of the size of SEs on the performance, since the values and patterns in the SE are learned with the neural network. In our experiments, the SEs of DMOP and DMCL were initialized with all-zero matrices. We chose the grid search strategy to search the optimal kernel size for DMOP and DMCL. The possible kernel sizes of DMOP were set to $[2 \times 2, 3 \times 3, 4 \times 4, 5 \times 5]$ and the possible kernel sizes of DMCL were set to $[2 \times 2, 3 \times 3, 4 \times 4, 5 \times 5]$. We repeated each combination for five times and report descriptive statistics such as *max*, *min*, *mean* and *std* in Table II. The results show that the SE sizes 2×2 and 3×3 have achieved the highest average F-measure of 85.9% and 86.7% on CTW1500 and Total-Text.

Moreover, if we keep the SE size of DMOP to be 2×2 , the performance decreases as the SE size of DMCL increases from 3×3 . Interpreting the results, we can see that the F-measure is sensitive to the kernel size: for DMCL, when the kernel size is greater than or equal to 4×4 , it may cause very close text instances to stick together and hence results in

TABLE IV
THE EFFECTIVENESS OF OUR PROPOSED DMOP AND DMCL ON A GCN BASED METHOD.

Datasets	OP	GCN	DMOP	DMCL	P (%)	R (%)	F (%)
CTW1500	×	✓	×	×	86.2	83.2	84.7
	✓	✓	×	×	86.9	82.9	84.9
	×	✓	✓	×	87.4	83.2	85.2
	×	×	✓	✓	89.0	83.2	86.0

a fluctuation of F-measure on CTW1500 and Total-Text. For kernel sizes less than 3×3 , sometimes the DMCL fails to deal with large gaps and holes. For DMOP, when the kernel size is equal to or larger than 5×5 , over-correction occurs, also resulting in a decrease in performance. When the kernel size is larger than or equal to 5×5 , it is more likely to use the DMOP as a feature extractor and have more significant impact on rectifying the former CNN features.

The results in [31], [21], [20], [32] show that when deep morphological operations are used as the classifier or feature extractor, the SE size is usually greater than or equal to 5×5 . Instead, in our approach, DMCL and DMOP are used to complement CNNs, and they work well for removing false detections, filling hole areas and smoothing the results of CNNs. A large SE tends to over-correct the detection results and is ineffective when filtering some local false detections.

3) *Comparison with the GCN-based Method:* We also compare the proposed DMOP and DMCL with a GCN-based method [7]. Here, we removed the DMOP module and replaced the DMCL with the GCN implementation in [8] from our framework in Fig. 3. We mentioned earlier that GCNs only perform link prediction and have no ability to deal with errors accumulated from visual representation. The results in Table IV show that GCN-based methods can also be improved by the OP and DMOP operations. Both OP and DMOP can remove some noise patterns but the latter one is guided by the loss. As shown in this table, with the help of DMOP, the GCN-based method can achieve a performance comparable to state-of-the-art (SOTA) with an F-measure of 85.2%. Moreover, the link prediction step of GCN can be fully replaced by DMCL to ensure the connectivity of text segments. Here, DMCL allows text segments to stretch along the most significant orientation of the text instance, which helps to alleviate the problem when text instances are broken in the middle due to the limitations of CNN features and their fixed geometric form.

4) *The Impact of the Residual Connection:* To verify the impact of the residual connection, we keep the kernel size and layers of DMOP and DMCL. Table III shows that when residual connection was removed, all of precision, recall and F-measure were decreased. This supports the claim that residual connection can alleviate the problem of over-correction.

5) *The Impact of the Number of Layers:* To verify the impact of the number of erosion and dilation layers in DMOP and DMCL on the performance, we chose the grid search strategy to search the optimal number of layers of erosion/dilation for DMOP and DMCL. Since the erosion and dilation operations often exhibit symmetrical appearance, the number of erosion and dilation layers in DMOP was selected from [2, 3, 4, 5] and the number of erosion and dilation layers in DMCL

TABLE V
THE IMPACT OF THE NUMBER OF LAYERS OF EROSION AND DILATION IN DMOP AND DMCL ON CTW1500.

Layer l	CTW1500			
DMOP, DMCL	Min F (%)	Max F (%)	Mean F (%)	Std dev F
$l = 2, l = 2$	84.7	85.2	84.9	0.17
$l = 2, l = 3$	85.2	85.5	85.3	0.12
$l = 2, l = 4$	85.6	86.0	85.9	0.16
$l = 2, l = 5$	85.4	85.7	85.5	0.10
$l = 3, l = 2$	84.1	84.9	84.6	0.30
$l = 3, l = 3$	85.1	85.4	85.2	0.11
$l = 3, l = 4$	85.3	85.5	85.4	0.06
$l = 3, l = 5$	85.0	85.4	85.2	0.16
$l = 4, l = 2$	84.0	84.4	84.2	0.16
$l = 4, l = 3$	84.2	84.8	84.6	0.22
$l = 4, l = 4$	84.7	85.0	84.8	0.11
$l = 4, l = 5$	84.6	84.9	84.7	0.12
$l = 5, l = 2$	83.6	84.1	83.8	0.18
$l = 5, l = 3$	83.4	83.9	83.6	0.17
$l = 5, l = 4$	83.8	84.3	83.9	0.17
$l = 5, l = 5$	83.3	83.8	83.6	0.21

TABLE VI
THE EFFECTIVENESS AND ROBUSTNESS OF THE PROPOSED DMOP/DMCL ON REGULARIZING FALSE POSITIVES AND FALSE NEGATIVES ON CTW1500.

DMOP/DMCL	TPs	FPS	FNs	P (%)	R (%)	F (%)
<i>Without</i>	2537	423	531	85.7	82.7	84.2
$2 \times 2, 3 \times 3$	2553 ↑	315 ↓	515 ↓	89.0 ↑	83.2↑	86.0 ↑
$2 \times 2, 4 \times 4$	2429↓	614↑	639↑	79.8↓	79.2↓	79.5↓
$2 \times 2, 5 \times 5$	2313↓	706↑	755↑	76.6↓	75.4↓	75.6↓
$3 \times 3, 3 \times 3$	2571↑	344↓	497↓	88.2↑	83.8 ↑	85.9↑
$4 \times 4, 3 \times 3$	2497↓	420↓	571↑	85.6↓	81.4↓	83.4↓
$5 \times 5, 3 \times 3$	2460↓	444↑	608↑	84.7↓	80.2↓	82.4↓

was selected from [2, 3, 4, 5]. We repeated five times for each combination and reported descriptive statistics such as *max*, *min*, *mean* and *std*. As shown in Table V, with the increase of the erosion and dilation layers in DMOP, there is a decreasing trend in the F-measure of the proposed method on CTW1500. Too many erosion and dilation layers may grossly change the feature representation of CNNs, but in our work, we hope the DMOP module does not over rectify the results of CNNs. For DMCL, four erosion and dilation layers often achieve better F-measure. Fewer than four layers of erosion and dilation may cause the DMCL module fail to handle large gaps.

6) *Discussion on Regularizing False Detections:* We further validate the robustness of our method in addressing false detections. In Figs. 7 and 8, we show some qualitative comparison of the intermediate results, illustrating how MorphText addresses the varied sizes of noise and large linkage problems. The visualization of the results show that the proposed kernel sizes of DMOP and DMCL are robust to different types of noise and relatively large gaps between text segments. Moreover, the various interfering text-like noise patterns, as well as the large holes or gaps between text segments that cause missing connections are the main causes of false detections. The results in Table VI show the effectiveness and robustness of our proposed DMOP and DMCL with kernel sizes of 2×2 and 3×3 for dealing with false positives (FPs) and false negatives (FNs). Since both FPs and FNs have decreased, we conclude that the DMOP/DMCL can reduce the

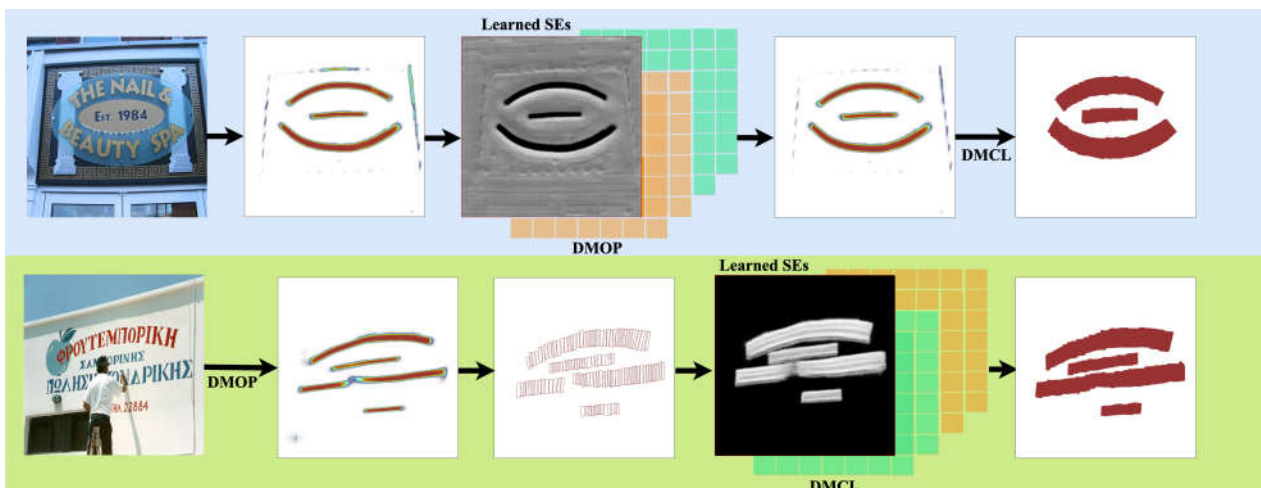


Fig. 7. Visualisation of the intermediate results of MorphText addressing noise patterns (top) and the connection problem between text segments (bottom).



Fig. 8. Qualitative comparisons with the SOTA bottom-up method [7] (top row) on handling varied sizes of interfering patterns and relative large gaps

false positive detections and address the connection problem between text segments. Moreover, an increase in the kernel size of DMOP/DMCL may have a negative impact on dealing with FPs and FNs.

D. Comparison on Benchmark Datasets

In this section, we present the experimental results conducted on the Total-Text, CTW1500, MSRA-TD500 and IC-DAR2017 datasets. Some of the detection results of our proposed method are visualized in Fig. 6. For fair comparison, networks with appropriate volumes such as VGG16 and ResNet50 were selected as the feature extraction backbones to exclude performance gains brought by the improvement of the backbones. Moreover, single-scale testing results are recorded to exclude performance gains brought by multi-scale testing.

Table VII compares the detection results of MorphText with those of the SOTA approaches on the four benchmark datasets, respectively. The “-” in the tables indicates that the comparative method has not reported their results on the dataset. The “EXT” in the tables indicates the comparative method has used external data for pre-training.

Comparison on arbitrary-shape datasets: As shown in Table VII, we first compare our results with the SOTA methods on CTW1500. Our method has achieved **90.0%**, **83.3%** and **86.5%** in precision, recall and F-measure, and outperforms all of the comparison methods. Compared with the pioneering bottom-up methods [1], [7], our model has improved the detection accuracy by 1.7 percent. This result demonstrates that our DMCL module can replace the complex GCN reasoning process to make inferences on the relationships between text segments. When the DMOP module was applied to suppress false detections, the precision of our method significantly improved to 90.0% on CTW1500. Our method remedies the error-prone post-processing steps in the bottom-up methods. The DMOP module introduces a further rectification under the guidance of the loss function, and it provides another way to remove false detections instead of introducing deformable or non-local convolutional blocks, as in the methods shown in [1], [5], [6], [46], [44]. Thanks to our proposed approach embedded with the deep morphological blocks, the results on CTW1500 demonstrate the superb effectiveness of MorphText when dealing with long curved texts.

As shown in Table VII, our method has achieved **90.6%**,

TABLE VII
RESULTS ON CTW1500, TOTAL-TEXT, MSRA-TD500 AND ICDAR2017. (§TOP-DOWN METHODS, †BOTTOM-UP METHODS)

Method	Venue	Backbone	CTW1500				Total-Text				MSRA-TD500				ICDAR2017			
			EXT	P(%)	R(%)	F1(%)	EXT	P(%)	R(%)	F1(%)	EXT	P(%)	R(%)	F1(%)	EXT	P(%)	R(%)	F1(%)
Wang et al.§[45]	CVPR'19	VGG16	×	80.1	80.2	80.1	×	80.9	76.2	78.5	×	85.2	82.1	83.6	-	-	-	-
LOMO§[2]	CVPR'19	ResNet50	✓	85.7	76.5	80.8	✓	88.6	75.7	81.6	-	-	-	-	✓	78.8	60.6	68.5
DB§[5]	AAAI'20	ResNet50	✓	86.9	80.2	83.4	✓	87.1	82.5	84.7	✓	91.5	79.2	84.9	✓	83.1	67.9	74.7
Dai et al.§[46]	TMM'21	ResNet50	×	86.2	80.4	83.2	×	85.4	81.2	83.2	-	-	-	-	×	79.5	66.8	72.6
MS-CAFA §[47]	TMM'21	ResNet50	×	85.7	85.1	85.4	✓	84.6	78.6	81.5	-	-	-	-	-	-	-	-
ContourNet§[10]	CVPR'20	ResNet50	×	85.7	84.0	84.8	×	86.9	83.9	85.4	-	-	-	-	-	-	-	-
TextPerceptron§[25]	AAAI'20	ResNet50	✓	87.5	81.9	84.6	✓	88.8	81.8	85.2	-	-	-	-	-	-	-	-
TextFuseNet§[8]	IJCAI'20	ResNet50	✓	85.0	85.8	85.4	✓	87.5	83.2	85.3	-	-	-	-	-	-	-	-
PCR §[19]	CVPR'21	DLA34 [48]	✓	87.2	82.3	84.7	✓	88.5	82.0	85.2	✓	90.8	83.5	87.0	-	-	-	-
FCENet§[6]	CVPR'21	ResNet50	×	87.6	83.4	85.5	×	89.3	82.5	85.8	-	-	-	-	-	-	-	-
TextSnake†[17]	ECCV'18	VGG16	✓	67.9	85.3	75.6	✓	82.7	74.5	78.4	✓	83.2	73.9	78.3	-	-	-	-
PSENet†[27]	CVPR'19	ResNet50	✓	84.8	79.7	82.2	✓	84.0	78.0	80.9	-	-	-	-	✓	73.7	68.2	70.9
TextRay†[16]	MM'20	ResNet50	×	82.8	80.4	81.6	×	83.5	77.8	80.6	-	-	-	-	-	-	-	-
OPOM†[44]	TMM'20	ResNet50	✓	85.1	80.8	82.9	✓	88.5	82.9	85.6	✓	86.0	83.4	84.7	✓	82.9	70.5	76.2
CRAFT†[15]	CVPR'19	VGG16	✓	86.0	81.1	83.5	✓	87.6	79.9	83.6	✓	88.2	78.2	82.9	✓	80.6	68.2	73.9
TextDragon†[39]	IJCV'19	VGG16	✓	84.5	82.8	83.6	✓	85.6	75.7	80.3	-	-	-	-	-	-	-	-
PuzzleNet†[9]	arXiv'20	ResNet50	✓	84.1	84.7	84.4	✓	-	-	-	✓	88.2	83.5	85.8	-	-	-	-
DRRG†[7]	CVPR'20	VGG16	✓	85.9	83.0	84.4	✓	86.5	84.9	85.7	✓	88.1	82.3	85.1	✓	74.5	61.0	67.3
ReLaText†[11]	PR'21	ResNet50	✓	86.2	83.3	84.8	✓	84.8	83.1	84.0	✓	90.5	83.2	86.7	-	-	-	-
Ours†	-	ResNet50	×	89.0	83.2	86.0	×	88.4	85.5	86.9	×	88.5	82.7	85.5	×	81.9	74.0	77.8
*Ours†	-	ResNet50	✓	90.0	83.3	86.5	✓	90.6	85.2	87.8	✓	90.7	83.5	87.0	✓	82.8	74.2	78.3



Fig. 9. Qualitative comparisons with the SOTA methods on challenging samples.

85.2% and **87.8%** in precision, recall and F-measure on Total-Text, and outperforms the SOTA methods [6] by 2.0 percent in F-measure. The DMCL module can also handle short curved texts in Total-Text flexibly, breaking the text instance when it is appropriate to do so according to the loss.

Fig. 9 qualitatively compares the visual detection results of the proposed MorphText obtained on some samples containing long and curved text instances with results obtained using some recent top-down and bottom-up methods, namely TextFuseNet [8] and DRRG [7]. As shown in this comparison, our method exhibits greater robustness to interfering text-like

noises patterns and is less prone to the connection problem thanks to the proposed DMOP and DMCL modules.

Our results obtained on these two most representative arbitrary-shape text detection datasets demonstrate that the proposed deep morphological modules can handle the two important issues in bottom-up methods in an end-to-end manner, *i.e.*, the unstable post-processing step and the missing connections. Our proposed DMOP module is able to remove the text-like objects from text segments, making DMCL's relational reasoning more reliable.

Comparison on multi-oriented datasets: On the dataset MSRA-TD500, as shown in Table VII, our method has also achieved a comparable SOTA result of an F-measure of 87.0%, compared with the best-performing method PCR [19]. But different from [19], which adopted the latest backbone [48], our MorphText only utilizes the original ResNet50 as the feature extraction backbone. Moreover, the modeling of our MorphText follows the regularity of texts, using the DMCL module to connect character-like text segments instead of introducing complex feature fusion process. Furthermore, different from the methods [9], [1] that used multi-scale training, our method adopts a fixed scale to train the model. The results on MSRA-TD500 prove that the DMCL module can also link non-Latin text segments, allowing MorphText to deal with long multi-oriented texts in an end-to-end manner.

The text detection results on the multi-lingual and multi-oriented dataset ICDAR2017 are shown in Table VII. As it shows, our MorphText has achieved **82.8%**, **74.2%** and **78.3%** in precision, recall and F-measure, so it also surpasses the best performing method [44] by 2.1% in F-measure. Text segments of non-Latin languages such as Chinese and Japanese are different to those of English, where non-Latin languages have no space between words. Therefore, connecting non-Latin text segments requires additional linkage information, since the intervals between non-Latin text segments vary considerably. Methods such as those in [15], [8], [1], [7] attempted to introduce a character-level annotation as well as GCNs to guarantee the connectivity. Different from these methods, our method can directly regularize text segments by the proposed deep morphological modules, so the multi-lingual text segments are less likely to affect the performance. The highest F-measure and recall rate show the robustness of our approach in detecting multi-lingual texts.

V. LIMITATION

Our method can alleviate the error accumulation problem and handle texts with larger word spaces, thanks to the newly proposed DMOP and DMCL modules. Some failure cases happen with low-contrast texts, object-like texts or extremely tiny texts such as the failure cases in the first five pictures in Fig. 10, which are also a challenge for the SOTA methods [8], [19], [6].

For some confusing text segments such as the sixth and seventh pictures in Fig. 10, it is very hard to determine the connectivity of text segments based purely on visual features, which is also challenging for the GCN-based methods [1], [9], [7]. Similar to [46], [47], our methods may also fail when text instances have large text overlapping with other texts, *e.g.*, in the eighth and ninth pictures. Since overlapped texts are very rare in the training datasets, our network may not learn the features of overlapped text.

VI. CONCLUSION

In this paper, we have introduced deep morphology into the field of arbitrary-shape text detection for the first time, as an effective way to tackle the error accumulation of false text segment detection and the missing connection problems that



Fig. 10. Some failure cases of the proposed MorphText, where the green bounding boxes indicate the detected results and the red bounding boxes highlight the failure areas.

prevent bottom-up text detection approaches from achieving their great potential for handling arbitrary-shape text. Two deep morphological modules have been designed and embedded into the network to regularize the text segments based on their patterns of regularity learned through training. Extensive experiments conducted on four widely-used benchmark datasets have demonstrated the effectiveness of our proposed approach. The resulting state-of-the-art performance shows that deep morphology can be used for the challenging task of arbitrary-shape text detection and that the proposed methods have demonstrated an effective approach for the task.

ACKNOWLEDGMENT

This research was supported by the National Key R&D Program of China (2018AAA0100300), UTS FEIT Research Scholarship, and China Natural Science Foundation Grant (ID: 61976037). We thank anonymous reviewers whose comments helped improve and clarify this manuscript.

REFERENCES

- [1] C. Ma, L. Sun, Z. Zhong, and Q. Huo, "Relatext: exploiting visual relationships for arbitrary-shaped scene text detection with graph convolutional networks," *Pattern Recognit.*, vol. 111, p. 107684, 2021.
- [2] C. Zhang, B. Liang, Z. Huang, M. En, J. Han, E. Ding, and X. Ding, "Look more than once: An accurate detector for text of arbitrary shapes," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2019, pp. 10 552–10 561.
- [3] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2018, pp. 7794–7803.
- [4] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei, "Deformable convolutional networks," in *Proc. IEEE Int. Conf. Comput. Vision*, 2017, pp. 764–773.
- [5] M. Liao, Z. Wan, C. Yao, K. Chen, and X. Bai, "Real-time scene text detection with differentiable binarization," in *Proc. AAAI Conf. Artif. Intell.*, vol. 34, no. 07, 2020, pp. 11 474–11 481.
- [6] Y. Zhu, J. Chen, L. Liang, Z. Kuang, L. Jin, and W. Zhang, "Fourier contour embedding for arbitrary-shaped text detection," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2021, pp. 3123–3131.
- [7] S.-X. Zhang, X. Zhu, J.-B. Hou, C. Liu, C. Yang, H. Wang, and X.-C. Yin, "Deep relational reasoning graph network for arbitrary shape text detection," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2020, pp. 9699–9708.

- [8] J. Ye, Z. Chen, J. Liu, and B. Du, "Textfusenet: Scene text detection with richer fused features." *Proc. Int. Joint Conf. Artif. Intell.*, 2020.
- [9] H. Liu, A. Guo, D. Jiang, Y. Hu, and B. Ren, "Puzzlenet: scene text detection by segment context graph learning," *arXiv preprint arXiv:2002.11371*, 2020.
- [10] Y. Wang, H. Xie, Z.-J. Zha, M. Xing, Z. Fu, and Y. Zhang, "Contournet: Taking a further step toward accurate arbitrary-shaped scene text detection," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2020, pp. 11 753–11 762.
- [11] S. Xiao, L. Peng, R. Yan, K. An, G. Yao, and J. Min, "Sequential deformation for accurate scene text detection," in *Proc. Eur. Conf. Comput. Vision*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds., 2020.
- [12] E. Xie, Y. Zang, S. Shao, G. Yu, C. Yao, and G. Li, "Scene text detection with supervised pyramid context network," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, no. 01, 2019, pp. 9038–9045.
- [13] Z. Liu, G. Lin, S. Yang, F. Liu, W. Lin, and W. L. Goh, "Towards robust curve text detection with conditional spatial expansion," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2019, pp. 7269–7278.
- [14] B. Shi, X. Bai, and S. Belongie, "Detecting oriented text in natural images by linking segments," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2017, pp. 2550–2558.
- [15] Y. Baek, B. Lee, D. Han, S. Yun, and H. Lee, "Character region awareness for text detection," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2019, pp. 9365–9374.
- [16] F. Wang, Y. Chen, F. Wu, and X. Li, "Textray: Contour-based geometric modeling for arbitrary-shaped scene text detection," in *Proc. ACM Int. Conf. Multimedia*, 2020, pp. 111–119.
- [17] S. Long, J. Ruan, W. Zhang, X. He, W. Wu, and C. Yao, "Textsnake: A flexible representation for detecting text of arbitrary shapes," in *Proc. Eur. Conf. Comput. Vision*, 2018, pp. 20–36.
- [18] Z. Tian, M. Shu, P. Lyu, R. Li, C. Zhou, X. Shen, and J. Jia, "Learning shape-aware embedding for scene text detection," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2019, pp. 4234–4243.
- [19] P. Dai, S. Zhang, H. Zhang, and X. Cao, "Progressive contour regression for arbitrary-shape scene text detection," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2021, pp. 7393–7402.
- [20] K. Nogueira *et al.*, "An introduction to deep morphological networks," *arXiv e-prints*, pp. arXiv–1906, 2019.
- [21] R. Mondal, S. S. Mukherjee, S. Santra, and B. Chanda, "Morphological network: How far can we go with morphological neurons?" *arXiv preprint arXiv:1901.00109*, 2019.
- [22] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2016, pp. 770–778.
- [23] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Proc. IEEE Int. Conf. Comput. Vision*, 2017, pp. 2961–2969.
- [24] X. Wang, Y. Jiang, Z. Luo, C.-L. Liu, H. Choi, and S. Kim, "Arbitrary shape scene text detection with adaptive text region representation," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2019, pp. 6449–6458.
- [25] L. Qiao, S. Tang, Z. Cheng, Y. Xu, Y. Niu, S. Pu, and F. Wu, "Text perception: Towards end-to-end arbitrary-shaped text spotting," in *Proc. AAAI Conf. Artif. Intell.*, vol. 34, no. 07, 2020, pp. 11 899–11 907.
- [26] W. Zhang, Y. Qiu, M. Liao, R. Zhang, X. Wei, and X. Bai, "Scene text detection with scribble line," in *Proc. IEEE Int. Conf. on Document Anal. and Recognit.* Springer, 2021, pp. 79–94.
- [27] W. Wang, E. Xie, X. Li, W. Hou, T. Lu, G. Yu, and S. Shao, "Shape robust text detection with progressive scale expansion network," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2019, pp. 9336–9345.
- [28] Y. Xu, Y. Wang, W. Zhou, Y. Wang, Z. Yang, and X. Bai, "Textfield: Learning a deep direction field for irregular scene text detection," *IEEE Trans. Image Processing*, vol. 28, no. 11, pp. 5566–5579, 2019.
- [29] J. Serra, "Image analysis and mathematical morphology," 1982.
- [30] E. Zamora and H. Sossa, "Dendrite morphological neurons trained by stochastic gradient descent," *Neurocomputing*, vol. 260, pp. 420–431, 2017.
- [31] R. Mondal, P. Purkait, S. Santra, and B. Chanda, "Morphological networks for image de-raining," in *International Conference on Discrete Geometry for Computer Imagery*. Springer, 2019, pp. 262–275.
- [32] G. Franchi, A. Fehri, and A. Yao, "Deep morphological networks," *Pattern Recognit.*, vol. 102, p. 107246, 2020.
- [33] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2017, pp. 2117–2125.
- [34] R. Mondal, M. S. Dey, and B. Chanda, "Image restoration by learning morphological opening-closing network," *Mathematical Morphology-Theory and Applications*, vol. 4, no. 1, pp. 87–107, 2020.
- [35] B. R. Vatti, "A generic solution to polygon clipping," *Communications of the ACM*, vol. 35, no. 7, pp. 56–63, 1992.
- [36] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *arXiv preprint arXiv:1609.02907*, 2016.
- [37] A. Shrivastava, A. Gupta, and R. Girshick, "Training region-based object detectors with online hard example mining," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2016, pp. 761–769.
- [38] A. Gupta, A. Vedaldi, and A. Zisserman, "Synthetic data for text localisation in natural images," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2016, pp. 2315–2324.
- [39] W. Feng, W. He, F. Yin, X.-Y. Zhang, and C.-L. Liu, "Textdragon: An end-to-end framework for arbitrary shaped text spotting," in *Proc. IEEE Int. Conf. Comput. Vision*, 2019, pp. 9076–9085.
- [40] Y. Liu, L. Jin, S. Zhang, C. Luo, and S. Zhang, "Curved scene text detection via transverse and longitudinal sequence connection," *Pattern Recognit.*, vol. 90, pp. 337–345, 2019.
- [41] C. K. Ch'ng and C. S. Chan, "Total-text: A comprehensive dataset for scene text detection and recognition," in *Proc. IEEE Int. Conf. on Document Anal. and Recognit.*, vol. 1. IEEE, 2017, pp. 935–942.
- [42] C. Yao, X. Bai, W. Liu, Y. Ma, and Z. Tu, "Detecting texts of arbitrary orientations in natural images," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.* IEEE, 2012, pp. 1083–1090.
- [43] N. Nayef, F. Yin, I. Bizid, H. Choi, Y. Feng, D. Karatzas, Z. Luo, U. Pal, C. Rigaud, J. Chazalon *et al.*, "Icdar2017 robust reading challenge on multi-lingual scene text detection and script identification-rrc-mlt," in *Proc. IEEE Int. Conf. on Document Anal. and Recognit.*, vol. 1. IEEE, 2017, pp. 1454–1459.
- [44] S. Zhang, Y. Liu, L. Jin, Z. Wei, and C. Shen, "Opmp: An omnidirectional pyramid mask proposal network for arbitrary-shape scene text detection," *IEEE Trans. Multimedia*, vol. 23, pp. 454–467, 2020.
- [45] X. Wang, Y. Jiang, Z. Luo, C.-L. Liu, H. Choi, and S. Kim, "Arbitrary shape scene text detection with adaptive text region representation," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2019, pp. 6449–6458.
- [46] P. Dai, Y. Li, H. Zhang, J. Li, and X. Cao, "Accurate scene text detection via scale-aware data augmentation and shape similarity constraint," *IEEE Trans. Multimedia*, 2021.
- [47] P. Dai, H. Zhang, and X. Cao, "Deep multi-scale context aware feature aggregation for curved scene text detection," *IEEE Trans. Multimedia*, vol. 22, no. 8, pp. 1969–1984, 2019.
- [48] F. Yu, D. Wang, E. Shelhamer, and T. Darrell, "Deep layer aggregation," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 2018, pp. 2403–2412.

Chengpei Xu is currently a Ph.D. candidate at the Faculty of Engineering and Information Technology (FEIT), University of Technology Sydney (UTS). He received his Master's degree from the University of New South Wales and his Bachelor's degree from the National University of Defense Technology. His research has focused on scene text detection and recognition, document layout analysis, rotating object detection, and multimodal data fusion.

Wenjing Jia received her Ph.D. degree in Computing Sciences from UTS in 2007. She is currently a Senior Lecturer at FEIT, UTS. Her research falls into the fields of image processing and analysis, computer vision and pattern recognition.

Ruomei Wang received the Ph.D. degree from The Hong Kong Polytechnic University, Hong Kong, in 2007. Since 1996, she has been with Sun Yat-sen University, where she is currently a Professor with the School of Data and Computer Science. Her research interests include computer graphics, computer aided design, and image processing.

Xiaonan Luo is currently a Professor with the School of Computer Science and Information Security, Guilin University of Electronic Technology, Guilin, China. He received the National Science Fund for Distinguished Young Scholars granted by the National Nature Science Foundation of China. His research interests include computer graphics, CAD, image processing, and mobile computing.

Xiangjian He received his Ph.D. degree in University of Technology Sydney, Australia, in 1999. He will soon be with the University of Nottingham Ningbo China and be the Director of Computer Vision and Intelligent Perception Lab. He is currently the Director of Computer Vision and Pattern Recognition Laboratory, at the Global Big Data Technologies Centre of the University of Technology Sydney, Australia. His research interests include image processing, network security, pattern recognition, computer vision and machine learning.