



HAL
open science

Improving Oracle Bone Characters Recognition via A CycleGAN-based Data Augmentation Method

Wei Wang, Ting Zhang, Xinxin Jin, Harold Mouchère, Xinguo Yu

► To cite this version:

Wei Wang, Ting Zhang, Xinxin Jin, Harold Mouchère, Xinguo Yu. Improving Oracle Bone Characters Recognition via A CycleGAN-based Data Augmentation Method. International Conference on Neural Information Processing (ICONIP), Nov 2022, New Delhi, India. <hal-03919404>

HAL Id: hal-03919404

<https://hal.science/hal-03919404v1>

Submitted on 2 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Improving Oracle Bone Characters Recognition via A CycleGAN-based Data Augmentation Method^{*}

Wei WANG¹, Ting ZHANG^{1✉}, Yiwen ZHAO¹, Xinxin JIN¹, Harold MOUCHERE², and Xinguo YU¹

¹ Faculty of Artificial Intelligence in Education, Central China Normal University, Wuhan, China

² Nantes Université, École Centrale Nantes, LS2N, UMR 6004, F-44000, Nantes, France

✉ting.zhang@ccnu.edu.cn

Abstract. Oracle bone inscription is the earliest writing system in China which contains rich information about the history of Shang dynasty. Automatically recognizing oracle bone characters is of great significance since it could promote the research on history, philology and archaeology. The proposed solutions for oracle bone characters recognition are mainly based on machine learning or deep learning algorithms which rely on a large number of supervised training data. However, the existing dataset suffers from the problem of severe class imbalance. In this work, we propose a CycleGAN-based data augmentation method to overcome the limitation. Via learning the mapping between the glyph images data domain and the real samples data domain, CycleGAN could generate oracle character images of high-quality. The quality is evaluated using the quantitative measure. Totally, 185362 samples are generated which could serve as a complementary to the existing dataset. With these generated samples, the state of the art results of recognition task on OBC306 are improved greatly in terms of mean accuracy and total accuracy.

Keywords: Characters recognition · Class imbalance · Data augmentation · CycleGAN.

1 Introduction

Oracle bone inscriptions, usually inscribed on turtle shells and animal bones, are the oldest hieroglyphs in China. They record the development of civilization of late Shang dynasty. The study of oracle bone inscriptions could promote the research on the origin of Chinese characters, as well as the research on the history of Shang dynasty. Since firstly discovered in the city of Anyang, China, around

^{*} This work is supported by National Natural Science Foundation of China (No. 62007014), China Post doctoral Science Foundation (No. 2019M652678) and the Fundamental Research Funds for the Central Universities(No. CCNU20ZT019).

160000 bones or shells (see Figure 1) were unearthed gradually from which about 4500 classes of characters were found however more than 2000 characters remain not deciphered so far. Currently, the recognition and decipherment of oracle bone characters majorly rely on the manual method by domain experts with high expertise which to a certain degree blocks the development of this field. Thus, developing automatic recognition methods is in need and of great significance.



Fig. 1: Examples of unearthed oracle bone inscriptions.

The state of the art methods for oracle bone characters recognition are mainly based on deep learning algorithms. This task is more difficult than other hand-written characters recognition tasks since it exhibits several challenges. (1) Large intra-class variance. As oracle bone characters were carved by different ancients from multiple regions over several historic periods without much standardization [1], the same category could have several variants. Practically in China, characters were unified for the first time in Qin dynasty, after Shang dynasty. (2) Class imbalance. Due to the scarcity of the unearthed bones or shells, few instances were found for some characters. The distribution of the instances with respect to each category is quite imbalanced, following the long-tail distribution [8]. (3) Low quality. Some inscriptions were damaged to some degree during the long burial period or the excavation process. Accordingly, there exist some noisy or broken character instances in the collected dataset [8].

Several oracle bone character datasets have been built up to now for the research aim such as Oracle-20K [7] (20039 instances belonging to 261 categories), Oracle-AYNU [6] (39062 instances belonging to 2583 categories) and OBC306 [8] (309551 instances belonging to 306 categories) among which OBC306 has the largest scale and is closer to the real case as it uses the published oracle prints as the raw materials. As stated previously, the existing oracle datasets are quite imbalanced due to the scarcity of the unearthed materials. To overcome this limitation, in this work we propose a data augmentation method for oracle characters using Cycle-Consistent Adversarial Network (CycleGAN)[12]. CycleGAN is trained to learn the mapping between the glyph images and the real samples to develop the capability of generating oracle images of high-quality.

The main contributions of this paper are as follows:

- We propose a feasible solution to generate oracle character images using CycleGAN by enabling it to learn the mapping between the glyph images collection and the real samples collection. Totally, 185362 samples are gener-

ated which will be released later as a supplementary to the existing dataset for boosting the research in this field.

- To our knowledge, it is the first attempt to use the GAN-based methods to generate oracle character images.
- With the generated samples, the state of the art results of recognition task on OBC306 are greatly improved in terms of mean accuracy and total accuracy.

The rest of this paper is organized as follows. The related works are reviewed shortly in Section II. Section III introduces the proposed method. The experiments and results are discussed in Section IV. Section V concludes the work and proposes several perspectives for future work.

2 Related works

2.1 Oracle bone characters recognition

The research on oracle bone characters recognition is quite limited. Before the deep learning era, the proposed methods [2] [3] [4] for oracles recognition mainly follows the traditional pipeline, being pre-processing, feature extraction and classification. The accuracy of the models relies much on features extracted. Researchers usually design features representing oracles based on graph theory and topology theory.

The idea proposed in [2] regards an oracle character as a non-directed graph and uses the topological features for classification. The method in [3] firstly transforms the inscriptions to labelled graphs, then extracts the quasi-adjacency matrix of the graphs as features, finally perform the recognition using an algorithm based on graph isomorphism. [4] takes the skeleton and line points as features which later are fed into a two-stage recognition process. Even though these manually designed features were reported to work well on specific datasets, they have a low generalization capability.

Owing to the advantage of deep convolutional neural networks (DCNNs) [5], feature representations could be learned from data automatically. These feature representations are robust, and have strong generalization capability. Guo et al. [7] develop a novel hierarchical representation which combines a Gabor-related low-level representation and a sparse-encoder-related mid-level representation for oracle bone characters. This hierarchical representation serves as a complementary to CNN-based models. One point should be noted that [7] discards oracle characters with less than 25 instances which means they skip from the problem of class imbalance. To overcome the limitations, [6] propose a deep metric learning based classification model which uses a CNN to map oracle images into an Euclidean space and then carries out the classification by the Nearest Neighbour (NN) rule. This method achieves the state of the art results on the public dataset Oracle-20K and Oracle-AYNU. Almost at the same time, a large-scale oracle bone character dataset was released which consists of 309551 samples. However, this dataset also suffers the problem of class imbalance by nature. To solve the imbalance problem, [19] proposed a mix-up augmentation strategy to

synthesize samples in the feature space by leveraging information from both majority and minority classes. Experimental results verify the effectiveness of the proposed method. However, [19] synthesize samples in the feature space which makes it difficult to evaluate the quality of the synthesized samples explicitly. Different from work [19], we propose to generate samples directly for minority classes with GAN-based methods.

2.2 Generative Adversarial Networks (GANs)

GAN [9] [10] was firstly proposed by Ian GoodFellow in 2014 and since then it has quickly become a hot research topic. GAN consists of two models, being a generator and a discriminator, which can be implemented by any form of differentiable system such as neural network. The generator aims to capture the distribution of real samples for new data generation while the discriminator struggles to discriminate generated data from the real data as accurately as possible [11]. GANs are widely applied in the field of image processing and computer vision, such as image super-resolution, image synthesis and manipulation, etc. Many interesting variants of GANs have been proposed since 2014.

Conditional GANs (cGANs) [13] extend GANs to a conditional model by conditioning the generator and discriminator on some extra information. With cGANs, samples can be generated conditioning on class labels, texts, etc. In [14], cGAN is used for image-to-image translation. The developed software is named as pix2pix. The generator aims to learn a mapping from the input image (instead of random noise) to the generated image. pix2pix successfully applied cGANs for image-to-image translation however the need for the aligned image pairs hinders its generalisation.

The limitation of pix2pix is well solved by Cycle-consistent GANs (CycleGAN)[12] which could learn to translate an image from one domain to another without paired examples. By exploiting the property that translation should be cycle consistent, CycleGAN introduce a cycle consistency loss in addition to the typical adversarial loss.

Even though the limitation of pix2pix seems not to be a hinder in the context of our task and a paired image-to-image translation network can be trained to generate higher quality images, our task benefits a lot from the unpaired image-to-image training. Because being different from other image generation tasks which pursue only the high quality of images, our main goal is to improve the recognition accuracy, and it's mostly influenced by the diversity of the data. CycleGAN [12] tries to learn to the mapping between two domain collections without aligned image pairs. This working mechanism is beneficial to the oracle character images generation task especially for minority classes which could utilize rich stroke variety information from majority classes. Therefore, we propose to generate oracle bone character images based on CycleGAN.

3 Method

To our knowledge, OBC306 is the existing largest dataset for oracle bone characters recognition. However, it suffers a severe problem of class imbalance due to the scarcity of the unearthed inscriptions. Our aim is to generate samples for minority classes in OBC306 using GAN-based methods. The vanilla GAN model [9] generates samples without designated class labels. The generated sample could belong to any class. However, our target is to generate samples belonging to designated classes. One possible way is to train a GAN model for each class. Unfortunately, training one model for each class requires a large amount of samples which is in conflict with minority classes. This contradiction makes the task based on the vanilla GAN model be challenging. The propose of cGAN [13] makes it possible to generate samples conditioning on specific class labels. cGAN is trained to learn the mapping from latent code to the corresponding image which is especially complex for a domain like oracle characters which have complicated glyphs. Furthermore, it also requires considerable number of samples from each class. *pix2pix* [14] is one successful application of cGAN in the image-to-image translation task. It uses aligned image pairs from two different domains (such as gray scale and color) to train the model to learn the mapping from one domain image to the other. The major limitation of *pix2pix* is its requirement for aligned image pairs. Even though this limitation seems not to be a hinder in the context of our task and a paired image-to-image translation network can be trained to generate higher quality images, our task benefits a lot from the unpaired image-to-image training. Because being different from other image generation tasks which pursue only the high quality of images, our main goal is to improve the recognition accuracy, and it's mostly influenced by the diversity of the data. CycleGAN [12] tries to learn to the mapping between two domain collections without aligned image pairs. This working mechanism is beneficial to the oracle character images generation task especially for minority classes which could utilize rich stroke variety information from majority classes. Based on the analysis above, A CycleGAN-based data augmentation method for improving oracle bone characters recognition is proposed in this work. Two data domains are assumed as glyph images and real character images respectively.

3.1 Statistics of OBC306

OBC306 [8] consists of 309551 instances belonging to 306 categories. There are 70 categories with more than 1000 samples, including 5 categories with more than 10000 images. 236 categories with less than 1000 samples, including 85 categories with less than 100 samples and 29 categories with only one sample. Figure 2 illustrates the distribution of data in OBC306. As can be seen, the distribution is quite imbalanced. We set 1000/category as the target of augmentation.

3.2 Extraction of the glyphs of oracle characters

We extract the glyph of oracle from the real samples. The process for extraction of the glyph is quite simple, being binarization, noise erasing and image inverse




Fig. 2: Illustration of data distribution in OBC306.

if required. Figure 3 presents several cases of the extracted glyph. The glyph images will be one of the image collections of CycleGAN.




Fig. 3: Some examples of the extracted glyph.

3.3 CycleGAN




Fig. 4: The framework of CycleGAN.

Given two image collections X and Y , CycleGAN aims to learn the mapping between X and Y . As illustrated in Figure 4, CycleGAN consists of two generators $G_1 : X \rightarrow Y$ and $G_2 : Y \rightarrow X$, and two discriminators D_X and D_Y . G_1 tries to generate images $G_1(x)$ which are indistinguishable from images $\in Y$ and D_Y tries to distinguish between generated samples $G_1(x)$ and real samples




Fig. 5: The network architecture of the generator used in this work.

$\in Y$. The same case is for G_2 and D_X . In order to constrain the learned mapping further, the consistency loss (Equation 1) is introduced in addition to the adversarial loss (Equation 2). The consistency loss term tries to make the model work as: $x \rightarrow G_1(x) \rightarrow G_2(G_1(x)) \approx x$ and $y \rightarrow G_2(y) \rightarrow G_1(G_2(y)) \approx y$. The full objective is shown as Equation 3.

$$\mathcal{L}_{cyc}(G_1, G_2) = \mathbb{E}_{x \sim p_{data}(x)} [\|G_2(G_1(x)) - x\|_1] + \mathbb{E}_{y \sim p_{data}(y)} [\|G_1(G_2(y)) - y\|_1] \tag{1}$$

$$\mathcal{L}_{GAN}(G_1, D_Y, X, Y) = \mathbb{E}_{y \sim p_{data}(y)} [\log D_Y(y)] + \mathbb{E}_{x \sim p_{data}(x)} [\log (1 - D_Y(G_1(x)))] \tag{2}$$

$$\mathcal{L}(G_1, G_2, D_X, D_Y) = \mathcal{L}_{GAN}(G_1, D_Y, X, Y) + \mathcal{L}_{GAN}(G_2, D_X, Y, X) + \lambda \mathcal{L}_{cyc}(G_1, G_2) \tag{3}$$

As CycleGAN does not require paired samples, we randomly select 1000 real images and 1000 glyph images to train a CycleGAN model. Both generators and discriminators are implemented with neural networks. We use the same network architecture with literature [12]. The generator consists of three convolution layers for extracting features from input images, nine residual blocks which could translate feature vectors from the source domain to the target domain, two deconvolution layers and one convolution layer for restoring low-level features from feature vectors and generating images following the distribution of the target domain. The architecture is presented in Figure 5. For the discriminator (Figure 6), we use five convolution layers among which the last one is used to produce a 1-dimensional output telling input images are real or fake.




Fig. 6: The network architecture of the discriminator used in this work.

3.4 Samples generation for categories with few instances

Our object is to generate instances for oracle characters with few real instances to solve the problem of class imbalance. 1000/category is the target of augmentation. However for some categories with few real instances, the corresponding glyph images are quite few also. To achieve the target, we enrich the glyph images set first via general augmentation methods such as rotation, noise generation, occlusion, etc. Then feed these glyph images to multiple trained CycleGAN models to generate oracle character images. For example, assume that we need to generate 900 instances for some category. First step is to get 30 glyph images via general augmentation methods if required and next is to feed these 30 glyph images to 30 trained CycleGAN models to yield 900 instances.

4 Experiments

4.1 Experiments on oracle character images generation

Settings As stated previously, we evaluate the proposed method on the OBC306 dataset. We randomly select 1000 real images and 1000 glyph images (unpaired) to train a CycleGAN model. The model is trained for 200 epochs with Adam optimizer [15]. For the first 100 epochs, the learning rate is set to 0.0002; it linearly decays to zero over the next 100 epochs.

Metrics GAN-train and GAN-test as defined in [17] are used as quantitative metrics to evaluate the quality of generated samples.

- **GAN-train** represents the accuracy of a classification model trained with the synthetic samples and evaluated on real samples.
- **GAN-test** is the accuracy of a classification model trained with real data, and evaluated on synthetic data.
- **GAN-base** denotes the accuracy of a baseline network trained and tested on real data [18].

Via comparing GAN-train accuracy and GAN-base accuracy, GAN-test accuracy and GAN-base accuracy, we can check the quality of the generated samples. Ideally, both GAN-train and GAN-test should be close to GAN-base accuracy.

Results We carry out samples augmentation for 236 categories with less than 1000 instances. Totally, there are 185362 samples generated. Figure 7 compares the generated samples with the real samples of the same category. From the perceptual point of view, there is still a discrepancy between the real ones and the generated ones.

We also take an automatic quantitative measure on the generated samples using GAN-train and GAN-test. As CycleGAN models are trained with samples randomly selected, in other words, with unbalanced data samples, we question that if the quality of generated samples is different or not among categories with




Fig. 7: Comparison of the real samples (a) and the generated samples (b).

variable numbers of real instances. Thus, experiments are carried out respectively on categories with different numbers of real instances. Specifically, we select 5 categories which have around 500 real instances to conduct Experiment *A* and 5 categories which have 10-100 real instances to conduct Experiment *B*. From Table 1, we can see in Experiment *A*, GAN-train is relatively close to GAN-base (87.4% VS 83.2%) which tells that CycleGAN models generate for the selected 5 categories high-quality images with similar diversity to the real data. In Experiment *B*, GAN-train is 9.5% less than GAN-base (81.6% VS 72.1%) which means the generated samples for the selected categories lost some patterns compared to the real data. For both cases, GAN-test is not far away from GAN-base and is not larger than GAN-base telling that the trained CycleGAN models do not have the problem of over-fitting. In conclusion, CycleGAN tends to generate samples of higher quality for categories with a lot of real instances.

Table 1: Results of GANtrain and GANtest experiments. *A*: experiment on categories with around 500 real instances; *B*: experiment on categories with 10-100 real instances

Experiment	GAN-base	GAN-train	GAN-test
<i>A</i>	87.4%	83.2%	84.3%
<i>B</i>	81.6%	72.1%	76.8%

4.2 Experiments on oracle characters recognition

Settings We adopt AlexNet [5], ResNet50 [16] and Inception-v4 [20] which have shown good performance on OBC306 as the classification models. For each category, 80% of the samples are used for training, and the remaining 20% are used for test. Classifiers are trained with Adam optimizer also.

To evaluate the contribution of generated samples to recognition task, we set three experiments. The three experiments use the same test data. In terms of training data, Experiment I uses samples from the original OBC306 dataset only.

Experiment II uses samples from the original OBC306 dataset + samples generated with traditional methods. Experiment III uses samples from the original OBC306 dataset + samples generated via CycleGAN. The numbers of training data for Experiment II & III are the same, larger than the number of training data in Experiment I .

Metrics To account for long-tail effects, we measure the performance of the model with regard to multiple metrics as same as [19].

- **Mean accuracy** denotes the accuracy that is calculated by averaging the accuracy of per class.
- **Total accuracy** refers to the overall accuracy of all the test samples on all classes.

Table 2: Results of oracle characters recognition on the OBC306 dataset.

Experiment	Model	Mean accuracy			Total accuracy
		Top 1	Top 5	Top 10	
I: Original OBC306	AlexNet [8]	65.24%	74.22%	79.15%	—
	ResNet50 [8]	68.77%	78.67%	81.32%	—
	Inception-v4 [8]	69.51%	78.11%	81.16%	86.41%
	Inception-v4 (277 classes)*	75.62%	—	—	89.81%
II: Traditional augmentation	AlexNet	66.41%	75.24%	80.10%	—
	ResNet50	71.89%	80.27%	81.74%	—
	Inception-v4	74.61%	82.92%	84.23%	89.92%
III: CycleGAN-based augmentation	AlexNet	70.17%	77.22%	82.71%	—
	ResNet50	77.63%	87.15%	91.83%	—
	Inception-v4	83.44%	91.39%	93.13%	92.02%
IV: Mix-up augmentation (277 classes)*	Inception-v4 [19]	79.02%	—	—	91.56%

* The results of these two rows are extracted directly from the published literature [19]. They remove 29 classes which have only one instance and carry out the experiments on the remaining 277 classes. The rest of the results in this table are produced by us.

Results Table 2 presents recognition accuracies in different experiments. From the statistics, we can see the samples generated by CycleGAN could improve the recognition accuracy dramatically: mean accuracy by approximately 14% (69.51% \rightarrow 83.44%), total accuracy by approximately 6% (86.41% \rightarrow 92.02%). However, the samples generated by traditional methods only improve the mean accuracy by around 5% (69.51% \rightarrow 74.61%) and total accuracy by around 3% (86.41% \rightarrow 89.92%). [19] proposed a mix-up augmentation strategy to synthesize samples by fusing features linearly from two examples. This method synthesizes

samples in the feature space which makes it difficult to evaluate the quality of the synthesized samples explicitly. Furthermore, our method outperforms the mix-up strategy in terms of both mean and total accuracy. While their total accuracy is improved by around 2% (89.81% \rightarrow 91.56%), our total accuracy is improved by around 6% (86.41% \rightarrow 92.02%).

Here is one point we need to clarify. For 29 classes with only one sample, we put them in the test set. Thus, there is no data from these classes to train CycleGAN models. But, we feed glyph images of these classes into trained CycleGANs to generate samples (1000 per class). These operations maybe influence the mean accuracy a lot. However, the total accuracy will not be affected much as there are around 60000 samples in the test set. Furthermore, the total accuracy is improved by around 6% which verifies the effectiveness of the proposed method.

To further analyse the effects of generated samples on the recognition accuracy of classes with variable numbers of real instances, we design a new recognition experiment. (1) Firstly randomly select 5 categories which have 1000-1100 real instances. (2) Secondly remove a part of the real instances to make these 5 categories be minority classes. (3) Use the traditional augmentation method and the CycleGAN-based method respectively to recover the original quantity. (4) ResNet50 is adopted as the classifier. Figure 8 shows the recognition results under different settings. As can be seen, when 90% real data removed, CycleGAN-based method exhibits significantly better performance than the traditional method. When 70% real data removed, CycleGAN-based method slightly outperforms the traditional method. When 50% real data removed, there is no obvious discrepancy between the two methods. Form this experiment, we could tell that the CycleGAN-based augmentation method has a good performance for minority classes. For classes with around 500 real instances, only a slight improvement is observed. The result of the case using 50% real data is close to the case using 100% real data which means the classifier learns the patterns of oracle characters with around 500 samples each category.




Fig. 8: Results of oracle characters recognition on specific classes.

5 Conclusion

In this work, we propose a CycleGAN-based data augmentation method to solve the problem of class imbalance in the existing oracle characters dataset. Totally, 185362 samples are generated which could be a supplementary to the existing dataset. With these generated samples, the SOTA results of recognition task on OBC306 are improved greatly in terms of mean accuracy and total accuracy.

However, some limitations exist in our work. Currently, we extract the glyph from real samples only and take some simple transformations if required. In fact, during the training process of CycleGAN, the features learned from real samples also guide the structure of glyph. Thus, extracting the glyph from generated samples could be an interesting attempt in the future. Another point is there is no attempt to solve the problem of large intra-class variance which will be left to the future.

References

1. Flad, R.-K.: Divination and power: A multi-regional view of the development of oracle bone divination in early China. *Current Anthropol*, 49(3), 403–437 (2008)
2. Li, F., Woo, P.-Y.: The coding principle and method for automatic recognition of jia gu wen characters. *International Journal of Human Computer Studies*, 53(2), 289–299 (2000)
3. Li, Q.-S., Yang, Y.-X., Wang, A.-M.: Recognition of inscriptions on bones or tortoise shells based on graph isomorphism. *Computer Applications in Engineering Education*, 47(8), 112–114 (2011)
4. Meng, L.: Two-Stage Recognition for Oracle Bone Inscriptions. In: *International Conference on Image Analysis and Processing*, pp. 672–682 (2017)
5. Krizhevsky, A., Sutskever, I., Hinton, G.-E.: Imagenet classification with deep convolutional neural networks. In: *International Conference on Neural Information Processing Systems* (2012)
6. Zhang, Y.-K., Zhang, H., Liu, Y.-G., Yang, Q., Liu, C.-L.: Oracle character recognition by nearest neighbor classification with deep metric learning. In: *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pp. 309–314 (2019)
7. Guo, J., Wang, C., Roman-Rangel, E., Chao, H., Rui, Y.: Building hierarchical representations for oracle character and sketch recognition. *IEEE Transactions on Image Processing*, 25(1), 104–118 (2015)
8. Huang, S., Wang, H., Liu, Y., Shi, X., Jin, L.: OBC306: A large-scale oracle bone character recognition dataset. In: *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pp. 681–688 (2019)
9. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *NIPS* (2014)
10. Zhao, J., Mathieu, M., LeCun, Y.: Energy-based generative adversarial network. In: *ICLR* (2017)
11. Gui, J., Sun, Z., Wen, Y., Tao, D., Ye, J.: A Review on Generative Adversarial Networks: Algorithms, Theory, and Applications. *IEEE Transactions on Knowledge and Data Engineering* (2021)

12. Zhu, J., Park, T., Isola, P., Efros, A.-A.: Unpaired Image-to-Image Translation Using Cycle-Consistent Adversarial Networks. In: 2017 IEEE International Conference on Computer Vision (ICCV), 2017, pp. 2242–2251 (2017)
13. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784 (2014)
14. Isola, P., Zhu, J.-Y., Zhou, T., Efros, A.-A.: Image-to-image translation with conditional adversarial networks. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 1125–1134 (2017)
15. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770–778 (2016)
17. Shmelkov, K., Schmid, C., Alahari, K.: How good is my GAN. In: European Conference on Computer Vision (ECCV), pp. 213–229 (2018)
18. Bissoto, A., Valle, E., Avila, S.: The Six Fronts of the Generative Adversarial Networks . arXiv preprint arXiv:1910.13076 (2019)
19. Li, J., Wang, Q.-F., Zhang, R., Huang, K.: Mix-Up Augmentation for Oracle Character Recognition with Imbalanced Data Distribution. In: Proceedings of International Conference on Document Analysis and Recognition (2021)
20. Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A.-A.: Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning. In: Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (2017)

Appendix




Fig. 9: Illustration of samples generated via CycleGAN. The samples in a row belong to the same category.