TextAdaIN: Paying Attention to Shortcut Learning in Text Recognizers

Oren Nuriel, Sharon Fogel, and Ron Litman

AWS AI Labs {onuriel, shafog, litmanr}@amazon.com

Abstract. Leveraging the characteristics of convolutional layers, neural networks are extremely effective for pattern recognition tasks. However in some cases, their decisions are based on unintended information leading to high performance on standard benchmarks but also to a lack of generalization to challenging testing conditions and unintuitive failures. Recent work has termed this "shortcut learning" and addressed its presence in multiple domains. In text recognition, we reveal another such shortcut, whereby recognizers overly depend on local image statistics. Motivated by this, we suggest an approach to regulate the reliance on local statistics that improves text recognition performance.

Our method, termed TextAdaIN, creates local distortions in the feature map which prevent the network from overfitting to local statistics. It does so by viewing each feature map as a sequence of elements and deliberately mismatching fine-grained feature statistics between elements in a mini-batch. Despite TextAdaIN's simplicity, extensive experiments show its effectiveness compared to other, more complicated methods. TextAdaIN achieves state-of-the-art results on standard handwritten text recognition benchmarks. It generalizes to multiple architectures and to the domain of scene text recognition. Furthermore, we demonstrate that integrating TextAdaIN improves robustness towards more challenging testing conditions.

Keywords: Text Recognition, Handwriting Recognition, Scene Text Recognition, Shortcut Learning, Regularization

1 Introduction

Reading someone else's handwriting is often a challenging task; some of the characters are unclear, the text is cursive, there is background clutter and the image quality can be low. When deciphering each character, we often rely on the surrounding area to compensate for the occasional obscurity of the text. The automation of reading text images has been a thriving field of research in computer vision for decades. Recent deep learning methods have significantly improved recognition results [6, 67, 42, 38, 2, 1].

However, previous works suggest that despite their super-human capabilities, deep learning methods are limited by their tendency to err even when introducing small (in some cases even invisible) modifications to the input. Numerous works have touched on this subject from various angles [10, 50, 59, 58, 23, 45] and propose tailored solutions. Geirhos *et al.* [22] view each of these as a symptom of the same underlying problem, and terms this phenomenon *shortcut learning - decision rules that perform well on*



Fig. 1: **Decoder attention maps.** Each example shows the input image (bottom), attention map (top) and model prediction (left). Each line in the attention map is a time step representing the attention per character prediction. (a) The baseline model, which uses local statistics as a shortcut, misinterprets the corrupted images. (b) Our proposed method which overcomes this shortcut, enhances performance on both standard and challenging testing conditions

independent and identically distributed (i.i.d.) test data but fail on tests that are outof-distribution (o.o.d). Shortcut learning occurs when there is an easy way to learn attributes which are highly correlated with the label (i.e., "it stands on a grass lawn \implies it is a cow"). Thus, shortcuts are inherent in the data [23, 28]. Yet, as they can depend on other various factors such as the architecture [15], and the optimization procedure [64, 16] as well, revealing an instance of shortcut learning is non-trivial task. The fact that a model learned a shortcut can be discovered at test time when it encounters examples that the shortcut heuristic fails on (i.e., a cow on the beach).

In this work, we reveal a shortcut pertaining to text recognizers, specifically, we reveal the unhealthy reliance of text recognizers on local statistics. Based on the observation that text recognizers operate on a local level, we hypothesize that they are susceptible to overly rely on local information which may even be indistinguishable to the naked eye (*i.e.*, a certain level of curvature is unique and highly correlated with the character "f"). As such in this case it is difficult to identify the exact shortcut heuristic, thus we provide intuition through quantitative and qualitative analyses. Fig. 1(a)illustrates the decoder's attention maps for a state-of-the-art [38] recognizer before and after applying local corruptions to the image. In the first row, additive Gaussian noise is applied, which distorts local information while maintaining semantic information. As a consequence, the model diverges completely. In the second row, motion blur is applied, artificially imitating object motion. In this case, in order to successfully decode the corrupted image, global context is necessary (imagine reading the blurred image character by character without access to the surrounding characters). As one can tell from the attention maps, the model fails to do so and is unable to correctly decode the image. An in-depth analysis reveals that this phenomenon is apparent in standard text recognizers across both scene text and handwritten domains, as further elaborated in Section 3.

To prevent the model from using the aforementioned shortcut, we propose a simple, yet powerful, technique for moderating the reliance on local statistics. The key idea of our approach is to probabilistically swap fine-grained feature statistics in a manner that



Fig. 2: **TextAdaIN**. TextAdaIN views each of the feature maps as a sequence of individual elements and swaps feature statistics between elements instead of entire images as in AdaIN. This feature-level distortion alleviates the model's tendency to rely on shortcuts in the form of local statistics

is adjusted to the task at hand. To that end, we propose TextAdaIN, a local variation of AdaIN [29], depicted in Fig. 2. In contrast to AdaIN which operates on the entire image, TextAdaIN splits the feature map into a sequence of elements and operates on each element independently. Furthermore, the normalization is performed over multiple dimensions. Both modifications increase the granularity level in which the statistics are modified and enable the usage of multiple donor images. Effectively, the representation space undergoes distortions derived from an induced distribution, namely other text images, at a sub-word level. Thus, forcing the encoder to account for information in surrounding areas as the local information cannot always be depended on. This is observed in Fig. 1(b), where in contrast to the baseline model, TextAdaIN successfully utilizes global context and properly decodes the images despite the corruptions.

We validate our method by comparing its performance with state-of-the-art approaches on several handwritten text benchmarks. TextAdaIN achieves state-of-the-art results, reducing the word error rate by **11.8%** and **16.4%** on IAM and RIMES, respectively. Furthermore, our method shows a consistent improvement across multiple architectures and in the domain of scene text recognition. Not only does our model surpass other, more complicated methods ([9, 20, 67, 42, 2]), but it is simple to implement and can effortlessly be integrated into any mini-batch training procedure. To summarize, the key contributions of our work are:

- 1. We reveal an instance of shortcut learning in text recognizers in the form of a heavy dependence on local statistics and suggest to regulate it.
- We introduce TextAdaIN, a simple yet effective normalization approach to remediate the reliance on local statistics in text recognizers.
- 3. Extensive experimental validation shows our method achieves state-of-the-art results on several popular handwritten text benchmarks. In addition, it is applicable to the domain of scene text and can be used independent of the chosen architecture.
- We demonstrate that our method leads to improved performance on challenging testing conditions.

2 Related Work

Shortcut learning. An extensive amount of research has been conducted at understanding the behaviour of neural networks. One behavioural aspect under investigation is their unintended solutions, whereby a decision is made based on misleading information, thus limiting the performance under general conditions. A typical example of this is an intriguing discovery made by Szegedy *et al.* [58]. They showed the susceptibility of neural networks to adversarial examples, which are samples that undergo minimal, or even unnoticeable, modifications capable of altering model predictions. More recent works show other ill-desired sensitivities, [23] for texture, [3] for color constancy, [45, 49] for global statistics and [8] for image background. These phenomena were classified by Geirhos *et al.* [22] as instances of the same underlying problem that they termed *shortcut learning*. So far, for shortcut learning in the domain of text recognition, we are only aware of [61], which exposed the decoder's high reliance on vocabulary.

Normalization and style transfer. Normalizing feature tensors is an effective and powerful technique that has been explored and developed over the years for a variety of tasks. Ioffe & Szegedy [30] were the first to introduce Batch Normalization, which inspired a series of normalization-based methods such as Layer Normalization [5], Instance Normalization [60], and Group Normalization [65].

The style of an image was characterized by [21] to be the statistics of activation maps in intermediate layers of convolutional neural networks. Instance Normalization (IN) [60] is a normalization layer which normalizes these statistics, removing the style representation from the activation map. Subsequently, Adaptive Instance Normalization (AdaIN) was proposed by [29] for real-time arbitrary style transfer. AdaIN changes the style of an image by incorporating statistics from an additional image.

Leveraging the benefits of this operation, Geirhos *et al.* [23] created Stylized-ImageNet, a stylized version of ImageNet. They demonstrated that classifiers trained on this dataset rely more on shape than on texture. More recently, Zhou *et al.* [68] proposed to probabilistically mix instance-level feature statistics during training across different domains, thus increasing the generalizability of the model. Furthermore, Nuriel *et al.* [45] demonstrated that a similar approach can reduce classifiers' reliance on global statistics, therefore increasing classification performance.

In this work we reveal a shortcut learning phenomenon in text recognizers. To remediate this problem, we also leverage AdaIN. However, instead of operating on a word (instance) level, we swap fine-grained statistics on a sub-word level. We do so by treating every few consecutive frames in the sequential feature map as individual elements.

Text Recognition. Text recognition has attracted considerable attention over the past few years. In particular, deep learning approaches have achieved remarkable results [51, 6, 38, 56, 18, 19, 1]. Still, current state-of-the-art methods struggle to train robust models when the amount of data is insufficient to capture the magnitude of styles. Various methods have been suggested to cope with this problem. Bhunia *et al.* [9] proposed an adversarial feature deformation module that learns ways to elastically warp extracted features, boosting the network's capability to learn highly informative features. Luo *et al.* [42] introduced an agent network that learns from the output of the recognizer. The

agent controls a set of fiducial points on the image and uses a spatial transformer to generate "harder" training samples.

Unlike previous methods, our method does not require additional data, new models or complex training paradigms. Instead, we suggest a normalization-based method adjusted to text images and to sequence-to-sequence approaches. TextAdaIN is extremely easy to implement and can fit into any encoder as part of a mini-batch training process.

Throughout this work, unless mentioned otherwise, we integrate our proposed method with a state-of-the-art recognizer named SCATTER [38]. The SCATTER architecture consists of four main components: spatial transformation, feature extraction, visual feature refinement and a selective-contextual refinement block. For further details about the baseline architecture, we refer the reader to Appendix A.

3 Method

3.1 Shortcut Learning in Text Recognizers

As elaborated on in [22], shortcuts are dependent on various factors, including the architecture of the model and the data. We therefore draw our motivation for suspecting local statistics as a potential instance of shortcut learning in text recognizers based on two key observations pertaining to both the architecture and the nature of the data. First, the characters appear across a series of frames along the width axis of the image. Therefore, text recognition approaches define the recognition task as

Table 1: **Text recognizers' performance in challenging testing conditions.** Mean accuracies of different state-of-the-art text recognizers while applying a series of subtle local image corruptions. The corruptions consist of local masking and pixel-wise distortions: Cutout, Dropout, Additive Guassian Noise, Elastic Transform and Motion Blur. Models are especially susceptible to learning shortcuts when the training data is limited, as in handwriting

Method	Scen	e Text	Handwritten
	Regular text	Irregular text	IAM
Baek et al. (CTC) [6]	88.7	72.9	80.6
+Local Corruption	69.8 (-18.9)	44.1 (-28.8)	40.4 (-40.2)
Baek et al. (Attn) [6]	92.0	77.4	82.7
+Local Corruption	74.5 (-17.5)	50.5 (-26.9)	46.2 (-36.5)
SCATTER [38]	93.5	82.1	85.7
+Local Corruption	76.8 (-16.7)	55.3 (-26.8)	54.8 (-30.9)

a sequence character classification problem. This leads us to the second observation, the model's predictions are mostly based on local information which lies in consecutive frames [52, 46]. Local information is referred to, in our work, as local feature statistics.

As illustrated in Figure 1, when applying slight local modifications to the image distorting local information in a way which does not affect human reading capabilities (see Appendix B), the baseline model is unable to predict the words correctly. Table 1 quantitatively exemplifies this phenomenon, whereby the models fail to generalize to images under challenging testing conditions in the form of local corruptions. The table shows the performance of three different text recognition architectures - Baek et al. [6] with CTC and attention heads and SCATTER [38]. All models have been tested on

the original scene-text and IAM datasets and on corrupted versions of these datasets. The corruptions used are the local-masking and pixel-wise distortions described in Table 6. All of the models suffer from a large decline in performance when adding the corruptions, and especially on handwriting where the training data is scarce.

An additional experiment validating text recognizers' sensitivity to local distortions is shown in Section 5.5, where we empirically demonstrate that text recognizers heavily rely on local information. Specifically, we show that profusely distorting global statistics has little to no effect on performance (unlike in classification models [45]). However, intensely distorting local statistics substantially degrades performance, which leads to the understanding that text recognizers do indeed rely on such information. The local statistics are an important cue, although they are not a determining factor. Thus, to regulate this ill-desired model bias, we propose TextAdaIN, a sequence-aware local variation of AdaIN.

3.2 AdaIN

To define AdaIN, we first begin with formally describing Instance Normalization (IN) [60]. Given an instance $x \in \mathbb{R}^{C \times H \times W}$, where C, H and W are the channels, height, and width respectively, IN is defined to be:

$$IN(x) = \gamma \left(\frac{x - \mu(x)}{\sigma(x)}\right) + \beta.$$

Where $\gamma, \beta \in \mathbb{R}^C$ are learned parameters and $\mu(x), \sigma(x) \in \mathbb{R}^C$ are calculated by taking the mean and standard deviation over H, W of x.

Adaptive Instance Normalization (AdaIN), proposed in [29], is built upon IN. Given two image representations, $x_a, x_b \in \mathbb{R}^{C \times H \times W}$, AdaIN shifts the statistics of the representation of x_a to the representation of x_b . This is done in two steps. First, Instance Normalization is applied on x_a to remove x_a 's style information.



Fig. 3: **TextAdaIN block.** (a) A standard residual block and (b) a TextAdaIN block. TextAdaIN is probabilistically employed after every Conv layer during training. *split* and *merge* refer to the mapping operations of reshaping each feature map and $\mu_{c,h}$, $\sigma_{c,h}$ refer to the feature statistics, as described in Section 3.3

Then, the normalized activation map is scaled and shifted to match x_b 's statistics. This operation is perceived to transfer x_b 's style onto x_a .

$$AdaIN_{c}(x_{a}, x_{b}) = \sigma(x_{b}) \left(\frac{x_{a} - \mu(x_{a})}{\sigma(x_{a})}\right) + \mu(x_{b}).$$
⁽²⁾

AdaIN_c denotes the standard AdaIN operation in which σ , μ are calculated over the spatial dimensions, resulting in shifting corresponding channel statistics.

3.3 TextAdaIN

We wish to design a method that produces model decision rules which are not overly reliant on local information. To this end, we propose a method which distorts local feature statistics during training. Leveraging the recent development proposed by [45, 68], we suggest two modifications to the vanilla AdaIN operation that correspond to the aforementioned observations:

- 1. Viewing the feature map as a sequence of individual elements and swapping the feature statistics between elements instead of entire images.
- 2. Modifying AdaIN to operate on two dimensions the height and channels of the feature map.

These modifications increase the granularity level in which statistics are calculated and modified, thus regulating the reliance on local statistics. In addition, the sequential view enables the utilization of multiple donor images. Both are crucial for our method's success, as shown in Section 5.

Formally, given an image representation $x \in \mathbb{R}^{C \times H \times W}$, we define $\mu_{c,h}$, $\sigma_{c,h}$ to be the following:

$$\mu_{c,h}(x) = \frac{1}{W} \sum_{w=1}^{W} x_{c,h,w} , \qquad (3)$$

$$\sigma_{c,h}(x) = \sqrt{\frac{1}{W} \sum_{w=1}^{W} (x_{c,h,w} - \mu_{c,h}(x))^2 + \epsilon} .$$
(4)

Therefore, a local variation of AdaIN_c can be defined as:

$$AdaIN_{c,h}(x_a, x_b) = \sigma_{c,h}(x_b) \left(\frac{x_a - \mu_{c,h}(x_a)}{\sigma_{c,h}(x_a)}\right) + \mu_{c,h}(x_b).$$
(5)

This variant of AdaIN swaps statistics for every corresponding channel and height, thus impacting the feature map's statistics at a higher level of granularity. We note that backpropagating gradients only occur through $\mu_{c,h}(x_a)$, $\sigma_{c,h}(x_a)$. The reason behind this is to avoid gradient flow from input image labels to the donor images. Given the definitions above, we formulate the TextAdaIN operation used during training. Let $X = \{x_i\}_{i=1}^B$ denote a mini-batch of *B* feature maps. We divide each sample x_i into *K* windows along its width. The result is a batch of elements pertaining to $B \cdot K$ windows. This operation can be defined as a mapping:

$$X \in \mathbb{R}^{B \times C \times H \times W} \to \widehat{X} \in \mathbb{R}^{B \cdot K \times C \times H \times \frac{W}{K}}.$$
(6)

Then, we employ a similar procedure to the one used in [45]. We randomly draw a permutation of the modified batch \hat{X} and apply AdaIN_{c,h} between the modified batch and the permuted one. Namely, let

$$\pi(\widehat{X}) = [\hat{x}_{\pi(1)}, \hat{x}_{\pi(2)} \dots \hat{x}_{\pi(B \cdot K)}]$$
(7)

unlabeled test for tra	aining	. Our	method	achie	eves s	tate-of-t	he-art	acros	ss all data
Mathad	IAM		RIMES			CVL			
Method	WER	CER	Average	WER	CER	Average	WER	CER	Average
Bluche et al. [11]	24.7	7.3	16.00	-	-	-	-	-	-
Bluche et al. [12]	24.6	7.9	16.25	-	-	-	-	-	-
Sueiras et al. [57]	23.8	8.8	16.30	-	-	-	-	-	-
Alonso et al. [4]	-	-	-	11.9	4.0	7.95	-	-	-
ScrabbleGAN [20]	23.6	-	-	11.3	-	-	22.9	-	-
SSDAN* [67]	22.2	8.5	15.35	-	-	-	-	-	-
Bhunia et al. [9]	17.2	8.4	12.80	10.5	6.4	8.45	-	-	-
Kang et al.* [33]	17.3	6.8	12.05	-	-	-	-	-	-
SeqCLR [2]	20.1	9.5	14.80	7.6	2.6	5.10	22.2	-	-
Luo et al. [42]	14.1	5.4	9.75	8.7	2.4	5.50	-	-	-
SCATTER [38]	14.4	6.4	10.40	6.7	2.3	4.50	22.3	18.9	20.6
+MixStyle [68]	14.3	6.3	10.30	6.8	2.4	4.60	22.3	18.5	20.40
+pAdaIN [45]	14.6	6.4	10.50	6.4	2.2	4.30	22.3	18.2	20.25
+TextAdaIN	12.7	5.8	9.25	5.6	1.9	3.70	21.8	18.0	19.90

Table 2: **Comparison to previous methods.** Word and character error rates (WER and CER) are measured on IAM, RIMES and CVL handwriting datasets. '*' indicates using the unlabeled test for training. Our method achieves state-of-the-art across all datasets

denote applying a permutation $\pi : [B \cdot K] \to [B \cdot K]$ on \widehat{X} . Then the output of TextAdaIN on the i^{th} window of \widehat{X} is defined by:

TextAdaIN
$$(\hat{x}_i) = \text{AdaIN}_{c,h}(\hat{x}_i, \hat{x}_{\pi(i)})$$
. (8)

Subsequently, the batch of windows is rearranged back to its original form using the inverse mapping operation.

TextAdaIN is applied batch-wise with probability p after every convolutional layer in the encoder, as illustrated in Fig. 3(b). The permutation π is sampled uniformly, and p is a hyperparameter fixed ahead of training. TextAdaIN is only applied during training and not during inference.

4 Experiments

In this section, we begin by comparing our method's performance against state-of-theart methods on several public handwriting datasets. Then, we demonstrate that TextAdaIN can be integrated into additional recognition architectures and applied to natural scene text images.

Datasets We conduct experiments on several public handwriting and scene text datasets. For handwriting, we consider the English datasets IAM [43] and CVL [36], and the French dataset RIMES [26]. For scene text, we train on the synthetic datasets Synth-Text [27] and MJSynth [31] and test on four real-world regular text datasets: IIT5K [44],

SVT [63], IC03 [41], IC13 [35] and three real-world irregular text datasets: ICDAR2015 [34], SVTP [47] and CUTE 80 [48]. We present samples from each dataset and include more details in Appendix C.

Metrics To evaluate recognition performance, word-level accuracy is measured. For handwritten text recognition state-of-the-art comparison, Word Error Rate (WER) and Character Error Rate (CER) are adopted, similar to the convention used in [57, 67, 2]. **Implementation Details** Unless mentioned otherwise, in all of our experiments, Tex-tAdaIN is fused into the backbone of SCATTER [38]. The experimental settings, including the optimizer, learning rate, image size, and training datasets, are identical to SCATTER [38]. Full implementation details are described in Appendix D.

4.1 Comparison to State-of-the-Art

In Table 2, we measure the accuracy of our proposed method on public handwritten text benchmarks. Our method achieves state-of-the-art results across all datasets. Compared to current state-of-the-art methods, incorporating TextAdaIN achieves a performance increase of +1.4 pp (85.9% vs. 87.3%) on IAM, +2.0 pp (92.4% vs. 94.4%) on RIMES and +0.4 pp (77.8% vs. 78.2%) on CVL. We wish to emphasize that previous methods, such as [9, 33, 42, 2], introduced complex modifications to the training phase, including adversarial learning and contrastive pre-training. In contrast, TextAdaIN can be easily implemented in a few lines of code and seamlessly fit into any mini-batch training procedure. In addition, the effect of applying MixStyle [68] and pAdaIN [45] is displayed in Table 2. Both have little to no effect on the results, indicating that handwritten text recognizers are already invariant to changes in global statistics. We refer the reader to Appendix E to see failure cases of the model.

Table 3: Generalization to new domains and architectures. Similar to [38], we show weighted (by size) average results on the regular and irregular scene text datasets. All experiments are reproduced. Integrating TextAdaIN results in consistent improvements

	Scen	e Text	Handwritten		
Method	Regular text	Irregular text	IAM	RIMES	
	5,529	3,010	17,990	7,734	
Baek et al. (CTC) [6]	88.7	72.9	80.6	87.8	
+TextAdaIN	89.5 (+ 0.8)	73.8 (+0.9)	81.5 (+ 0.9)	90.7 (+2.9)	
Baek et al. (Attn) [6]	92.0	77.4	82.7	90.2	
+TextAdaIN	92.2 (+ 0.2)	77.7 (+0.3)	84.1 (+1.4)	93.0 (+ 2.8)	
SCATTER [38]	93.6	83.0	85.7	93.3	
+TextAdaIN	94.2 (+ 0.6)	83.4 (+ 0.4)	87.3 (+1.6)	94.4 (+1.1)	
AbiNet [18]	93.9	82.0	85.4	92.0	
+TextAdaIN	94.2 (+ 0.3)	82.8 (+ 0. 8)	86.3 (+ 0.9)	93.0 (+1.0)	

4.2 Generalization of Proposed Method

In this subsection, we explore our method's transferability to both the domain of scene text and to different recognition architectures. In addition to the SCATTER and AbiNet [18] architectures, we utilize the Baek *et al.* [6] framework, which can describe the building blocks of many text recognizers, including [52, 55, 51, 53, 40, 62, 13, 14, 67, 38, 20, 66]. We choose to present TextAdaIN's performance when integrated into Baek *et al.* [6] framework while employing either a CTC [24] or an attention decoder [14, 53]. As in [38], weighted (by size) average word accuracy is adopted where regular and irregular text datasets are distinguished.

In Table 3, we present the reproduced performance of the above methods. TextAdaIN shows consistent improvement on both scene and handwritten text benchmarks independent of the chosen architecture. The results are competitive with recent stateof-the-art methods.

5 Ablation Study

In this section, we conduct a series of experiments to further understand the performance improvements and analyze the impact of our key contributions. For this purpose, we adopt the IAM dataset, and the baseline model is the SCATTER architecture [38]. Similar to our previous experiments, implementation details are described in Appendix D.

We begin by exploring the different variants of AdaIN, demonstrating that our method significantly outperforms other variants. Then, we show the compatibility of our method with other augmentation pipelines demonstrating its unique benefit as a complimentary method. To analyze the performance as a function of the granularity level, we then measure the

Table 4: **AdaIN dimensions.** Accuracy on the IAM dataset when applying the AdaIN operation over different dimensions. Our method, Tex-tAdaIN, shows significant improvement upon the baseline of **+1.6** pp

Method	AdaIN Dim	Windows	IAM Accuracy
Baseline	-	-	85.7
AdaIN	С	-	85.4
	<u>w</u>		85.7
	Н	-	85.6
	H,W	-	85.7
	C,W	-	85.8
A do IN Variationa	C,H	-	85.9
Adam variations	Ē		85.9
	W	5	85.5
	Н	5	85.9
	H,W	5	85.8
	C,W	5	85.9
TextAdaIN	C,H		87.3

accuracy while varying the number of windows. Subsequently, we demonstrate the reliance of text recognizers on local feature statistics by increasing the hyperparameter p. Lastly, we perform an analysis of our method's robustness towards challenging testing conditions, demonstrating its effectiveness on relieving the model of this shortcut.



TextAdaIN: Paying Attention to Shortcut Learning in Text Recognizers 11

Fig. 4: **Information encompassed in each dimension.** Intensity maps of the features after the first convolution layer of the encoder are visualized. Each dimension of the feature space represents different information, hence, operating on different dimensions can influence the model accordingly. TextAdaIN has two noticeable effects: (1) injecting background distortions drawn from an induced distribution and (2) introducing masking on a local scale. Both regulate the reliance on local feature statistics, as further investigated in Section 5.4

5.1 AdaIN Variations

Each dimension of the feature space represents different information about the input. Hence, modifying AdaIN to operate across different dimensions can influence the model accordingly. Our method is employed over both the channel and height dimensions between different elements of samples in a batch. Each element consists of a pre-defined number of consecutive frames in the sequential feature map. As seen in Table 4, TextAdaIN significantly improves recognition performance as opposed to all other AdaIN variations.

To better understand the information encompassed in each of the dimensions, we visualize the feature maps of a trained baseline model in Fig. 4. For this purpose, we apply PCA on the spatial dimensions of the first convolutional layer's output, thus obtaining an $H \times W \times 1$ intensity map. An image depicting the spatial intensities is displayed after normalization.

As depicted in Fig. 4, applying $AdaIN_c$, the vanilla variation, has almost no effect on the learned features. This is in alignment with the quantitative results indicating that text recognizers are relatively invariant to changes in global statistics. As for $AdaIN_{c,w}$, modifying individual vertical frames introduces subtle changes to the feature map. The network can easily compensate for these distortions, leading to minimal impact on the training process. Interestingly, applying $AdaIN_{h,w}$ injects text from the donor image into the feature map. This phenomenon originates from shifting each corresponding spatial location in the representation space. Clearly, without the modification of the labels, this effect will not improve the performance.

TextAdaIN has two major effects, as visualized in Fig. 4. The first is injecting local perturbations drawn from an induced distribution into the feature space. The distribution is induced from the manner in which TextAdaIN operates, providing a correct balance between the coarse to fine distortion level. Namely, TextAdaIN's impact is more local than AdaIN_c, yet more global in the sequence dimension than AdaIN_{c,w}. Therefore, the impact aligns with both the nature of the data and sequence-to-sequence approaches.

Occasionally, statistics of smooth areas (without text) are injected into regions of the feature space which represent text. This generates the second effect of local masking,

A	American Antion Dineline	IA	Μ	RIMES	
Architecture	Augmentation Pipeline	×	\checkmark	×	\checkmark
	X	84.6	86.5	91.6	94.2
	SCATTER [38]	85.7	87.3	93.3	94.4
SCATTER [38]	VisionLAN [19], ABINet [18]	86.3	87.3	93.4	94.7
	RandAug [17]	86.4	87.0	93.5	94.4
	Luo <i>et al.</i> [42]	85.7	87.2	93.5	94.7

Table 5: TextAdaIN's contribution is complimentary to data augmentations effect. Word accuracy for augmentation pipelines with (X) and without (\checkmark) TextAdaIN

in which part of the textual features undergo masking. We hypothesize that this forces the model to rely on semi-semantic information, which compensates for the missing visual cues. This was partially observed by Aberdam *et al.* [2] while applying horizontal cropping and in the context of speech recognition by Baevski *et al.* [7]. As this analysis was performed on the feature space, we also show the influence of TextAdaIN on the input space in Section 5.4. Moreover, to see the importance of the induced distribution causing these effects, see Appendix F.

5.2 Compatibility with Augmentation Strategies

Similarly to TextAdaIN, data augmentation strategies can be used to relieve a model's propensity to shortcuts and increase their robustness. Hence, they have a similar effect, yet they operate at different levels. One at the input space and the other at the feature space. To validate the compatibility of our method with augmentation strategies, we apply a variety of augmentation strategies on top of the same architecture [38] with and without TextAdaIN. We choose augmentation strategies ranging from other text recognition frameworks [19, 18], pseudo character-level augmentations [42] and state-of-the-art augmentation pipelines [17]. Table 5 displays the word accuracy for each augmentation pipeline with and without TextAdaIN. Applying TextAdaIN in conjunction with other augmentation strategies consistently improves performance and thus the contribution of TextAdaIN without any augmentation outperforms all other independent augmentation pipelines, indicating the effectiveness and impact of our method over the most common regularization technique.

5.3 Number of Windows

TextAdaIN splits the feature map into windows along the width axis. Each window is perceived as an individual element in the AdaIN operation. As the features vary in size at different layers, we define K to represent the number of elements created per sample. Thus, K determines the window size at each layer. Modifying K has several different effects. For example, it controls the granularity level in which the statistics are calculated and modified and the number of donors. Therefore, an optimal value of K

Table 6: **Reducing the gap in challenging testing conditions.** To demonstrate that TextAdaIN indeed regulates the usage of shortcuts, we evaluate it on corrupted versions of IAM. The corruptions are divided into three different categories based on their impact type. The normalized gap provides evidence for improved robustness, especially on local corruptions

		Local Masking		Pixel-wis	Geometric			
Corruption Type	None	Dropout	Cutout	Additive Gaussian Noise	Elastic Transform	Motion Blur	Shear & Rotate	Perspective
Baseline TextAdaIN	85.7 87.3	51.4 57.7	53.7 59.6	62.0 72.0	66.7 72.4	70.5 73.6	76.8 78.9	74.0 75.2
Gap Normalized Gap	+1.6 1.00	+6.3 3.94	+5.9 3.69	+10 6.25	+5.7 3.56	+3.1 1.94	+2.1 1.31	+1.2 0.75

can be found to balance the different effects. In Appendix G, we plot the performance as a function of K. The best result is achieved when using K = 5. We note that the average length of English words is 4.7 characters. Thus, when K = 5, the statistics are normalized per character on average.

5.4 Challenging Testing Conditions

Notice that our approach leads to a similar boost on each of the datasets across all tested architectures and augmentation settings (for example $\approx +1\%$ on IAM). This suggests that this gap indeed stems from the shortcuts in the data and is not model dependent. Thus, we expect our method to perform similarly even on future and better-performing architectures. Moreover, we remind that one can detect shortcut learning at test time, given an image that violates the shortcut hypothesis (i.e., an "f" with a different level of curvature than in the train). Thus, there are potentially even more undiscovered shortcuts than our improvements on standard test benchmarks reveal. To better demonstrate the existence of unrevealed shortcut learning, we perturb the test data in a way that changes its local style without modifying the semantic content. If the models are independent of the local style, one would expect that they will be robust to such modifications. Nevertheless, the significant performance drop (Table 1) reveals that these models are highly prone to shortcut learning, much more than discovered on the unperturbed test sets. To demonstrate that TextAdaIN indeed improves robustness towards challenging testing conditions, in Table 6, we evaluate its performance on several types of corruptions, comparing it to the baseline model.

The corruptions are divided into three categories: (1) local masking, (2) pixel-wise methods and (3) geometric transformations. For each corruption, we display the gap between the performance of the baseline versus the TextAdaIN model. To accentuate the improvement provided by our method, we also display the normalized gap - the ratio between the gap on the corrupted data and the gap on the original data. This normalization removes the performance advantage of TextAdaIN on the original data. The results indicate that TextAdaIN improves the model's robustness towards *o.o.d* testing scenario. For example, the gap on additive noise is 10%. We note that the normalized

gap, which represents TextAdaIN's robustness gains, is substantially higher in localbased corruptions rather than geometric corruptions which are applied globally. This provides further evidence that TextAdaIN regulates the reliance on local statistics in text recognizers. Despite the benefits of our method, Table 6 indicates that there are potentially even more undiscovered shortcuts as a significant gap exists when comparing to the original unperturbed test set.

5.5 Reliance on Local Statistics

In this subsection, we wish to further assert text recognizers' reliance on local statistics. Nuriel et al. [45] observed that applying AdaIN at high values of p resulted in a significant degradation of classification performance. We can thus infer that image classifiers depend on global statistics. For text recognizers, as shown in Table 7, this is not the case. Increasing the value of p, when applying AdaIN, only slightly affects the results. This indicates that global statistics are less significant in text recognition. In contrast, applying TextAdaIN with a high value of p decreases performance substantially. This implies that profusely applying TextAdaIN distorts important information that the model relies on. Therefore, text recognizers are prone to develop an unintended shortcut solution in the form of local statistics. If applied correctly, with the right p, TextAdaIN can alleviate this shortcut.

Table 7: **Probability of applying TextAdaIN.** Applying AdaIN at high values of p has little to no effect on the performance. In contrast, profusely applying TextAdaIN distorts important information that the model relies on.

Method	p	IAM Acc.
Baseline	-	85.7
	0.01	85.4
AdaIN	0.1	85.9
	0.25	85.8
	0.001	86.1
	0.01	87.3
TextAdaIN	0.05	86.8
	0.1	78.0
	0.25	74.7

6 Conclusion

Text recognizers leverage convolutional layers to extract rich visual features, and hence are extremely powerful. However, in this work, we expose their propensity towards learning an unintended "shortcut" strategy, whereby they overly rely on local statistics. Consequently, exhibiting a sensitivity towards subtle modifications that preserve image content. To relieve text recognizers' shortcut learning, we introduce TextAdaIN, a normalization-based method which distorts the feature space in a local manner and effectively regulates the reliance on local statistics. Our method achieves state-of-the-art results on handwritten text recognition benchmarks and improves robustness towards challenging testing conditions. TextAdaIN is also applicable to various recognition architectures and to the domain of scene text images. Furthermore, it can be implemented simply in a few lines of code and effortlessly integrated into a mini-batch training procedure.

By taking into account the nature of the data and the architectural characteristics, this shortcut was exposed. Yet, as we have seen, the prevalence of shortcuts is still at large. As future work, we wish to explore other shortcuts pertaining to text recognizers.

Bibliography

- Aberdam, A., Ganz, R., Mazor, S., Litman, R.: Multimodal semi-supervised learning for text recognition. arXiv preprint arXiv:2205.03873 (2022) 1, 4, 20
- [2] Aberdam, A., Litman, R., Tsiper, S., Anschel, O., Slossberg, R., Mazor, S., Manmatha, R., Perona, P.: Sequence-to-sequence contrastive learning for text recognition. arXiv preprint arXiv:2012.10873 (2020) 1, 3, 8, 9, 12, 20, 22
- [3] Afifi, M., Brown, M.S.: What else can fool deep learning? addressing color constancy errors on deep neural network performance. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 243–252 (2019) 4
- [4] Alonso, E., Moysset, B., Messina, R.: Adversarial generation of handwritten text images conditioned on sequences. In: 2019 International Conference on Document Analysis and Recognition (ICDAR). pp. 481–486. IEEE (2019) 8
- [5] Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016) 4
- [6] Baek, J., Kim, G., Lee, J., Park, S., Han, D., Yun, S., Oh, S.J., Lee, H.: What is wrong with scene text recognition model comparisons? dataset and model analysis. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 4715–4723 (2019) 1, 4, 5, 9, 10, 20, 22, 23
- [7] Baevski, A., Zhou, H., Mohamed, A., Auli, M.: wav2vec 2.0: A framework for self-supervised learning of speech representations. arXiv preprint arXiv:2006.11477 (2020) 12
- [8] Beery, S., Van Horn, G., Perona, P.: Recognition in terra incognita. In: Proceedings of the European conference on computer vision (ECCV). pp. 456–473 (2018) 4
- [9] Bhunia, A.K., Das, A., Bhunia, A.K., Kishore, P.S.R., Roy, P.P.: Handwriting recognition in low-resource scripts using adversarial learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4767–4776 (2019) 3, 4, 8, 9
- [10] Bickel, S., Brückner, M., Scheffer, T.: Discriminative learning under covariate shift. Journal of Machine Learning Research 10(9) (2009) 1
- [11] Bluche, T.: Deep neural networks for large vocabulary handwritten text recognition. Ph.D. thesis, Paris 11 (2015) 8
- [12] Bluche, T.: Joint line segmentation and transcription for end-to-end handwritten paragraph recognition. In: Advances in Neural Information Processing Systems. pp. 838–846 (2016) 8
- [13] Cheng, Z., Bai, F., Xu, Y., Zheng, G., Pu, S., Zhou, S.: Focusing attention: Towards accurate text recognition in natural images. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 5076–5084 (2017) 10
- [14] Cheng, Z., Bai, F., Xu, Y., Zheng, G., Pu, S., Zhou, S.: Focusing attention: Towards accurate text recognition in natural images. In: Proceedings of the IEEE international conference on computer vision. pp. 5076–5084 (2017) 10, 20
- [15] d'Ascoli, S., Sagun, L., Bruna, J., Biroli, G.: Finding the needle in the haystack with convolutions: on the benefits of architectural bias. arXiv preprint arXiv:1906.06766 (2019) 2

- 16 Nuriel et al.
- [16] De Palma, G., Kiani, B.T., Lloyd, S.: Deep neural networks are biased towards simple functions. arXiv preprint arXiv:1812.10156 (2018) 2
- [17] *et al.*, C.: Randaugment: Practical automated data augmentation with a reduced search space. In: CVPR 12
- [18] et al., F.: Read like humans: Autonomous, bidirectional and iterative language modeling for scene text recognition 4, 9, 10, 12, 23
- [19] et al., W.: From two to one: A new scene text recognizer with visual language modeling network. In: CVPR 4, 12
- [20] Fogel, S., Averbuch-Elor, H., Cohen, S., Mazor, S., Litman, R.: Scrabblegan: Semi-supervised varying length handwritten text generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4324–4333 (2020) 3, 8, 10, 22
- [21] Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2414–2423 (2016) 4
- [22] Geirhos, R., Jacobsen, J.H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., Wichmann, F.A.: Shortcut learning in deep neural networks. Nature Machine Intelligence 2(11), 665–673 (2020) 1, 4, 5
- [23] Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel, W.: Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In: International Conference on Learning Representations (2019), https://openreview.net/forum?id=Bygh9j09KX 1, 2, 4
- [24] Graves, A., Fernández, S., Gomez, F., Schmidhuber, J.: Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In: Proceedings of the 23rd international conference on Machine learning. pp. 369–376 (2006) 10
- [25] Graves, A., Fernández, S., Gomez, F., Schmidhuber, J.: Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In: Proceedings of the 23rd international conference on Machine learning. pp. 369–376. ACM (2006) 20
- [26] Grosicki, E., El Abed, H.: Icdar 2009 handwriting recognition competition. In: 2009 10th International Conference on Document Analysis and Recognition. pp. 1398–1402. IEEE (2009) 8, 22
- [27] Gupta, A., Vedaldi, A., Zisserman, A.: Synthetic data for text localisation in natural images. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2315–2324 (2016) 8, 22
- [28] Hermann, K.L., Chen, T., Kornblith, S.: The origins and prevalence of texture bias in convolutional neural networks. arXiv preprint arXiv:1911.09071 (2019) 2
- [29] Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1501–1510 (2017) 3, 4, 6
- [30] Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint arXiv:1502.03167 (2015) 4
- [31] Jaderberg, M., Simonyan, K., Vedaldi, A., Zisserman, A.: Synthetic data and artificial neural networks for natural scene text recognition. arXiv preprint arXiv:1406.2227 (2014) 8, 22

- [32] Jung, A.B., Wada, K., Crall, J., Tanaka, S., Graving, J., Reinders, C., Yadav, S., Banerjee, J., Vecsei, G., Kraft, A., Rui, Z., Borovec, J., Vallentin, C., Zhydenko, S., Pfeiffer, K., Cook, B., Fernández, I., De Rainville, F.M., Weng, C.H., Ayala-Acevedo, A., Meudec, R., Laporte, M., et al.: imgaug. https://github. com/aleju/imgaug (2020), online; accessed 01-Feb-2020 20
- [33] Kang, L., Rusiñol, M., Fornés, A., Riba, P., Villegas, M.: Unsupervised adaptation for synthetic-to-real handwritten word recognition. In: 2020 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 3491–3500. IEEE (2020) 8, 9
- [34] Karatzas, D., Gomez-Bigorda, L., Nicolaou, A., Ghosh, S., Bagdanov, A., Iwamura, M., Matas, J., Neumann, L., Chandrasekhar, V.R., Lu, S., et al.: Icdar 2015 competition on robust reading. In: 2015 13th International Conference on Document Analysis and Recognition (ICDAR). pp. 1156–1160. IEEE (2015) 9, 22
- [35] Karatzas, D., Shafait, F., Uchida, S., Iwamura, M., i Bigorda, L.G., Mestre, S.R., Mas, J., Mota, D.F., Almazan, J.A., De Las Heras, L.P.: Icdar 2013 robust reading competition. In: 2013 12th International Conference on Document Analysis and Recognition. pp. 1484–1493. IEEE (2013) 9, 22
- [36] Kleber, F., Fiel, S., Diem, M., Sablatnig, R.: Cvl-database: An off-line database for writer retrieval, writer identification and word spotting. In: 2013 12th international conference on document analysis and recognition. pp. 560–564. IEEE (2013) 8, 22
- [37] Li, H., Wang, P., Shen, C., Zhang, G.: Show, attend and read: A simple and strong baseline for irregular text recognition. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 8610–8617 (2019) 22
- [38] Litman, R., Anschel, O., Tsiper, S., Litman, R., Mazor, S., Manmatha, R.: Scatter: Selective context attentional scene text recognizer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020) 1, 2, 4, 5, 8, 9, 10, 12, 20, 22, 23
- [39] Liu, W., Chen, C., Wong, K., Su, Z., Han, J.: Star-net: A spatial attention residue network for scene text recognition. In: BMVC (2016) 20
- [40] Liu, W., Chen, C., Wong, K.Y.K., Su, Z., Han, J.: Star-net: A spatial attention residue network for scene text recognition. In: BMVC. vol. 2, p. 7 (2016) 10
- [41] Lucas, S.M., Panaretos, A., Sosa, L., Tang, A., Wong, S., Young, R.: Icdar 2003 robust reading competitions. In: Seventh International Conference on Document Analysis and Recognition, 2003. Proceedings. pp. 682–687. Citeseer (2003) 9, 22
- [42] Luo, C., Zhu, Y., Jin, L., Wang, Y.: Learn to augment: Joint data augmentation and network optimization for text recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13746–13755 (2020) 1, 3, 4, 8, 9, 12
- [43] Marti, U.V., Bunke, H.: The iam-database: an english sentence database for offline handwriting recognition. International Journal on Document Analysis and Recognition 5(1), 39–46 (2002) 8, 21
- [44] Mishra, A., Alahari, K., Jawahar, C.: Scene text recognition using higher order language priors (2012) 8, 22
- [45] Nuriel, O., Benaim, S., Wolf, L.: Permuted adain: Reducing the bias towards global statistics in image classification. In: Proceedings of the IEEE/CVF Con-

ference on Computer Vision and Pattern Recognition. pp. 9482–9491 (2021) 1, 4, 6, 7, 8, 9, 14

- [46] Qiao, Z., Zhou, Y., Yang, D., Zhou, Y., Wang, W.: Seed: Semantics enhanced encoder-decoder framework for scene text recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13528– 13537 (2020) 5
- [47] Quy Phan, T., Shivakumara, P., Tian, S., Lim Tan, C.: Recognizing text with perspective distortion in natural scenes. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 569–576 (2013) 9, 22
- [48] Risnumawan, A., Shivakumara, P., Chan, C.S., Tan, C.L.: A robust arbitrary text detection system for natural scene images. Expert Systems with Applications 41(18), 8027–8048 (2014) 9, 22
- [49] Sakai, T., Nagao, M., Kanade, T.: Picture structure and its processing the case of human-face photographs. In: Proceedings of Joint Conference of Electrical Engineers of Japan (October 1971) 4
- [50] Schölkopf, B., Janzing, D., Peters, J., Sgouritsa, E., Zhang, K., Mooij, J.: On causal and anticausal learning. arXiv preprint arXiv:1206.6471 (2012) 1
- [51] Shi, B., Bai, X., Yao, C.: An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. IEEE transactions on pattern analysis and machine intelligence **39**(11), 2298–2304 (2016) **4**, 10
- [52] Shi, B., Bai, X., Yao, C.: An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. IEEE transactions on pattern analysis and machine intelligence 39(11), 2298–2304 (2016) 5, 10
- [53] Shi, B., Wang, X., Lyu, P., Yao, C., Bai, X.: Robust scene text recognition with automatic rectification. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4168–4176 (2016) 10, 20
- [54] Shi, B., Yang, M., Wang, X., Lyu, P., Yao, C., Bai, X.: Aster: An attentional scene text recognizer with flexible rectification. IEEE transactions on pattern analysis and machine intelligence (2018) 22
- [55] Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014) 10
- [56] Slossberg, R., Anschel, O., Markovitz, A., Litman, R., Aberdam, A., Tsiper, S., Mazor, S., Wu, J., Manmatha, R.: On calibration of scene-text recognition models. arXiv preprint arXiv:2012.12643 (2020) 4
- [57] Sueiras, J., Ruiz, V., Sanchez, A., Velez, J.F.: Offline continuous handwriting recognition using sequence to sequence neural networks. Neurocomputing 289, 119–128 (2018) 8, 9
- [58] Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R.: Intriguing properties of neural networks. arXiv preprint arXiv:1312.6199 (2013) 1, 4
- [59] Torralba, A., Efros, A.A.: Unbiased look at dataset bias. In: CVPR 2011. pp. 1521–1528. IEEE (2011) 1
- [60] Ulyanov, D., Vedaldi, A., Lempitsky, V.: Instance normalization: The missing ingredient for fast stylization. arXiv preprint arXiv:1607.08022 (2016) 4, 6

- [61] Wan, Z., Zhang, J., Zhang, L., Luo, J., Yao, C.: On vocabulary reliance in scene text recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11425–11434 (2020) 4
- [62] Wang, J., Hu, X.: Gated recurrent convolution neural network for ocr. In: Advances in Neural Information Processing Systems. pp. 335–344 (2017) 10
- [63] Wang, K., Babenko, B., Belongie, S.: End-to-end scene text recognition. In: 2011 International Conference on Computer Vision. pp. 1457–1464. IEEE (2011) 9, 22
- [64] Wu, L., Zhu, Z., et al.: Towards understanding generalization of deep learning: Perspective of loss landscapes. arXiv preprint arXiv:1706.10239 (2017) 2
- [65] Wu, Y., He, K.: Group normalization. In: Proceedings of the European conference on computer vision (ECCV). pp. 3–19 (2018) 4
- [66] Yousef, M., Bishop, T.E.: Origaminet: Weakly-supervised, segmentation-free, one-step, full page text recognition by learning to unfold. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14710– 14719 (2020) 10
- [67] Zhang, Y., Nie, S., Liu, W., Xu, X., Zhang, D., Shen, H.T.: Sequence-to-sequence domain adaptation network for robust text image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2740– 2749 (2019) 1, 3, 8, 9, 10
- [68] Zhou, K., Yang, Y., Qiao, Y., Xiang, T.: Domain generalization with mixstyle. In: ICLR (2021) 4, 7, 8, 9