

# Diving into the Depths of Spotting Text in Multi-Domain Noisy Scenes

Alloy Das<sup>§1</sup>, Sanket Biswas<sup>§2</sup>, Umapada Pal<sup>1</sup> and Josep Lladós<sup>2</sup>

**Abstract**—When used in a real-world noisy environment, the capacity to generalize to multiple domains is essential for any autonomous scene text spotting system. However, existing state-of-the-art methods employ pretraining and fine-tuning strategies on natural scene datasets, which do not exploit the feature interaction across other complex domains. In this work, we explore and investigate the problem of *domain-agnostic scene text spotting*, i.e., training a model on multi-domain source data such that it can directly generalize to target domains rather than being specialized for a specific domain or scenario. In this regard, we present the community a text spotting validation benchmark called *Under-Water Text (UWT)* for noisy underwater scenes to establish an important case study. Moreover, we also design an efficient super-resolution based end-to-end transformer baseline called *DA-TextSpotter* which achieves comparable or superior performance over existing text spotting architectures for both regular and arbitrary-shaped scene text spotting benchmarks in terms of both accuracy and model efficiency. The dataset, code and pre-trained models have been released in our [Github](#).

## I. INTRODUCTION

Contextual cues from scenes are often used by humans for recognizing text under degraded scene conditions suffering from low-resolution, blurs, distortions, and occluded objects. Despite its success, current scene text recognition (STR) benchmarks [22], [5], [32] in the literature fail to justify the robustness and effectiveness of state-of-the-art (SOTA) text spotting models under real-world noisy domains, especially underwater scenes. Capturing clear underwater scene images is practically unfeasible, mostly due to the effect caused by colour scatter in addition to colour cast from varying light attenuation in different wavelengths [1], [25]. The lack of a public benchmark led us to develop a real-world *Under-Water Text (UWT)* spotting validation benchmark to enhance the progress of end-to-end text spotting systems for multi-domain noisy and complex scenes.

Given two visual domains of a completely different nature, it often leads to reason that each would require a different architecture. However, multi-domain learning (MDL) was introduced in [41] to utilize the same computational pipeline to learn incrementally a set of parameters for every added domain, while retaining the performance of already learned datasets [13]. In this work, we explore a new domain-agnostic framework to investigate the ability of a scene text spotting model to generalize to a wider range

<sup>1</sup>Alloy Das and Umapada Pal are with Computer Vision and Pattern Recognition Unit, Indian Statistical Institute, Kolkata, India [alloyuit@gmail.com](mailto:alloyuit@gmail.com), [umapada@isical.ac.in](mailto:umapada@isical.ac.in)

<sup>2</sup>Sanket Biswas and Josep Lladós are with the Computer Vision Center, Computer Science Department, Universitat Autònoma de Barcelona, Barcelona, Spain [sbiswas](mailto:sbiswas@cvc.uab.es), [josep](mailto:josep@cvc.uab.es)

<sup>§</sup>Equal contribution



Fig. 1. **Qualitative comparison with SOTA approaches:** The above figure shows a sample test image from the TotalText dataset and the end-to-end text spotting performance of DA-TextSpotter compared with others.

of domain complexities rather than being specialized for a specific domain. The key intuition behind the utility of MDL to DA-TextSpotter is to improve the model robustness against different variations of noises as demonstrated in [46], [47]. To demonstrate this important property, we propose a transformer-based scene-text spotting pipeline called *Domain Agnostic-TextSpotter (DA-TextSpotter)* which can also generalize to new domains and scenarios (in this case, from natural scene domain to underwater) without requiring any re-training or fine-tuning.

Intelligent reading systems applied to underwater scenes get affected by reduced visibility and contrast, due to wavelength-dependent absorption and scattering caused by variations of refractive indices associated with temperature, optical turbulence, and salinity microstructures present in lakes and oceans [17]. To mitigate this issue and to significantly improve the capabilities of underwater robotic systems [57], DA-TextSpotter introduces an enhancer unit as a pre-processing step which utilizes the Enhanced-Super-Resolution Generative Adversarial Networks (ESR-GAN) [54], [56] to enhance and improve the visual quality of the input images coming from either natural scenes or

underwater. A Swin-Transformer [35] backbone has been adapted to the proposed framework for visual feature extraction. A significant performance boost is obtained compared to ResNets [15] or simple Vision Transformers (ViTs) [9] owing to the better fine-grained-prediction of text instances with its shifted-window-based attention strategy. The rest of the pipeline follows a single encoder-dual decoder transformer unit, following a similar scheme to TESTR [59]. A further in-depth study with a domain-generalization [37] setting was also conducted to observe the performance of DA-TextSpotter compared to leading SOTA text spotting approaches. With an overall assessment, DA-TextSpotter outperforms its nearest competitors [18], [10], [59] in both text detection and end-to-end recognition tasks for both natural scene and underwater domains.

The key contributions of the paper can be summarized in three-fold: (1) A novel scene text-spotting framework which is domain-agnostic and achieves superior performance on both natural scene and underwater image benchmarks. (2) A simple, efficient and flexible text-spotting pipeline, which introduces a super-resolution-based enhancement unit for dealing with more adverse noisy images, along with a visual feature extraction branch that introduces both local and global attention over patches. (3) A new public evaluation benchmark for text spotting in underwater scenes which helps to justify the essential findings of this work, along with future scopes for the text recognition community.

## II. RELATED WORK

**Multi-Domain Learning.** Multi-Domain Learning (MDL) was analysed in Joshi *et. al.* [21] to probe some important research questions on taking advantage of multiple domain labels to improve learning. Rebuffi *et. al.* [41] then coined it as the task of learning a single visual representation from multiple label spaces of different visual domains, without any forgetting [27]. Rosenfeld *et. al.* [45] further extended this idea of multi-domain towards a multi-task setting, where different tasks are to be performed on the same domain. Liu *et. al.* [34] trained a network to align and match the feature distributions from multiple spaces for image classification and more recent studies [48], [23] in document classification. Motivated by the aforementioned works, we use a simple concept of MDL towards scene text spotting across multi-domain complex scenes to gain some model robustness by inherently learning domain-agnostic parameters from diverse noisy domains [46].

**Image Super-resolution.** The image super-resolution field has evolved since the inception of SRCNN [8] used by super-resolution GAN [24] models to upsample and enhance low-resolution images and restore them. The ESR-GAN [56], [54] has been one of the landmark super-resolution-based approaches that use a relativistic GAN discriminator [44] and a real-world fully synthetic dataset to train their model. This work uses the power of Real-ESRGAN's [54] to be used as an enhancement unit in the proposed DA-TextSpotter framework. It becomes really

important, especially for text spotting in our proposed UWT benchmark. Previously, Banerjee *et. al.* [3] proposed a text detection strategy for underwater scene images, but not applied to multi-domain end-to-end text spotting.

**End-to-End Text Spotting.** Scene text spotting is considered as joint modelling of text detection and recognition modules, especially challenging in autonomous driving systems [42], [12], [4]. Previous approaches [50], [29] addressed the training process in a two-stage manner (detection and recognition) separately and later joined them for inference. The first attempt to joint detection and recognition simultaneously was done through RoI operations [26], [16], [30] during training, although they had sub-par performance in detecting arbitrary-shaped text. Segmentation-based techniques [36], [28] based on Mask-RCNN [14] became more popular later and achieved decent performance on arbitrary-shaped text instances. However, they were computationally more expensive and PAN++ [52] improved its efficiency by using a faster detector to multiply the segmentation maps with the features to suppress the background as followed in other approaches [53], [40]. To avoid further post-processing of the mask representation, MANGO [38] proposed a mask attention module to use more global features and give it more robustness. Other representations for curved texts are parametric Bézier curves used by ABCNet [31], [33] or coordinates of bounding-box guided polygon vertices as in TESTR [59]. In this work, we follow a polygonal bounding-box representation for arbitrary-shaped text.

Modelling linguistic knowledge for end-to-end STR was first proposed in [11] which was later extended to an end-to-end scene text spotting framework in [10] by using bidirectional and iterative language modelling to spot text regions. Transformers [49] became popular in recent times [19], [58] in STR tasks due to its global modelling capabilities and better parallelization ability. Later, ViTs [9] has been used to perform faster and more efficient STR [2]. Recently, self-supervised ViT pretraining [47] to learn different kinds of scene degradations has been employed to achieve the best STR performance. For scene text spotting, however, the relationships between different text instances are critical since they might share common background textures and text styles. Current state-of-the-art methods like SwinTextSpotter [18] and TESTR [59] propose a joint optimization of detection and recognition with different strategies. While SwinTextSpotter uses a Swin transformer [18] backbone coupled with Feature Pyramid Networks (FPNs) to spot smaller text instances, TESTR [59] relies on multi-head deformable attention [60] to achieve a similar objective. In this work, we strive towards efficiency by using a tiny variant of the Swin feature backbone coupled with a deformable attention module as in [6] for the text spotting unit.

## III. METHODOLOGY

Our method proposes a domain-agnostic and efficient framework called *DA-TextSpotter* for scene text spotting that mainly addresses unified text detection and recognition

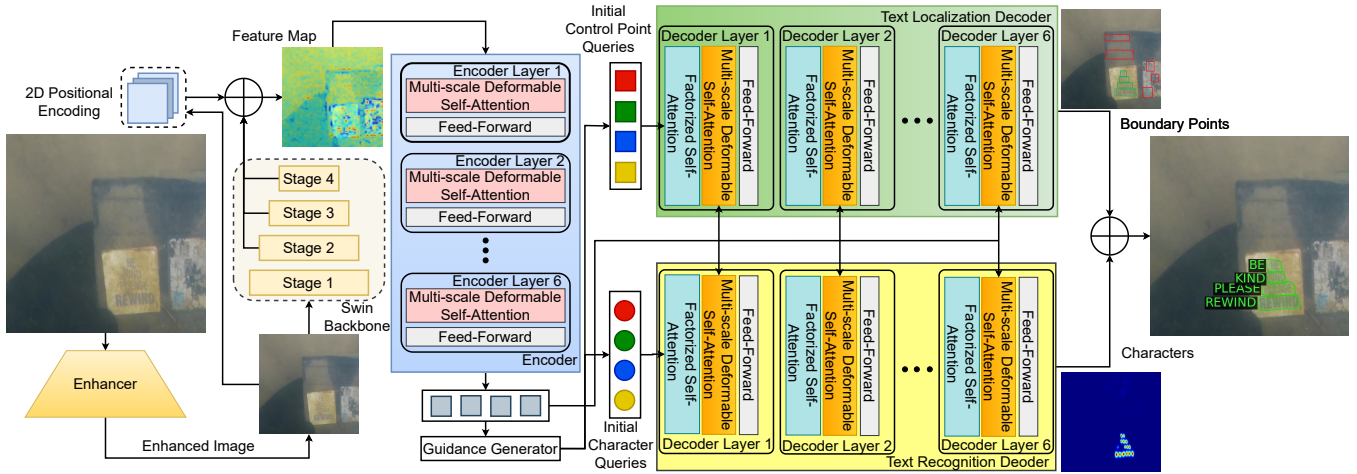


Fig. 2. The overall framework of DA-TextSpotter consisting of Super-Resolution, Feature Extraction and Text Spotting units sequentially.

in noisy domains such as underwater scenes having low-resolution, high contrast, occluded objects and so on. The primary goal of our proposed network is to learn some domain-agnostic feature representation by pre-training on both natural and underwater scene source data. We hereby introduce the architectural pipeline of our method, describe the training strategies, discuss model efficiency, and finally discuss some key insights.

### A. Model Architecture

The overall architectural pipeline as illustrated in Fig.2 consists of three major components: (1) a super-resolution unit that helps to pre-process and enhance the input scene images; (2) a feature extraction unit based on a Swin-Transformer [35] backbone to collect scene-context aware information; (3) a text spotting unit which unifies both detection and recognition modules for the desired output.



Fig. 3. Illustration of the effectiveness of the Super-Resolution unit for noise removal on underwater (First case) and natural (second case) scenes.

**Super-resolution Unit.** The super-resolution unit mainly consists of a pre-trained ESR-GAN [54] backbone which has been used for denoising and upsampling the low-resolution and noisy input scene images. The backbone primarily consists of a Residual in Residual Dense Block (RRDB) which combines a multi-level residual network and dense connection without Batch Normalization. In the discriminator, this network uses the relativistic GAN [20] which helps it to generate more realistic images. Using contextual and perceptual losses, the neural network is pushed to the

natural or underwater image manifold using a discriminator network trained to differentiate between super-resolved images and original photo-realistic images, whereas adversarial loss pushes the neural network to the natural image manifold. The above details can be quantified in terms of eq. 1 where  $L_G$  denotes the overall loss computed in the generator while  $L_1$  and  $L_{per}$  denote the contextual and perceptual losses.

$$L_G = L_{per} + \lambda L_G^{Ra} + \eta L_1, \quad (1)$$

**Feature Extraction Unit.** It is hard to connect remote features with vanilla convolutions since they operate locally at fixed sizes (e.g.,  $3 \times 3$ ). Text spotting requires modelling the relationship between different texts since scene texts from the same image have strong similarities, such as backgrounds, textures and text styles. For our backbone, we used a small and efficient variant of Swin-Transformer [35] denoted as Swin-tiny. Considering the blanks between words in a line of text, the receptive field should be large enough to help distinguish whether adjacent texts belong to the same line. In Fig.4 we have illustrated how our Swin-tiny backbone generates a better-localized representation over the text than the standard Resnet-50 model. The final output from the encoder’s last layer is then propagated to the next phase.

**Text Spotting Unit.** The text-spotting unit is mainly composed of two primary model components: (1) Text Localization Decoder and (2) Text Recognition Decoder. It follows the transformer decoder pipeline from [59]. Accordingly, we formulate our problem as a set prediction problem, to predict a set consisting of a set of point-character tuples, for a particular image. We formulate it as  $X = \{(S^{(i)}, R^{(i)})\}_{i=1}^K$ . Where  $i$  is the index of each instances,  $S^{(i)} = (s_1^{(i)}, \dots, s_N^{(i)})$  is the coordinates of  $N$  control points, and  $R^{(i)} = (r_1^{(i)}, \dots, r_N^{(i)})$  is the  $M$  characters of the text. The location decoder will detect (predict  $S^{(i)}$ ) while the character decoder will recognise (predict  $R^{(i)}$ ) the text in a unified manner.

**Text Localization Decoder.** For the location decoder, we convert the queries to composite queries which predict multiple control points for a single instance. We have such  $Q$  queries,

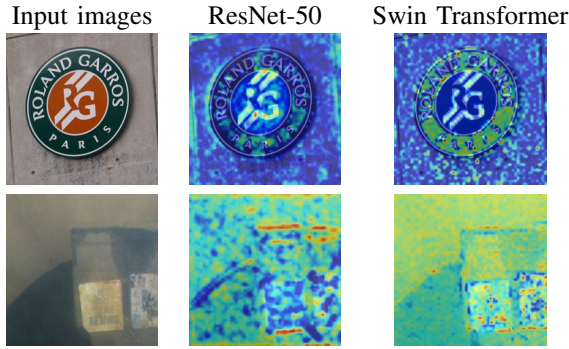


Fig. 4. Illustrating the **advantage of Swin-Transformer** over the ResNet-50 as a backbone generating better-localized representation.

each corresponding to a text instance, as  $S^{(i)}$ . Each query has many sub-queries  $s_j$ , where  $S^{(i)} = (s_1^{(i)}, \dots, s_N^{(i)})$ . Then the initial control points are passed through the location decoder with multiple layers, followed by a classification head that predicts the confidence from the final control points alongside a two-channel regression head generating the normalized coordinates for each. Here the control points are the polygon points starting from the top left corner in clockwise order.

*Text Recognition Decoder.* The character decoder is almost the same as the location decoder, only we change the control points queries with character queries  $R^{(i)}$ . Both queries,  $S^{(i)}$  and  $R^{(i)}$  with the same index belong to the same text instance. So, during the prediction, each decoder predicts the control points and characters for the same instance. In the end, a classification head predicts the multiple character classes based on the final character queries.

### B. Training Strategies

The training strategies proposed for the DA-TextSpotter framework have been elucidated in the following two subsections:

**Domain-agnostic Pretraining.** In this work, we propose a novel training strategy for pretraining which incorporates multiple domain source data from natural scenes and underwater images. The key objective is for the model to learn both domain-independent and domain-specific parameters simultaneously during the pre-training phase which helps them to improve generalization capacity during the fine-tuning phase. Experimental studies have been carried out to effectively explore this strategy in the next section.

**Learning Objectives.** The overall losses used in this work can be summarised under the encoder  $\mathcal{L}_{\text{enc}}$  and decoder  $\mathcal{L}_{\text{dec}}$  blocks shown in eqns. 2 and 3 respectively,

$$\mathcal{L}_{\text{enc}} = \sum_i \left( \lambda_{\text{cls}} \mathcal{L}_{\text{cls}}^{(i)} + \lambda_{\text{coord}} \mathcal{L}_{\text{coord}}^{(i)} + \lambda_{\text{glou}} \mathcal{L}_{\text{glou}}^{(i)} \right) \quad (2)$$

$$\mathcal{L}_{\text{dec}} = \sum_j \left( \lambda_{\text{cls}} \mathcal{L}_{\text{cls}}^{(j)} + \lambda_{\text{coord}} \mathcal{L}_{\text{coord}}^{(j)} + \lambda_{\text{char}} \mathcal{L}_{\text{char}}^{(j)} \right) \quad (3)$$

Where  $\mathcal{L}_{\text{cls}}^{(j)}$  is focal loss for classification of text instances.  $\mathcal{L}_{\text{coord}}^{(j)}$  is L-1 loss used for control point coordinate regression.  $\mathcal{L}_{\text{char}}^{(j)}$  is Cross entropy loss for character classification.  $\mathcal{L}_{\text{glou}}$  is the generalized IoU loss defined in [43] for bounding box regression.  $\lambda_{\text{cls}}$ ,  $\lambda_{\text{char}}$ ,  $\lambda_{\text{coord}}$ ,  $\lambda_{\text{cls}}$ , and  $\lambda_{\text{glou}}$  are the weighting factor for the losses.

## IV. EXPERIMENTS

### A. Datasets and Evaluation

The UWT dataset <sup>1</sup> has been proposed in this work to evaluate the performance of the proposed DA-TextSpotter framework. The sample images have been snapped from YouTube videos broadcasting the retrieval of lost items in rivers, seas, and pools. On average, there is just one text instance for each image. All the text instances are in English. Overall, the dataset consists of 250 total images with dimensions of  $400 \times 400$  pixels, split into 200 images for training and 50 images for testing. For natural scene benchmarks, CTW1500 [32] and TotalText [5] has been used for arbitrary-shaped text, while ICDAR15 [22] has been used for regular text. The Precision (P), Recall (R) and F1 (F) metrics has been adapted for text detection while end-to-end recognition score use the None (without using lexica) and Full (using lexica) settings.

### B. Experiments on Text Spotting

In this subsection, we highlight the final text spotting results on the different benchmarks.

**Arbitrary-shaped Text.** For the irregular text we have used two datasets Total-Text and CTW1500 as described in Section IV-A. We show our quantitative results in Table I. In terms of text detection, we can see our method outperforms TESTR and Swintextspotter in terms of Hmean in the Total-Text dataset. In the case of the CTW1500 dataset, the same surpluses these two methods in terms of all the metrics. For text spotting our supervised method outperform all the methods in terms of Full lexicon metrics. In the Total-Text dataset, we outperformed TESTR by almost 1%. For CTW1500 we outperform TESTR, Swintetextspotter, and Abinet ++ by 3%, almost 6%, and almost 2% respectively. Qualitative results are shown in Fig 5. In the first column, we show the results of Total-text and in the second column, we show the results of CTW1500. Here we can see our method can spot both the regular and irregular text in these datasets. In summary, we can say that the qualitative and quantitative results represent the effectiveness of our model.

**Regular Text.** ICDAR15 has been used as a regular text benchmark. The performance of DA-TextSpotter compared to SOTA approaches is shown in Table II. For Text detection, we achieved a 2% gain in precision which improved slightly in the F-measure over TESTR. In the text spotting task, our method achieved the best result in the ‘‘Strong’’ type outperforming TESTR, Swintextspotter, and ABINet++ by almost 1%, 2%, and 0.5% respectively. In Fig 5, the 3rd

<sup>1</sup>The dataset will be made publicly available upon acceptance.

TABLE I  
TEXT SPOTTING PERFORMANCE ON TOTAL-TEXT AND CTW1500. “NONE” REFERS TO RECOGNITION WITHOUT LEXICON.  
THE “FULL” LEXICON CONTAINS ALL THE WORDS IN THE TEST SET

Methods	Total-Text					CTW1500				
	Detection			End-to-end		Detection			End-to-end	
	P	R	F	None	Full	P	R	F	None	Full
ABCNet[31]	-	-	-	64.2	75.7	-	-	81.4	45.2	74.1
Text Perceptron[39]	88.8	81.8	85.2	69.7	78.3	87.5	81.9	84.6	57.0	-
ABCNet v2[33]	90.2	84.1	87.0	70.4	78.1	85.6	83.8	84.7	57.5	77.2
MANGO[38]	-	-	-	72.9	83.6	-	-	-	<b>58.9</b>	78.7
TESTR[59]	<b>93.4</b>	81.4	86.90	73.25	83.3	89.7	83.1	86.3	53.3	79.9
Swintextspotter[18]	-	-	88.0	<b>74.3</b>	84.1	-	-	88.0	51.8	77.0
<b>DA-TextSpotter</b>	91.13	<b>86.31</b>	<b>88.66</b>	71.9	<b>85.01</b>	90.98	<b>85.53</b>	88.17	53.96	82.15
<b>DA-TextSpotter (w/o UWT)</b>	91.21	85.77	88.41	72.56	84.44	<b>91.45</b>	85.16	<b>88.19</b>	56.02	<b>82.91</b>




Fig. 5. Illustration of our method on different datasets. 1st column from Total-Text, 2nd column from CTW1500, 3rd column from ICDAR15, 4th and 5th columns from UWT. [Use 200% zoom for better visualization of the qualitative results]

TABLE II  
TEXT SPOTTING PERFORMANCE ON THE ICDAR-15 DATASET. “S”, “W”, “G” REPRESENTS RECOGNITION WITH “STRONG”, “WEAK”, “GENERIC”, LEXICA, RESPECTIVELY

Methods	Detection			End-to-end		
	P	R	F	S	W	G
	Unconstrained[40]	89.4	87.5	87.5	83.4	79.9
Text Perceptron[39]	89.4	82.5	87.1	80.5	76.6	65.1
ABCNet v2 [33]	90.4	86.0	88.1	82.7	78.5	73.0
MANGO [38]	-	-	-	81.8	78.9	67.3
PGNet[51]	91.8	84.8	88.2	83.3	78.3	63.5
TESTR[59]	90.3	<b>89.7</b>	90.0	85.2	79.4	73.6
Swintextspotter[18]	-	-	-	83.9	77.3	70.5
Abinet++[10]	-	-	-	86.1	81.9	<b>77.8</b>
<b>DA-TextSpotter</b>	<b>92.60</b>	88.59	<b>90.55</b>	85.92	80.44	74.16
<b>DA-TextSpotter(w/o UWT)</b>	91.49	89.07	90.27	<b>86.78</b>	<b>81.90</b>	75.82

column shows how our method performs on ICDAR15. UWT. In UWT we compare it with SOTA methods and our previous method. We can see almost 5%, 3%, and 4% improvement compared with Banerjee *et al.* in terms of detection Precision, Recall, and F-measure respectively. Table III shows how we outperform the SOTA methods in terms of all the metrics. We also show some qualitative results of our method in Fig 5 (last two columns).

### C. Domain Generalization

Normally it is seen that if we train our model with one type of image and after that if we train it with noisy images it loses its performance drastically its called the catastrophic forgetting phenomenon of deep learning neural network. To address this phenomenon we perform zero-shot experiments.

TABLE III  
TEXT SPOTTING PERFORMANCE OF THE PROPOSED AND THE SOTA SYSTEMS ON UWT.

Methods	Detection			End-to-end	#Params
	P	R	F	None	
Banerjee <i>et al.</i> [3]	90.25	45.37	60.38	-	-
TESTR[59]	92.24	33.86	49.54	29.63	<b>55M</b>
Swintextspotter [18]	83.21	34.49	48.77	29.08	150M
<b>DA-TextSpotter</b>	<b>95.65</b>	<b>48.73</b>	<b>64.57</b>	<b>64.15</b>	60M

TABLE IV  
DOMAIN GENERALIZATION SETTINGS SHOWING THE EFFECTIVENESS OF ADDING UWT

Dataset	Method	Detection			End-to-end
		P	R	F	None
ICDAR2015	DA-TextSpotter(w/o UWT)	83.41	84.74	84.07	52.20
	DA-TextSpotter	<b>91.08</b>	<b>84.59</b>	<b>87.72</b>	<b>84.02</b>
Total-Text	DA-TextSpotter(w/o UWT)	88.42	67.25	76.40	59.77
	DA-TextSpotter	<b>91.76</b>	<b>76.42</b>	<b>83.39</b>	<b>64.67</b>
CTW1500	DA-TextSpotter(w/o UWT)	39.19	36.56	37.71	0
	DA-TextSpotter	<b>41.69</b>	<b>44.35</b>	<b>42.98</b>	<b>0</b>

First, we train our model with Synthtext and Icdar17 dataset namely Pre-train Scene Text. We evaluate it on the SOTA datasets, Total-text, CTW1500, and ICDAR15. After that, we again train it with Underwater images. As it consists of more noisy images. We expected it will degrade the performance of the model. But in Table IV it improves the result. In that case, the End-to-End result of CTW1500 is 0 because it is text-line based annotated, and in the pretraining, our model

only learns word-level annotation.

TABLE V

ABLATION FOR DIFFERENT FEATURE EXTRACTION BACKBONES. WHERE “P” STANDS FOR PRECISION, “R” STANDS FOR RECALL, “F” STANDS FOR F-MEASURE RESPECTIVELY, AND “NONE” REFERS NO LEXICON HAS BEEN USED.

Methods	Detection			End-to-end
	P	R	F	None
Resnet-50 [15]	88.87	76.47	82.20	60.06
ViT-T [9]	90.17	72.90	80.62	59.40
<b>Swin-T [35]</b>	<b>93.59</b>	<b>75.20</b>	<b>83.40</b>	<b>67.06</b>

TABLE VI

RESULTS FOR DIFFERENT ENHANCEMENT METHODS ON UWT. WHERE “P” STANDS FOR PRECISION, “R” STANDS FOR RECALL, “F” STANDS FOR F-MEASURE RESPECTIVELY, AND “NONE” REFERS NO LEXICON HAS BEEN USED.

Methods	Detection			End-to-end
	P	R	F	None
Without Enhancement	68.06	31.01	42.61	16.09
SRGAN [24]	85.54	40.76	55.21	56.32
ESRGAN[56]	90.20	42.56	57.83	58.32
Real-ESRGAN[54]	<b>95.65</b>	<b>48.73</b>	<b>64.57</b>	<b>64.15</b>

#### D. Ablation Studies

To understand the significance of the different components of the DA-TextSpotter framework, we conduct ablation studies demonstrated as follows:

**Effectiveness of Swin Transformer.** The usage of Swin-Tiny [35] feature extraction backbone in the DA-TextSpotter framework shows a significant performance gain over ViT-Tiny [9] and Resnet-50 [15] backbones in both detection and recognition metrics. Almost a 4% change in detection precision and a 1% change in the detection F-measure over Resnet-50 is observed. Further, we get almost a massive 7% improvement in the End-to-End recognition results when it is run on the Total-Text [5] dataset as shown in Table V. This justifies that the window-based local attention computed with Swin-T can be effective for specially text regions with smaller local boundaries. Fig. 4 illustrates how hierarchical feature maps help for feature extraction, especially for smaller objects depicted in the underwater images.

**Effectiveness of Super-resolution Unit.** To justify the importance of the Super-resolution unit, we conducted an ablation on our proposed UWT benchmark which contains the most complex low-resolution images. Upsampling the images with the different Super-resolution GAN variants which include SR-GAN [24], ESR-GAN [56], and Real-ESRGAN [54]. The performance of Real-ESRGAN is extremely high compared to other approaches because of the pre-trained OutdoorSceneTraining [55] dataset, giving an almost 7% improvement over both detection F-measure and end-to-end metrics over second-best ESRGAN as shown in

TABLE VII

REPRESENTATION QUALITY. FINE-TUNING PERFORMANCE OF A PRE-TRAINED MODEL USING FROZEN BACKBONES. ALSO, ‘NONE’ REFERS TO NO LEXICON HAS BEEN USED.

Architectural blocks			Detection			End-to-End
Image Encoder	Location Decoder	Character Decoder	P	R	F	None
✗	✗	✓	92.09	84.64	88.21	70.76
✓	✗	✗	92.05	80.49	85.88	68.9
✗	✓	✗	90.22	85.82	87.69	69.75
✗	✗	✓	<b>93.29</b>	82.88	87.7	72.54
✓	✗	✓	93.23	80.08	86.57	<b>73.39</b>
✓	✓	✗	92.88	84.28	88.37	70.74
✓	✓	✓	91.21	<b>85.77</b>	<b>88.41</b>	72.56

Table VI. Also enhancing the images shows a substantial gain of **22%** and **48%** for detection and end-to-end metrics respectively. Fig. 3 illustrates the impact of enhancement unit on reading under complexities where text is illegible for even human vision.

**Representation Quality.** To evaluate the representation quality of the features learned during pre-training and the need for the encoder and dual decoder blocks in the DA-TextSpotter pipeline we did an exhaustive evaluation by freezing and unfreezing them. We performed all the possible combinations of frozen and unfrozen decoder and encoder to show the representation quality. All the results are illustrated in Table VII performed on the Total-Text dataset in the fine-tuning configuration. The consistent performances on both detection and recognition tasks for all the different configurations show the model’s robustness.

#### V. CONCLUSION AND FUTURE WORK

This work demonstrates the capability of a scene text spotting model to improve its robustness and generalization capability when trained across multiple domains, in this case, natural scenes and underwater. The super-resolution unit introduced in the DA-TextSpotter framework boosts model performance in more adverse and noisy underwater images. DA-TextSpotter also brings efficiency while using a tiny Swin-transformer backbone for extracting visual cues, both at the local and global levels of hierarchies. Finally, multi-domain learning could help a model to learn domain-agnostic parameters, without forgetting already learned domain-specific parameters. The future challenge includes going deeper into domain-incremental settings and using faster attention modules [7] to improve text spotting performance.

#### ACKNOWLEDGEMENT

This work acknowledges IDEAS - Technology Innovation Hub (ISI Kolkata), the Spanish projects GRAIL PID2021-126808OB-I00, DocAI 2021-SGR-01559, the CERCA Program / Generalitat de Catalunya, the FCT-19-15244, and PhD Scholarship from AGAUR 2023 FI-3-00223 and special thanks to Saumik Bhattacharya, Subhankar Ghosh, Prasun Roy and Kunal Biswas for their insightful discussions.

## REFERENCES

- [1] Derya Akkaynak, Tali Treibitz, Tom Shlesinger, Yossi Loya, Raz Tamir, and David Iluz. What is the space of attenuation coefficients in underwater computer vision? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4931–4940, 2017.
- [2] Rowel Atienza. Vision transformer for fast and efficient scene text recognition. In *Document Analysis and Recognition-ICDAR 2021: 16th International Conference, Lausanne, Switzerland, September 5–10, 2021, Proceedings, Part I 16*, pages 319–334. Springer, 2021.
- [3] Ayan Banerjee, Palaiahnakote Shivakumara, Soumyajit Pal, Umapada Pal, and Cheng-Lin Liu. Dct-dwt-fft based method for text detection in underwater images. In *Pattern Recognition: 6th Asian Conference, ACPR 2021, Jeju Island, South Korea, November 9–12, 2021, Revised Selected Papers, Part II*, pages 218–233. Springer, 2022.
- [4] Carl Case, Bipin Suresh, Adam Coates, and Andrew Y Ng. Autonomous sign reading for semantic mapping. In *2011 IEEE International Conference on Robotics and Automation*, pages 3297–3303. IEEE, 2011.
- [5] Chee-Kheng Ch'ng, Chee Seng Chan, and Cheng-Lin Liu. Total-text: toward orientation robustness in scene text detection. *International Journal on Document Analysis and Recognition (IJ DAR)*, 23(1):31–52, 2020.
- [6] Alloy Das, Sanket Biswas, Ayan Banerjee, Josep Lladós, Umapada Pal, and Saumik Bhattacharya. Harnessing the power of multilingual datasets for pre-training: Towards enhancing text spotting performance. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 718–728, 2024.
- [7] Alloy Das, Prasun Roy, Saumik Bhattacharya, Subhankar Ghosh, Umapada Pal, and Michael Blumenstein. Fast: Font-agnostic scene text editing. *arXiv preprint arXiv:2308.02905*, 2023.
- [8] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part IV 13*, pages 184–199. Springer, 2014.
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [10] Shancheng Fang, Zhendong Mao, Hongtao Xie, Yuxin Wang, Chenggang Yan, and Yongdong Zhang. Abinet++: Autonomous, bidirectional and iterative language modeling for scene text spotting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [11] Shancheng Fang, Hongtao Xie, Yuxin Wang, Zhendong Mao, and Yongdong Zhang. Read like humans: Autonomous, bidirectional and iterative language modeling for scene text recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7098–7107, 2021.
- [12] Sergi Garcia-Bordils, George Tom, Sangeeth Reddy, Minesh Mathew, Marçal Rusiñol, CV Jawahar, and Dimosthenis Karatzas. Read while you drive-multilingual text tracking on the road. In *Document Analysis Systems: 15th IAPR International Workshop, DAS 2022, La Rochelle, France, May 22–25, 2022, Proceedings*, pages 756–770. Springer, 2022.
- [13] Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013.
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [16] Tong He, Zhi Tian, Weilin Huang, Chunhua Shen, Yu Qiao, and Changming Sun. An end-to-end textspotter with explicit alignment and attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5020–5029, 2018.
- [17] Weilin Hou, Sarah Woods, Ewa Jarosz, Wesley Goode, and Alan Weidemann. Optical turbulence on underwater image degradation in natural environments. *Applied optics*, 51(14):2678–2686, 2012.
- [18] Mingxin Huang, Yuliang Liu, Zhenghao Peng, Chongyu Liu, Dahua Lin, Shenggao Zhu, Nicholas Yuan, Kai Ding, and Lianwen Jin. Swintextspotter: Scene text spotting via better synergy between text detection and text recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4593–4603, 2022.
- [19] Yili Huang, Chengyu Gu, Shilin Wang, Zheng Huang, Kai Chen, and Hui Autonomous Region. Spatial aggregation for scene text recognition. In *Proceedings of the 32nd British Machine Vision Conference*, 2021.
- [20] Alexia Jolicoeur-Martineau. The relativistic discriminator: a key element missing from standard gan. *arXiv preprint arXiv:1807.00734*, 2018.
- [21] Mahesh Joshi, Mark Dredze, William Cohen, and Carolyn Rose. Multi-domain learning: when do domains matter? In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 1302–1312, 2012.
- [22] Dimosthenis Karatzas, Lluís Gomez-Bigorda, Angelos Nicolaou, Suman Ghosh, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, et al. Icdar 2015 competition on robust reading. In *2015 13th international conference on document analysis and recognition (ICDAR)*, pages 1156–1160. IEEE, 2015.
- [23] Stefan Larson, Yi Yang Gordon Lim, Yutong Ai, David Kuang, and Kevin Leach. Evaluating out-of-distribution performance on document image classifiers. *Advances in Neural Information Processing Systems*, 35:11673–11685, 2022.
- [24] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.
- [25] Chongyi Li, Chunle Guo, Wenqi Ren, Runmin Cong, Junhui Hou, Sam Kwong, and Dacheng Tao. An underwater image enhancement benchmark dataset and beyond. *IEEE Transactions on Image Processing*, 29:4376–4389, 2019.
- [26] Hui Li, Peng Wang, and Chunhua Shen. Towards end-to-end text spotting with convolutional recurrent neural networks. In *Proceedings of the IEEE international conference on computer vision*, pages 5238–5246, 2017.
- [27] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.
- [28] Minghui Liao, Guan Pang, Jing Huang, Tal Hassner, and Xiang Bai. Mask textspotter v3: Segmentation proposal network for robust scene text spotting. In *European Conference on Computer Vision*, pages 706–722. Springer, 2020.
- [29] Minghui Liao, Baoguang Shi, Xiang Bai, Xinggang Wang, and Wenyu Liu. Textboxes: A fast text detector with a single deep neural network. In *Thirty-first AAAI conference on artificial intelligence*, 2017.
- [30] Xuebo Liu, Ding Liang, Shi Yan, Dagui Chen, Yu Qiao, and Junjie Yan. Fots: Fast oriented text spotting with a unified network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5676–5685, 2018.
- [31] Yuliang Liu, Hao Chen, Chunhua Shen, Tong He, Lianwen Jin, and Liangwei Wang. Abcnet: Real-time scene text spotting with adaptive bezier-curve network. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9809–9818, 2020.
- [32] Yuliang Liu, Lianwen Jin, Shuaitao Zhang, Canjie Luo, and Sheng Zhang. Curved scene text detection via transverse and longitudinal sequence connection. *Pattern Recognition*, 90:337–345, 2019.
- [33] Yuliang Liu, Chunhua Shen, Lianwen Jin, Tong He, Peng Chen, Chongyu Liu, and Hao Chen. Abcnet v2: Adaptive bezier-curve network for real-time end-to-end text spotting. *arXiv preprint arXiv:2105.03620*, 2021.
- [34] Yajing Liu, Xinmei Tian, Ya Li, Zhiwei Xiong, and Feng Wu. Compact feature learning for multi-domain image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7193–7201, 2019.
- [35] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021.
- [36] Pengyuan Lyu, Minghui Liao, Cong Yao, Wenhao Wu, and Xiang Bai. Mask textspotter: An end-to-end trainable neural network for spotting text with arbitrary shapes. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 67–83, 2018.
- [37] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain

- generalization via invariant feature representation. In *International conference on machine learning*, pages 10–18. PMLR, 2013.
- [38] Liang Qiao, Ying Chen, Zhanzhan Cheng, Yunlu Xu, Yi Niu, Shiliang Pu, and Fei Wu. Mango: A mask attention guided one-stage scene text spotter. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 2467–2476, 2021.
- [39] Liang Qiao, Sanli Tang, Zhanzhan Cheng, Yunlu Xu, Yi Niu, Shiliang Pu, and Fei Wu. Text perceptron: Towards end-to-end arbitrary-shaped text spotting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 11899–11907, 2020.
- [40] Siyang Qin, Alessandro Bissacco, Michalis Raptis, Yasuhisa Fujii, and Ying Xiao. Towards unconstrained end-to-end text spotting. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4704–4714, 2019.
- [41] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. *Advances in neural information processing systems*, 30, 2017.
- [42] Sangeeth Reddy, Minesh Mathew, Lluís Gomez, Marçal Rusinol, Dimosthenis Karatzas, and CV Jawahar. Roadtext-1k: Text detection & recognition dataset for driving videos. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11074–11080. IEEE, 2020.
- [43] Hamid Rezatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 658–666, 2019.
- [44] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [45] Amir Rosenfeld and John K Tsotsos. Incremental learning through deep adaptation. *IEEE transactions on pattern analysis and machine intelligence*, 42(3):651–663, 2018.
- [46] Prasun Roy, Subhankar Ghosh, Saumik Bhattacharya, and Umapada Pal. Effects of degradations on deep neural network architectures. *arXiv preprint arXiv:1807.10108*, 2018.
- [47] Mohamed Ali Souibgui, Sanket Biswas, Andres Mafla, Ali Furkan Biten, Alicia Fornés, Yousri Kessentini, Josep Lladós, Lluís Gomez, and Dimosthenis Karatzas. Text-diae: Degradation invariant autoencoders for text recognition and document enhancement. *arXiv preprint arXiv:2203.04814*, 2022.
- [48] Jordy Van Landeghem, Sanket Biswas, Matthew Blaschko, and Marie-Francine Moens. Beyond document page classification: Design, datasets, and challenges. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2962–2972, 2024.
- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [50] Kai Wang, Boris Babenko, and Serge Belongie. End-to-end scene text recognition. In *2011 International conference on computer vision*, pages 1457–1464. IEEE, 2011.
- [51] Pengfei Wang, Chengquan Zhang, Fei Qi, Shanshan Liu, Xiaoqiang Zhang, Pengyuan Lyu, Junyu Han, Jingtuo Liu, Errui Ding, and Guangming Shi. Pgnnet: Real-time arbitrarily-shaped text spotting with point gathering network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 2782–2790, 2021.
- [52] Wenhai Wang, Enze Xie, Xiang Li, Xuebo Liu, Ding Liang, Yang Zhibo, Tong Lu, and Chunhua Shen. Pan++: Towards efficient and accurate end-to-end spotting of arbitrarily-shaped text. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [53] Xiaobing Wang, Yingying Jiang, Zhenbo Luo, Cheng-Lin Liu, Hyunsoo Choi, and Sungjin Kim. Arbitrary shape scene text detection with adaptive text region representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6449–6458, 2019.
- [54] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1905–1914, 2021.
- [55] Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Recovering realistic texture in image super-resolution by deep spatial feature transform. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 606–615, 2018.
- [56] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European conference on computer vision (ECCV) workshops*, pages 0–0, 2018.
- [57] Junjie Wen, Jinqiang Cui, Zhenjun Zhao, Ruixin Yan, Zhi Gao, Lihua Dou, and Ben M Chen. Syreanet: A physically guided underwater image enhancement framework integrating synthetic and real images. *International Conference on Robotics and Automation*, 2023.
- [58] Chuhui Xue, Jiaxing Huang, Wenqing Zhang, Shijian Lu, Changhu Wang, and Song Bai. Image-to-character-to-word transformers for accurate scene text recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [59] Xiang Zhang, Yongwen Su, Subarna Tripathi, and Zhuowen Tu. Text spotting transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9519–9528, 2022.
- [60] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020.