

GA-DAN: Geometry-Aware Domain Adaptation Network for Scene Text Detection and Recognition

Fangneng Zhan, Chuhui Xue, and Shijian Lu

Nanyang Technological University

{fnzhan, Shijian.Lu}@ntu.edu.sg, XUEC0003@e.ntu.edu.sg

Abstract

Recent adversarial learning research has achieved very impressive progress for modelling cross-domain data shifts in appearance space but its counterpart in modelling cross-domain shifts in geometry space lags far behind. This paper presents an innovative Geometry-Aware Domain Adaptation Network (GA-DAN) that is capable of modelling cross-domain shifts concurrently in both geometry space and appearance space and realistically converting images across domains with very different characteristics. In the proposed GA-DAN, a novel multi-modal spatial learning technique is designed which converts a source-domain image into multiple images of different spatial views as in the target domain. A new disentangled cycle-consistency loss is introduced which balances the cycle consistency in appearance and geometry spaces and improves the learning of the whole network greatly. The proposed GA-DAN has been evaluated for the classic scene text detection and recognition tasks, and experiments show that the domain-adapted images achieve superior scene text detection and recognition performance while applied to network training.

1. Introduction

A large amount of labelled or annotated images is critical for training robust and accurate deep neural network (DNN) models, but collecting and annotating large datasets are often extremely expensive. In addition, state-of-the-art DNN models usually assume that images in the training and inference stages are collected under similar conditions which often experience clear performance drops while applied to images from different domains. Such lack of scalability and transferability makes collection and annotation even more expensive while dealing with images collected under different conditions from different domains. Unsupervised Do-

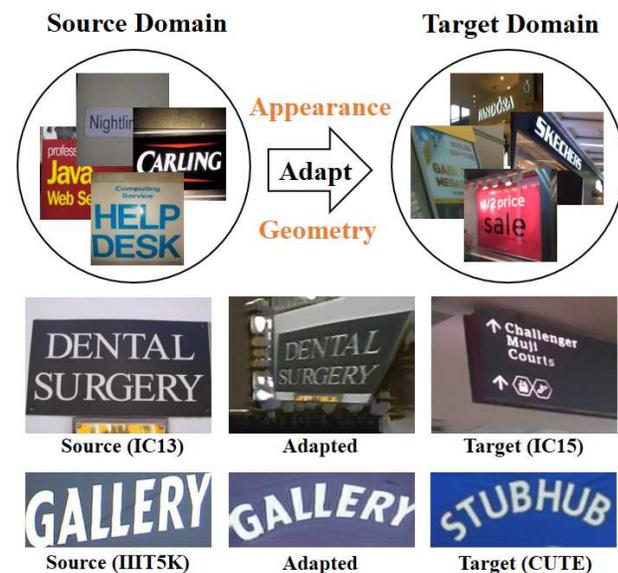


Figure 1. Domain adaptation by the proposed GA-DAN: For scene text images with clear shifts from the **Source Domain** to the **Target Domain**, GA-DAN models the domain shifts in appearance and geometry spaces simultaneously and generates **Adapted images** with high-fidelity in both appearance and geometry spaces.

main Adaptation (DA), which transfers images and features from a source domain to a target domain, has achieved very impressive performance especially with the recent advances of Generative Adversarial Networks (GANs) [11]. Different DA techniques have been developed and applied to different computer vision problems successfully such as style transfer, image synthesis, etc.

State-of-the-art DA still faces various problems. In particular, most existing systems focus on learning feature shift in appearance space only whereas the feature shift in geometry space is largely neglected. On the other hand, images from different domains often differ in both appearance and geometry spaces. Take various texts in scenes as an ex-

ample. They could suffer from motion blurs in appearance space and perspective distortion in geometry space concurrently as shown in the target domain images in Fig. 1, and both are essential features for learning robust and accurate scene text detectors and recognizers. As a result, existing techniques often suffer from a clear performance drop when source and target domains have clear geometry discrepancy as observed for images and videos from different domains.

We design an innovative Geometry-Aware Domain Adaptation Network (GA-DAN), an end-to-end trainable network that learns and models domain shifts in appearance and geometry spaces simultaneously as illustrated in the last two rows of Fig. 1. One unique feature of the proposed GA-DAN is a multi-modal spatial learning structure that learns multiple spatial transformations and converts a source-domain image into multiple target-domain images realistically as illustrated in Figs. 4 and 5. In addition, a novel disentangled cycle-consistency loss is designed which guides the GA-DAN learning towards optimal transfer concurrently in both geometry and appearance spaces. The GA-DAN takes the cycle structure as illustrated in Fig. 2, where the spatial modules (S_X and S_Y) model the feature shifts in geometry space and the generators (G_X and G_Y) complete the blank as introduced by the spatial transformation and model the feature shifts in appearance space. The discriminators discriminate not only ‘fake image’ and ‘real image’ but also ‘fake transformation’ and ‘real transformation’, leading to optimal modelling of domain and feature shifts in geometry and appearance spaces.

The contributions of this work are threefold. First, it designs a novel network that models domain shifts in geometry and appearance spaces simultaneously. To the best of our knowledge, this is the first network that performs domain adaptation in both spaces concurrently. Second, it designs an innovative multi-modal spatial learning mechanism and introduces a spatial transformation discriminator to achieve multi-modal adaptation in geometry space. Third, it designs a disentangled cycle-consistency loss that balances the cycle-consistency for concurrent adaptation in appearance and geometry spaces and can also be applied to generic domain adaptation.

2. Related Works

2.1. Domain Adaptation

Domain adaptation is an emerging research topic that aims to address domain shift and dataset bias [48, 61]. Existing techniques can be broadly classified into two categories. The first category focuses on minimizing discrepancies between the source domain and the target domain in the feature space. For example, [34] explored Maximum Mean Discrepancies (MMD) and Joint MMD distance across domains over fully-connected layers. [55] studied feature

adaptation by minimizing the correlation distance and then extended it to deep architectures [56]. [4] modelled domain-specific features to encourage networks to learn domain-invariant features. [10, 62] improved feature adaptation by designing various adversarial objectives.

The second category adopts Generative Adversarial Nets (GANs) [11] to perform pixel-level adaptation via continuous adversarial learning between generators and discriminators which has achieved great success in image generation [8, 45, 75], image composition [30, 74, 70] and image-to-image translation [79, 19, 52]. Different approaches have been investigated to address pixel-level image transfer by enforcing consistency in the embedding space. [57] translates a rendering image to a real image by using conditional GANs. [3] studies an unsupervised approach to learn pixel-level transfer across domains. [31] proposes an unsupervised image-to-image translation framework using a shared-latent space. [9] introduces an inference model that jointly learns a generation network and an inference network. More recently, CycleGAN [79] and its variants [68, 26] achieve very impressive image translation by using cycle-consistency loss. [16] proposes a cycle-consistent adversarial model that adapts at both pixel and feature levels.

2.2. Scene Text Detection and Recognition

Automated detection and recognition of various texts in scenes has attracted increasing interests as witnessed by increasing benchmarking competitions [25, 51]. Different detection techniques have been proposed from those earlier using hand-crafted features [42, 36] to the recent using DNNs [77, 22, 69, 60, 72, 65]. Different detection approaches have also been explored including character based [18, 59, 23, 15, 17], word-based [22, 28, 33, 14, 58, 78, 32, 63, 38, 39, 44, 73, 7, 35, 29] and the recent line-based [76]. Meanwhile, different scene text recognition techniques have been developed from the earlier recognizing characters directly [20, 67, 47, 1, 12, 21] to the recent recognizing words or text lines using recurrent neural network (RNN), [49, 53, 54, 50] and attention models [27, 5, 71].

Similar to other detection and recognition tasks, training accurate and robust scene text detectors and recognizers requires a large amount of annotated training images. On the other hand, most existing datasets such as ICDAR2015 [25] and Total-Text [6] contain a few hundred or thousand training images only which has become one major factor that impedes the advance of scene text detection and recognition research. The proposed domain adaptation technique addresses this challenge by transferring existing annotated scene text images to a new target domain, hence alleviate the image collection and annotation efforts greatly.

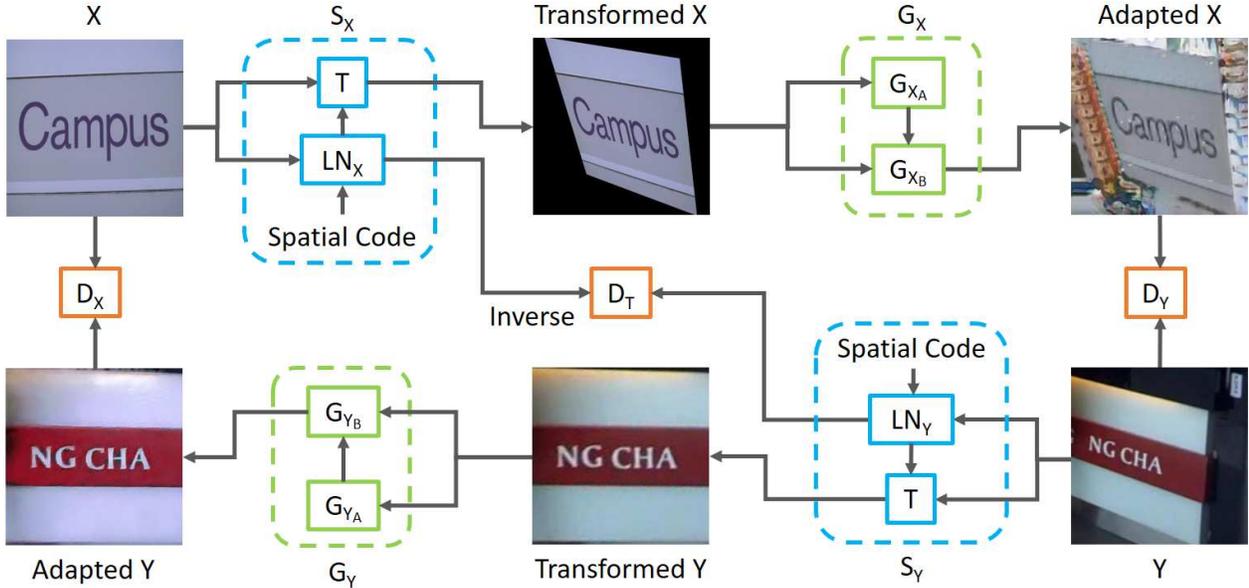


Figure 2. The structure of the proposed GA-DAN: S_X (or S_Y) represents the spatial modules as enclosed in blue-color boxes which consist of *Spatial Code*, transformation module T and localization network LN_X (or LN_Y) that predict transformation matrix and transform input images. G_X (or G_Y) denote generators consisting of G_{X_A} (or G_{Y_A}) and G_{X_B} (or G_{Y_B}) as enclosed in green-color boxes that complete the background and translate the image style, respectively. D_X , D_Y and D_T within orange-color boxes denote different discriminators.

3. Methodology

We propose an innovative geometry-aware domain adaptation network (GA-DAN) that performs multi-modal domain adaptation concurrently in both spatial and appearance spaces as shown in Fig. 2. Detailed network architecture, multi-modal spatial learning and adversarial training strategy will be presented in the following three subsections.

3.1. GA-DAN Architecture

The GA-DAN consists of spatial modules, generators and discriminators as enclosed within blue-color, green-color and orange-color boxes, respectively, as shown in Fig. 2. The overall network is designed in a cycle structure, where the mappings between the source domain X and the target domain Y are learned by sub-modules $X \rightarrow Adapted X$ ($X \rightarrow Y$) and $Y \rightarrow Adapted Y$ ($Y \rightarrow X$), respectively. In the $X \rightarrow Y$ mapping, the spatial module S_X transforms images in X to new images in *Transformed X* that has similar spatial styles as Y . The generator G_X then completes the blank as introduced by the spatial transformation and translates the completed images to new images in *Adapted X* that has similar appearance as Y . A discriminator D_Y attempts to distinguish *Adapted X* and Y which drives S_X and G_X to learn better spatial and appearance mapping from X to Y . Similar processes happen in the $Y \rightarrow X$ mapping as well.

The spatial modules S_X (as well as S_Y) has a localization network LN_X and a transformation module T for domain adaptation in geometry space, more details to be pre-

sented in the following subsection. The generator G_X (as well as G_Y) consists of two sub-generators G_{X_A} and G_{X_B} for adaptation in appearance space. In particular, the spatial module S_X will produce a binary map with 1 denoting pixels transformed from the original image and 0 for padded black background (not shown but can be inferred from the sample image in *Transformed X* in Fig. 2). Under the guidance of the binary map, G_{X_A} will learn from Y for new contents to complete the black background of the transformed image (as in *Transformed X*), and G_{X_B} further adapts the completed images to have similar appearance as Y as illustrated in Fig. 2. Our study shows that the *Adapted X* is quite blurry if a single generator is used to complete the black background and adapt the appearance. The use of the two dedicated generators G_{X_A} and G_{X_B} for background completing and appearance adaptation helps greatly for realistic adaptation in appearance space.

Note directly concatenating an appearance-transfer GAN (e.g. CycleGAN [79]) and a geometry-transfer GAN (e.g. ST-GAN [30]) does not perform well for simultaneous image adaptation in geometry and appearance spaces. Due to the co-existence of spatial and appearance shifts between the source and target domain images, the discriminator of the geometry-transfer GAN (or appearance-transfer GAN) will be confused by the appearance (or geometry) shift which leads to poor adversarial learning outcome. Our GA-DAN is an end-to-end trainable network that coordinates the learning in geometry and appearance spaces simultaneously and drives the network for optimal adaptation in both

spaces, more details to be presented in Section 3.3.

Table 1. Localization network LN_X and LN_Y within the multi-modal spatial learning shown in Fig. 2, N denotes the number of parameters.

Layers	Out Size	Configurations
Block1	128×128	$3 \times 3 \text{ conv}, 16, 2 \times 2 \text{ pool}$
Block2	64×64	$3 \times 3 \text{ conv}, 32, 2 \times 2 \text{ pool}$
Block3	32×32	$3 \times 3 \text{ conv}, 64, 2 \times 2 \text{ pool}$
Block4	16×16	$3 \times 3 \text{ conv}, 128, 2 \times 2 \text{ pool}$
Block5	8×8	$3 \times 3 \text{ conv}, 128, 2 \times 2 \text{ pool}$
FC1	512	-
FC2	N	-

3.2. Multi-Modal Spatial Learning

To generate images with different spatial views and features (similar to images in the target domain), we design a multi-modal spatial learning structure that learns multi-modal spatial transformations and maps a source-domain image to multiple target-domain images with different spatial views. Specifically, the multi-modal spatial learning first samples *Spatial Code* (i.e., random vectors) from normal distributions and then regresses it to predict spatial transformation matrix (according to the spatial features of the input image) by using a localization network LN_X (or LN_Y) as shown in Table 1. With the predicted transformation matrix that could be affine, homography or thin plate spline [2], the input image can be transformed to a new image with a different spatial view by T which performs actual transformation. Multiple new spatial views can be generated by running GA-DAN and sampling the *Spatial Code* multiple times, leading to the proposed multi-modal spatial mapping as illustrated in Figs. 4 and 5.

The multi-modal spatial learning as guided by D_X and D_Y tends to be unstable and hard to converge as the concurrent learning in geometry and appearance spaces is over-flexible and often entangled with each other. We address this issue by including a new discriminator D_T as shown in Fig. 2 which imposes certain constraints to the cyclic spatial learning and accordingly leads to more stable and efficient learning of the whole network. As shown in Fig. 2, S_X predicts a transformation matrix H_{XY} for mapping from domain X to domain Y , and S_Y predicts another transformation matrix H_{YX} for mapping from domain Y to domain X . The inverse matrix H_{XY}^{-1} can be derived from H_{XY} and it should be in the same transformation domain with the H_{YX} . D_T thus attempts to discriminates H_{XY}^{-1} and H_{YX} which drives the spatial transformations in two inverse directions to learn from each other. It bridges the spatial learning in opposite directions and imposes extra constraints in the geometry spaces, greatly improving the learning efficiency and learning stability of the whole network.

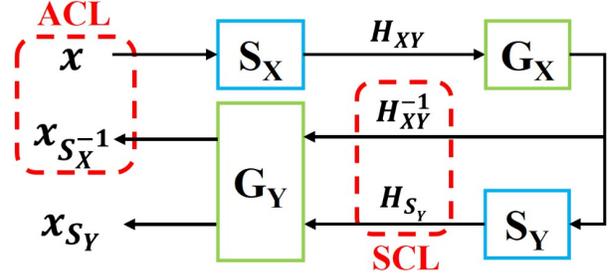


Figure 3. Illustration of the disentangled cycle-consistency loss: S_X , S_Y , G_X and G_Y denotes the spatial modules and generators, respectively, as shown in Fig. 2. x , $x_{S_X^{-1}}$ and x_{S_Y} refer to the input images in domain X , reconstructed image by the inverse transformation of S_X and reconstructed image by S_Y . H_{XY} and H_{S_Y} refer to the predicted transformation matrices by S_X and S_Y . ACL and SCL denote appearance cycle-consistency loss and spatial cycle-consistency loss which are obtained by calculating the L1 loss between $(x, x_{S_X^{-1}})$ and (H_{XY}^{-1}, H_{S_Y}) , respectively.

3.3. Adversarial Training

Due to the adaptation in geometry and appearance spaces, the adversarial learning needs to coordinate the minimization of cycle-consistency loss in both spaces properly. In addition, the adversarial learning also needs to take care of the new discriminator D_T as shown in Fig. 2. We design an innovative disentangled cycle-consistency loss and adversarial objective to tackle these challenges, more details to be described in the following two subsections.

Disentangled Cycle-Consistency Loss. We design a disentangled cycle-consistency loss that decomposes the cycle-consistency loss into a spatial cycle-consistency loss (SCL) and an appearance cycle-consistency loss (ACL) and balances their weights during learning. With spatial transformation involved, a small shift (due to inaccurate prediction of the spatial transformation matrix) in geometry space will lead to a very large cycle-consistency loss which can easily override the ACL and ruin the learning of the whole network. The decomposition of the cycle-consistency loss into ACL and SCL helps to address this issue effectively.

As shown in Fig. 3, the image x is fed into S_X to predict the transformation matrix H_{XY} and the transformed image is then fed to G_X for translation in appearance space. The translated image will be recovered in two different manners. First, it will be transformed by the inverse of H_{XY} (i.e. H_{XY}^{-1}) and further translated by G_Y to generate $x_{S_X^{-1}}$. Second, it will be passed to S_Y to predict the transformation matrix H_{S_Y} so as to be transformed by the estimated H_{S_Y} and further translated by G_Y to produce x_{S_Y} . Note the *Spatial Code* in S_X and S_Y are identical here so that the input image can be recovered in geometry space.

The $x_{S_X^{-1}}$ can be fully recovered from x in geometry space since the recovering matrix H_{XY}^{-1} is simply the in-

verse of H_{XY} . But $x_{S_X^{-1}}$ is different from x in appearance space. The ACL can thus be computed by L1 loss between x and $x_{S_X^{-1}}$ (only appearance difference exists) as follows:

$$ACL_X = E_{x \sim X} [\|x_{S_X^{-1}} - x\|] \quad (1)$$

Though $x_{S_X^{-1}}$ and x_{S_Y} differ only in geometry space, SCL cannot be obtained by computing L1 loss between them because a minor shift in geometry space will lead to a very large L1 loss. To ensure the spatial cycle-consistency, we obtain the SCL by directly computing the L1 loss of the transformation matrix H_{XY}^{-1} and H_{S_Y} as follows:

$$SCL = E_{x \sim X} [\|H_{XY}^{-1} - H_{S_Y}\|] \quad (2)$$

Further, the bordering regions of the original image may be lost by the spatial transformation which could affect the training seriously. While adapting an image from domain X to domain Y , the adaptation should ensure that all image information within the domain X is well preserved. Given the binary transformation map m from S_X , we can directly apply the inverse transformation H_{XY}^{-1} to m to obtain $m_{H_{XY}^{-1}}$. As the missing region by the spatial transformation will not be recovered, a region missing loss (RML) is defined for better preserving the transformed image as follows:

$$RML = E_{x \sim X} [\|m_{H_{XY}^{-1}} - m\|] \quad (3)$$

The overall cycle-consistency loss in the domain X can thus be formed as follows:

$$L_{cyc} = \lambda_{acl} ACL + \lambda_{scl} SCL + RML \quad (4)$$

where λ_{acl} and λ_{scl} are the weights of ACL and SCL .

Adversarial Objective. The adversarial objective of the mapping $X \rightarrow Y$ can be defined by:

$$\begin{aligned} L_{GAN} = & E_{y \sim Y} [\log D_Y(y)] \\ & + E_{x \sim X} [\log(1 - D_Y(G_X(S_X(x))))] \\ & + E_{y \sim Y} [\log D_T(H_{YX}^{-1})] + E_{x \sim X} [\log(1 - D_T(H_{XY}))] \end{aligned} \quad (5)$$

where H_{XY} and H_{YX}^{-1} are the transformation matrix for $X \rightarrow Y$ and the inverse transformation for $Y \rightarrow X$. S_X and G_X aim to minimize this objective while D_Y and D_T try to maximize it, i.e. $\min S_X G_X \max D_Y D_T L_{GAN}$. The objective of the mapping $Y \rightarrow X$ can be obtained similarly. Note to ensure that the translated image preserves features of the original image, an identity loss is included as follows:

$$L_{idt} = E_{x \sim X} [\|G_X(S_X(x)) * m - S_X(x) * m\|] \quad (6)$$

where m refers to the binary mask as produced by S_X .

4. Experiments

The proposed image adaptation technique has been evaluated over the scene text detection and recognition tasks.

4.1. Datasets

The experiments involve seven publicly available scene text detection and recognition datasets as listed:

ICDAR2013 [24] is used in the Robust Reading Competition in the International Conference on Document Analysis and Recognition (ICDAR) 2013. The images explicitly focused around the text content of interest. It contains 848 word images for network training and 1095 for testing.

ICDAR2015 [25] is used in the Robust Reading Competition under ICDAR2015. It contains incidental scene text images that appears in the scene without taking any specific prior action to improve its positioning / quality in the frame.

MSRA-TD500 [66] dataset consists of 500 natural images (300 for training, 200 for test), which are taken from indoor and outdoor scenes using a pocket camera. The indoor images mainly capture signs, doorplates and caution plates while the outdoor images mostly capture guide boards and billboards with complex background.

IIT5K [41] has 2000 training images and 3000 test images that are cropped from scene texts and born-digital images. Each word in this dataset has a 50-word lexicon and a 1000-word lexicon, where each lexicon consists of a ground-truth word and a set of randomly picked words.

SVT [64] is collected from the Google Street View images that were used for scene text detection research. 647 words images are cropped from 249 street view images and most cropped texts are almost horizontal.

SVTP [43] has 639 word images that are cropped from the SVT images. Most images in this dataset suffer from perspective distortion which are purposely selected for evaluation of scene text recognition under perspective views.

CUTE [46] has 288 word images most of which are curved. All words are cropped from the CUTE dataset which contains 80 scene text images that are originally collected for scene text detection research.

4.2. Scene Text Detection

The proposed GA-DAN is evaluated by the performance of the scene text detectors that are trained by using its adapted images. In evaluations, the training set of ICDAR2013 (IC13) is used as the source dataset and the training sets of ICDAR2015 (IC15) and MSRA-TD500 (MT) are used as the target datasets which contain very different images as compared with those in IC13. GA-DAN generates two sets of images ‘AD-IC13’ and ‘10-AD-IC13’ for each of the two target datasets. The ‘AD-IC13’ is generated by 1-to-1 adaptation where each IC13 image is transformed to a single image that has similar geometry and appearance

Table 2. Scene text detection over the test images of the target datasets ICDAR2015 and MSRA-TD500: ‘IC13’, ‘Target’, ‘AD-IC13’ and ‘10-AD-IC13’ denote the dataset ICDAR2013, target dataset (ICDAR2015 or MSRA-TD500), 1-to-1 adapted ICDAR2013 and 1-to-10 adapted ICDAR2013, respectively. ‘SynthText’ refers to 800K synthetic images as reported in [13].

Method	ICDAR2015			MSRA-TD500		
	Recall	Precision	F-score	Recall	Precision	F-score
RRD [29] [SynthText + Target]	79.0	85.6	82.2	73.0	87.0	79.0
TextSnake [35] [SynthText + Target]	80.4	84.9	82.6	73.9	83.2	78.3
EAST [IC13]	43.7	68.2	53.3	34.9	71.2	46.8
EAST [AD-IC13]	59.6	69.9	64.4	51.5	67.7	58.5
EAST [10-AD-IC13]	71.6	67.3	69.4	55.8	69.9	62.1
EAST [Target]	76.9	81.1	79.0	64.4	73.8	68.7
EAST [IC13 + Target]	77.0	83.2	80.0	66.2	74.8	70.3
EAST [AD-IC13 + Target]	79.2	83.7	81.4	67.7	77.5	72.3
EAST [10-AD-IC13 + Target]	81.6	85.6	83.5	71.1	80.5	75.5

as the target dataset. The ‘10-AD-IC13’ is produced by 1-to-10 adaptation where each IC13 image is transformed to 10 adapted images by sampling 10 different spatial codes. Scene text detector EAST [78] is adopted for evaluation.

Table 2 shows quantitative results on the test set of two target datasets. Seven EAST models are trained for each target dataset by using different training images including 1) [IC13]: the training set of IC13, 2) [AD-IC13]: the 1-to-1 adapted IC13, 3) [10-AD-IC13]: the 1-to-10 adapted IC13, 4) [Target]: the training set of each target dataset, 5) [IC13 + Target]: the combination of the IC13 training set and the training set of each target dataset, 6) [AD-IC13 + Target]: the combination of AD-IC13 and the training set of each target dataset, and 7) [10-AD-IC13 + Target]: the combination of 10-AD-IC13 and the training set of each target dataset.

As Table 2 shows, the effectiveness of GA-DAN adapted images can be observed from three aspects. First, **EAST [AD-IC13]** outperforms **EAST [IC13]** by f-scores of 11.1% (53.3% → 64.4%) and 11.7% (46.8% → 58.5%) on the target datasets IC15 and MT, respectively, demonstrating the effectiveness of GA-DAN in adapting images from IC13 to IC15 and MT. Second, **EAST [10-AD-IC13]** improves **EAST [AD-IC13]** by f-scores of 5% (64.4% → 69.4%) and 3.6% (58.5% → 62.1%) on IC15 and MT, respectively. This shows the effectiveness of the multi-modal spatial learning that transforms a source-domain image to multiple target-domain images that are complementary with different spatial views. Third, **EAST [10-AD-IC13+Target]** improves **EAST [IC13+Target]** by a f-score of 3.5% (80.0% → 83.5%) and 5.2% (70.3% → 75.5%) on IC15 and MT, respectively. This shows that the adapted im-

ages are clearly more useful when combined with the training images of the target datasets for model training.

In addition, **EAST [10-AD-IC13+Target]** achieves state-of-the-art performance on the dataset IC15 by including only 2.3K GA-DAN adapted images (from 230 IC13 training images). As a comparison, TextSnake and RRD use 800K synthetic images in ‘SynthText’ [20] and they are also more advanced scene text detectors. Though the ‘10-AD-IC13’ is much smaller than SynthText, it contributes more to the detection improvement largely because of the large domain shifts between SynthText and IC15. For the target dataset MT, the f-score of **EAST [10-AD-IC13+Target]** is slightly lower than that of state-of-the-art detectors TextSnake and RRD, largely because the domain shifts between MT and SynthText are relatively small and the much larger amount of images in SynthText help more on the performance improvement. We believe higher f-score can be achieved when a higher number of GA-DAN adapted images are included in model training.

Table 3 shows the detection performance of different domain adaptation methods when EAST are trained by using their adapted images from IC13 to IC15 (the Baseline is trained using the original IC13 training images). Note for CycleGAN we adopt patch-wise training to minimize the effect of geometry differences in adversarial training. As ST-GAN is originally for image composition, we adapt it to achieve image translation in geometry space and restrict the transformation parameters to avoid boundary losing in testing phase. As Table 3 shows, all three adaptation models GA-DAN, CycleGAN and ST-GAN outperform Baseline clearly, and GA-DAN achieves clearly better f-score



Figure 4. Comparing our GA-DAN with state-of-the-art adaptation methods: The first and last columns show source-domain (IC13) and target-domain (IC15) images. GA-DAN_1, GA-DAN_2 and GA-DAN_3 show three GA-DAN adapted images of different spatial views.

(64.4% vs. 57.2% and 57.6%), demonstrating its superiority in adapting more realistic images. We also evaluate a new model ST-GAN + CycleGAN that directly concatenates ST-GAN and CycleGAN for adaptation in both geometry and appearance spaces. It shows that our GA-DAN still performs better by a large margin in f-score (64.4% vs. 60.8%), demonstrating its advantages in concurrent learning of geometric and appearance features.

Fig. 4 compares our GA-DAN with several state-of-the-art image adaptation methods. As Fig. 4 shows, GA-DAN adapts in both appearance and geometric spaces realistically, whereas SimGAN and CycleGAN can only adapt appearance features and ST-GAN can only adapt geometric features. In addition, GA-DAN_1, GA-DAN_2 and GA-DAN_3 show three GA-DAN adapted images with different spatial views, demonstrating the effectiveness of our proposed multi-modal spatial learning.

Table 3. Scene text detection on the IC15 test images: The detection models are trained using the adapted IC13 training images (from IC13 to IC15) by different adaptation methods as listed. (Baseline is trained by using the original IC13 training images)

Method	Recall	Precision	F-score
CycleGAN [79]	50.3	66.3	57.2
ST-GAN [30]	52.9	63.4	57.6
ST-GAN + CycleGAN	57.3	64.7	60.8
Baseline	43.7	68.2	53.3
GA-DAN	59.6	69.9	64.4

4.3. Scene Text Recognition

For scene text recognition experiment, we select the CUTE [46] and SVTP [43] as the target datasets. As current scene text recognition datasets are all too small, we combine all images from datasets IC13 [24], IIIT5K [41] and SVT [64] as the source dataset denoted by ‘COMB’. As scene texts in CUTE and SVTP are most curved or in perspective views but most COMB texts are horizontal, we use the thin spline for spatial transformation which is flexible for various spatial transformations.

Table 4 shows recognition accuracy when COMB images are adapted by different adaptation methods and then used to train the scene text recognition model: MORAN [37]. As Table 4 shows, **GA-DAN [WM]** (GA-DAN with 1-to-1 spatial transformation) outperforms other adaptation methods with a large margin. Additionally, most compared adaptation methods do not show clear improvement over the Baseline (trained by using the original COMB images without adaptation). In particular, CycleGAN and CyCADA improve the accuracy by 1.0% and 1.3% only for CUTE because they only adapt in appearance space but the main discrepancy between COMB and CUTE is in geometry space. CoGAN and UNIT tend to over-adapt the text appearance which may even change the text semantics and make texts unrecognizable.

Table 4 also shows the ablation study results. Two GA-DAN models are trained for image adaptation. The first model is a complete GA-DAN with all newly designed features and components includes. The second is **GA-DAN [WD]** which is trained with a normal instead of disentangled cycle-consistency loss. For fair comparison, the region missing loss is also included in **GA-DAN [WD]**. For the complete GA-DAN, three sets of adapted images are gen-



Figure 5. Comparing GA-DAN with state-of-the-art adaptation methods: Rows 1-2 show adaptation from COMB to CUTE, Rows 3-4 show adaptation from COMB to SVTP. GA-DAN_1, GA-DAN_2 and GA-DAN_3 show GA-DAN adapted images of different spatial views.

Table 4. Ablation study and comparisons with state-of-the-art adaptation methods: Recognition models are trained by different adaptations of the COMB images to the target domains CUTE and SVTP (Baseline uses the original COMB images and ‘Random’ applies random spatial transformation in adaptation).

Methods	COMB-CUTE	COMB-SVTP
SimGAN [40]	30.7	42.6
UNIT [31]	28.7	40.8
CoGAN [68]	28.3	40.2
DualGAN [40]	31.5	42.7
CycleGAN [79]	31.9	43.0
CyCADA [16]	32.2	43.6
Baseline	30.9	42.5
Random	28.8	42.7
GA-DAN [WD]	32.6	44.9
GA-DAN [WA]	36.1	45.2
GA-DAN [WM]	38.2	47.1
GA-DAN [10 AD]	43.1	51.7

erated to train the recognition model. The first set is **GA-DAN [WA]** that just takes the output of S_X without appearance adaptation as shown in Fig. 2. The second set is **GA-DAN [WM]** that performs 1-to-1 adaptation and transforms each source-domain image into a single target-domain image. The third set is **GA-DAN [10 AD]** that performs 1-to-10 adaptation and transforms each source-domain image into 10 target-domain images. As Table 4 shows, **GA-DAN**

[WA] clearly outperforms Baseline and ‘Random’ (adapted using a random spatial transformation matrix) as well as state-of-the-art adaptation methods, showing the superiority of our spatial module in learning correct and accurate spatial transformations. **GA-DAN [WD]** outperforms state-of-the-art methods but clearly performs worse than **GA-DAN [WM]**, demonstrating the effectiveness of the proposed disentangled cycle-consistency loss. **GA-DAN [10 AD]** outperforms **GA-DAN [WM]** clearly, demonstrating the effectiveness of our proposed multi-modal spatial learning.

Fig. 5 compares our GA-DAN with several state-of-the-art adaptation methods. As Fig. 5 shows, GA-DAN adapts in both appearance and geometry spaces realistically whereas CycleGAN and SimGAN can only adapt in appearance space. In addition, GA-DAN_1, GA-DAN_2 and GA-DAN_3 show that the proposed GA-DAN is capable of transforming a source-domain image to multiple target-domain images of different spatial views.

5. Conclusions

This paper presents a geometry-aware domain adaptation network that achieves domain adaptation in geometry and appearance spaces simultaneously. A multi-modal spatial learning technique is proposed which can generate multiple adapted images with different spatial views. A novel disentangle cycle-consistency loss is designed which greatly improves the stability and concurrent learning in both geometry and appearance spaces. The proposed network has been validated over scene text detection and recognition tasks and experiments show the superiority of the adapted images while applied to train deep networks.

6. Acknowledgement

This work is funded by the Ministry of Education (MOE), Singapore, under the project ‘‘A semi-supervised learning approach for accurate and robust detection of texts in scenes’’ (RG128/17 (S)).

References

- [1] Jon Almazan, Albert Gordo, Alicia Forns, and Ernest Valveny. Word spotting and recognition with embedded attributes. *TPAMI*, 36(12):2552–2566, 2014.
- [2] Fred L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *TPAMI*, 11(6), 1989.
- [3] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *CVPR*, 2017.
- [4] Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. In *NIPS*, 2016.
- [5] Zhazhan Cheng, Fan Bai, Yunlu Xu, Gang Zheng, Shiliang Pu, and Shuigeng Zhou. Focusing attention: Towards accurate text recognition in natural images. In *ICCV*, pages 5076–5084, 2017.
- [6] Chee Kheng Chng and Chee Seng Chan. Total-text: A comprehensive dataset for scene text detection and recognition. In *ICDAR*, pages 935–942, 2017.
- [7] Dan Deng, Haifeng Liu, Xuelong Li, and Deng Cai. Pixellink: Detecting scene text via instance segmentation. In *AAAI*, 2018.
- [8] Emily Denton, Soumith Chintala, Arthur Szlam, and Rob Fergus. Deep generative image models using a laplacian pyramid of adversarial networks. In *NIPS*, 2015.
- [9] Vincent Dumoulin, Ishmael Belghazi, Ben Poole, Alex Lamb, Martin Arjovsky, Olivier Mastropietro, and Aaron Courville. Adversarially learned inference. *arXiv:1606.00704*, 2016.
- [10] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 325–333, 2015.
- [11] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. In *NIPS*, pages 2672–2680, 2014.
- [12] Albert Gordo. Supervised mid-level features for word image representation. In *CVPR*, 2015.
- [13] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. Synthetic data for text localisation in natural images. In *CVPR*, 2016.
- [14] Pan He, Weilin Huang, Tong He, Qile Zhu, Yu Qiao, and Xiaolin Li. Single shot text detector with regional attention. *arXiv:1709.00138*, 2017.
- [15] Tong He, Weilin Huang, Yu Qiao, and Jian Yao. Text-attentional convolutional neural network for scene text detection. *TIP*, 25(6):2529–2541, 2016.
- [16] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Alexei A. Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *ICML*, 2018.
- [17] Han Hu, Chengquan Zhang, Yuxuan Luo, Yuzhuo Wang, Junyu Han, and Errui Ding. Wordsup: Exploiting word annotations for character based text detection. In *ICCV*, Oct 2017.
- [18] Weilin Huang, Yu Qiao, and Xiaoou Tang. Robust scene text detection with convolution neural network induced mser trees. In *ECCV*, pages 497–511, 2014.
- [19] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. In *CVPR*, 2017.
- [20] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Synthetic data and artificial neural networks for natural scene text recognition. *arXiv preprint arXiv:1406.2227*, 2014.
- [21] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep structured output learning for unconstrained text recognition. In *ICLR*, 2015.
- [22] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Reading text in the wild with convolutional neural networks. *IJCV*, 116(1):1–20, 2016.
- [23] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Deep features for text spotting. In *ECCV*, pages 512–528, 2014.
- [24] Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluís Gomez i Bigorda, Sergi Robles Mestre, Joan Mas, David Fernandez Mota, Jon Almazan Almazan, and Lluís Pere de las Heras. Icdar 2013 robust reading competition. In *ICDAR*, pages 1484–1493, 2013.
- [25] Dimosthenis Karatzas¹, Lluís Gomez-Bigorda¹, Angelos Nicolaou¹, Suman Ghosh¹, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, Faisal Shafait, Seiichi Uchida, and Ernest Valveny. Icdar 2015 competition on robust reading. In *ICDAR*, pages 1156–1160, 2015.
- [26] Taeksoo Kim, Moonsu Cha, Hyunsoo Kim, Jung Kwon Lee, and Jiwon Kim. Learning to discover cross-domain relations with generative adversarial networks. In *ICML*, 2017.
- [27] Chen-Yu Lee and Simon Osindero. Recursive recurrent nets with attention modeling for ocr in the wild. In *CVPR*, pages 2231–2239, 2016.
- [28] Minghui Liao, Baoguang Shi, Xiang Bai, Xinggang Wang, and Wenyu Liu. Textboxes: A fast text detector with a single deep neural network. In *AAAI*, pages 4161–4167, 2017.
- [29] Minghui Liao, Zhen Zhu, Baoguang Shi, Guisong Xia, and Xiang Bai. Rotation-sensitive regression for oriented scene text detection. In *CVPR*, pages 5909–5918, 2018.
- [30] Chen-Hsuan Lin, Ersin Yumer, Oliver Wang, Eli Shechtman, and Simon Lucey. St-gan: Spatial transformer generative adversarial networks for image compositing. In *CVPR*, 2018.
- [31] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NIPS*, 2017.
- [32] Xuebo Liu, Ding Liang, Shi Yan, Dagui Chen, Yu Qiao, and Junjie Yan. Fots: Fast oriented text spotting with a unified network. In *CVPR*, pages 5676–5685, 2018.
- [33] Yuliang Liu and Lianwen Jin. Deep matching prior network: Toward tighter multi-oriented text detection. In *CVPR*, 2017.
- [34] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I. Jordan. Deep transfer learning with joint adaptation networks. In *ICML*, 2017.
- [35] Shangbang Long, Jiaqiang Ruan, Wenjie Zhang, Xin He, Wenhao Wu, and Cong Yao. Textsnake: A flexible repre-

- sentation for detecting text of arbitrary shapes. In *ECCV*, pages 20–36, 2018.
- [36] Shijian Lu, Tao Chen, Shangxuan Tian, Joo-Hwee Lim, and Chew-Lim Tan. Scene text extraction based on edges and support vector regression. *IJDAR*, 18(2):125–135, 2015.
- [37] Canjie Luo, Lianwen Jin, and Zenghui Sun. Moran: A multi-object rectified attention network for scene text recognition. In *Pattern Recognition*, volume 90, pages 109–118, 2019.
- [38] 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 *ECCV*, 2018.
- [39] Pengyuan Lyu, Cong Yao, Wenhao Wu, Shuicheng Yan, and Xiang Bai. Multi-oriented scene text detection via corner localization and region segmentation. In *CVPR*, pages 7553–7563, 2018.
- [40] Oncel Tuzel Ming-Yu Liu. Coupled generative adversarial networks. In *NIPS*, 2016.
- [41] Anand Mishra, Karteek Alahari, and C.V. Jawahar. Scene text recognition using higher order language priors. In *BMVC*, 2012.
- [42] Lukas Neumann and Jiri Matas. Real-time scene text localization and recognition. In *CVPR*, pages 3538–3545, 2012.
- [43] Trung Quy Phan, Palaiahnakote Shivakumara, Shangxuan Tian, and Chew Lim Tan. Recognizing text with perspective distortion in natural scenes. In *ICCV*, 2013.
- [44] Andrei Polzounov, Artsiom Ablavatski, Sergio Escalera, Shijian Lu, and Jianfei Cai. Wordfence: Text detection in natural images with border awareness. In *ICIP*, pages 1222–1226. IEEE, 2017.
- [45] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, 2016.
- [46] Anhar Risnumawan, Palaiahankote Shivakumara, Chee Seng Chan, and Chew Lim Tan. A robust arbitrary text detection system for natural scene images. *Expert Syst. Appl.*, 41(18):8027–8048, 2014.
- [47] Jose A. Rodriguez-Serrano, Albert Gordo, and Florent Perronnin. Label embedding: A frugal baseline for text recognition. *IJCV*, 2015.
- [48] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *ECCV*, pages 325–333, 2010.
- [49] Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *TPAMI*, 39(11):2298–2304, 2017.
- [50] Baoguang Shi, Xinggang Wang, Pengyuan Lyu, Cong Yao, and Xiang Bai. Robust scene text recognition with automatic rectification. In *CVPR*, 2016.
- [51] Baoguang Shi, Cong Yao, Minghui Liao, Mingkun Yang, Pei Xu, Linyan Cui, Serge Belongie, Shijian Lu, and Xiang Bai. Icdar2017 competition on reading chinese text in the wild (rctw-17). In *ICDAR*, volume 01, pages 1429–1434, 2017.
- [52] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Josh Susskind, Wenda Wang, and Russ Webb. Learning from simulated and unsupervised images through adversarial training. In *CVPR*, 2017.
- [53] Bolan Su and Shijian Lu. Accurate scene text recognition based on recurrent neural network. In *ACCV*, 2014.
- [54] Bolan Su and Shijian Lu. Accurate recognition of words in scenes without character segmentation using recurrent neural network. *PR*, 2017.
- [55] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *AAAI*, 2016.
- [56] Baochen Sun and Kate Saenko. Deep coral: correlation alignment for deep domain adaptation. In *ICCV workshop*, 2016.
- [57] Yaniv Taigman, Adam Polyak, and Lior Wolf. Unsupervised cross-domain image generation. In *ICLR*, 2017.
- [58] Shangxuan Tian, Ujjwal Bhattacharya, Shijian Lu, Bolan Su, Qingqing Wang, Xiaohua Wei, Yue Lu, and Chew Lim Tan. Multilingual scene character recognition with cooccurrence of histogram of oriented gradients. *PR*, pages 125–134, 2016.
- [59] Shangxuan Tian, Yifeng Pan, Chang Huang, Shijian Lu, Kai Yu, and Chew Lim Tan. Text flow: A unified text detection system in natural scene images. In *ICCV*, pages 4651–4659, 2015.
- [60] Zhi Tian, Weilin Huang, Tong He, Pan He, and Yu Qiao. Detecting text in natural image with connectionist text proposal network. In *ECCV*, pages 56–72, 2016.
- [61] Antonio Torralba and Alexei A. Efros. Unbiased look at dataset bias. In *CVPR*, 2011.
- [62] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *CVPR*, 2017.
- [63] Fangfang Wang, Liming Zhao, Xi Li, Xinchao Wang, and Dacheng Tao. Geometry-aware scene text detection with instance transformation network. In *CVPR*, June 2018.
- [64] Kai Wang, Boris Babenko, and Serge Belongie. End-to-end scene text recognition. In *ICCV*, 2011.
- [65] Chuhui Xue, Shijian Lu, and Fangneng Zhan. Accurate scene text detection through border semantics awareness and bootstrapping. In *ECCV*, pages 370–387, 2018.
- [66] Cong Yao, Xiang Bai, Wenyu Liu, Yi Ma, and Zhuowen Tu. Detecting texts of arbitrary orientations in natural images. In *CVPR*, 2012.
- [67] Cong Yao, Xiang Bai, Baoguang Shi, and Wenyu Liu. Strokelets: A learned multi-scale representation for scene text recognition. In *CVPR*, 2014.
- [68] Zili Yi, Hao Zhang, Ping Tan, and Minglun Gong. Dualgan: Unsupervised dual learning for image-to-image translation. In *ICCV*, 2017.
- [69] Xu-Cheng Yin, Wei-Yi Pei, Jun Zhang, and Hong-Wei Hao. Multiorientation scene text detection with adaptive clustering. *TPAMI*, 37(9):1930–1937, 2015.
- [70] Fangneng Zhan, Jiaying Huang, and Shijian Lu. Adaptive composition gan towards realistic image synthesis. *arXiv preprint arXiv:1905.04693*, 2019.
- [71] Fangneng Zhan and Shijian Lu. Esir: End-to-end scene text recognition via iterative image rectification. In *CVPR*, pages 2059–2068, 2019.
- [72] Fangneng Zhan, Shijian Lu, and Chuhui Xue. Verisimilar image synthesis for accurate detection and recognition of texts in scenes. In *ECCV*, pages 257–273, 2018.

- [73] Fangneng Zhan, Hongyuan Zhu, and Shijian Lu. Scene text synthesis for efficient and effective deep network training. *arXiv preprint arXiv:1901.09193*, 2019.
- [74] Fangneng Zhan, Hongyuan Zhu, and Shijian Lu. Spatial fusion gan for image synthesis. In *CVPR*, pages 3653–3662, 2019.
- [75] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *ICCV*, 2017.
- [76] Zheng Zhang, Wei Shen, Cong Yao, and Xiang Bai. Symmetry-based text line detection in natural scenes. In *CVPR*, pages 2558–2567, 2015.
- [77] Zheng Zhang, Chengquan Zhang, Wei Shen, Cong Yao, Wenyu Liu, and Xiang Bai. Multi-oriented text detection with fully convolutional networks. In *CVPR*, pages 4159–4167, 2016.
- [78] Xinyu Zhou, Cong Yao, He Wen, Yuzhi Wang, Shuchang Zhou, Weiran He, and Jiajun Liang. East: An efficient and accurate scene text detector. In *CVPR*, 2017.
- [79] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017.