

Two-Step Data Augmentation for Masked Face Detection and Recognition: Turning Fake Masks to Real

Yan Yang Aaren¹, George Bebis¹ and Mircea Nicolescu¹

¹*Department of Computer Science and Engineering, University of Nevada, Reno, 1664 N Virginia St, Reno, NV 89557, USA*
yanyang@nevada.unr.edu, {bebis, mircea}@cse.unr.edu

Keywords: Generative Adversarial Networks, Data Augmentation, Face Recognition

Abstract: The COVID-19 spread raised urgent requirements for masked face recognition and detection tasks. However, the current masked face datasets are insufficient. To alleviate the limitation of data, we proposed a two-step data augmentation that combines rule-based mask warping with unpaired image-to-image translation. Our qualitative evaluations showed that our method achieved noticeable improvements compared to the rule-based warping alone and complemented results from other state-of-the-art GAN-based generation methods, such as IAMGAN. The non-mask change loss and the noise input we used to improve training showed effectiveness. We also provided an analysis of potential future directions based on observations of our experiments.

1 INTRODUCTION

Computer Vision tasks like recognition, detection, classification, etc., performed on occluded human faces, existed even before the COVID-19 outbreak (Ge et al., 2017). The spread of COVID-19 has urgently imposed performance and robustness requirements to such applications. However, multiple latest masked-face detection or recognition works (Singh et al., 2021; Montero et al., 2021) claim their models only as starting points of future transfer learning with more data, instead of final results, or at least recognize the data insufficiency problem. The research community has built mature datasets of dominantly full faces, but masked face datasets are still under construction.

For full faces, recognition tasks use single-face image sets, with every single identity assigned to multiple facial images (Masi et al., 2018), and detection algorithms work on multi-face scene images and learn bounding box locations together with an optional class label for each box (Jain and Learned-Miller, 2010; Yang and Jiachun, 2018). To improve the masked-face datasets to facilitate the same learning approaches, we either collect and annotate raw data or generate artificial images to augment existing data (Wang et al., 2020a). In this paper, we focus on generating artificially masked faces.

Recent works address this problem by modifying unmasked faces into masked. Some use rules to warp masks onto faces (Wang et al., 2020b; Anwar and Raychowdhury, 2020; Cabani et al., 2021). One uses Neural Network (NN) to translate unmasked faces into masked (Geng et al., 2020). The rule-based methods provide realistic mask textures and completely avoid the risk of distorting other parts of

faces. However, they often result in bad transitions between masks and background faces, and the lighting on masks often looks unreal. Some rule-based algorithms (Anwar and Raychowdhury, 2020; Wang et al., 2021) achieved mask diversity by defining different mask image options, but this diversity is limited to the predefined mask types. In contrast, NN methods learn to avoid facial distortions and generate mask textures, often to a reasonable extent but never with promise. They provide natural transition, realistic details, and sometimes more diversity in compensation.

We propose a two-stage approach combining a rule-based method and image-to-image translation (I2I). After applying rule-based methods to full faces, we use an I2I model to translate rule-generated masks into more realistic ones. Rule-generated mask regions are calculated to serve as ground truth attention areas, from which we designed an extra loss to restrict I2I modifications only to mask regions. The rest of the paper will call the raw data “full-face” images, the faces with rule-based masks “rule-based mask” images, and the final outputs “realistic mask” images.

The applications of fake mask methods include not only masked face recognition/classification but also detection that requires multi-face images. By extracting bounding boxes in multi-face images, converting a portion into masked, and overlaying masked faces on original boxes, we transform multi-face images to serve masked face detection training.

2 RELATED WORK

Real-World Datasets. Traditional full-face datasets embody a large number of faces with high variations in demographics, head rotations, facial landmarks,

occlusion degrees, facial expressions, etc. (Sagonas et al., 2016; Karras et al., 2018)¹ Annotations to these datasets include categories mentioned above, as well as subject identities for facial recognition, bounding boxes for facial detection, and so on.

On the other hand, datasets emphasizing masked faces are limited in the quantity of data and the variation of features. Wang et al. (2020b) developed two different datasets based on real-world images. Real-World Masked Face Dataset (RMFD) (Wang et al., 2020b) for recognition tasks contains 5,000 masked and 90,000 normal faces belonging to 525 people. Masked Face Detection Dataset (MFDD) (Wang et al., 2020b) contains 24,711 masked face images for detection tasks, which is currently not publicly accessible². MAsked FAcEs (MAFA) by Ge et al. (2017) with 30,811 internet images and 35,806 masked faces is the largest real-world masked face dataset to our best knowledge. It is annotated with relatively abundant information such as face bounding boxes, mask bounding boxes, mask types, face orientations, occlusion degrees, gender, race, and more, but MAFA alone is not comparable with all the diversity of currently available full-face datasets. Moreover, no identity information is provided to MAFA, so it is only for face detection tasks. Besides RMFD, recognition tasks may also use Masked Faces in Real World for Face Recognition (MFR2) (Anwar and Raychowdhury, 2020), a small set of 269 images belonging to 53 identities. Geng et al. (2020) proposed a two-part Masked Face Segmentation and Recognition (MFSR) dataset. 9,742 masked images were collected from the web and annotated with mask regions in the first part. The second part contains 11,615 faces, masked or non-masked, for 704 real-world identities and 300 internet-obtained identities.

Artificially Masked Faces. Some researchers generated artificial data to augment existing real-world masked face datasets. Stemmed from RMFD and MFDD, the same lab proposed Simulated Masked Face Recognition Dataset (SMFRD) (Wang et al., 2020b), using a naive copy-and-paste method to put cartoon mask images onto existing face recognition datasets. Anwar and Raychowdhury (2020) and Cabani et al. (2021) separately used more sophisticated methods to warp mask images onto faces based on detected facial landmarks (Sagonas et al., 2016). Anwar and Raychowdhury (2020) provided their method as a MaskTheFace tool for both single- and multi-face images. Cabani et al. (2021) published

a single-face dataset, MaskedFace-Net, applying their rule-based method to Flickr-Faces-HQ (Karras et al., 2018). MaskedFace-Net consists of Correctly and Incorrectly Masked Face Dataset (CMFD and IMFD) with about 70,000 synthesized images each. CMFD is used as our model input. While our paper was in progress, Wang et al. (2021) published MLFW (Masked LFW), which enhanced the landmark-based warping by extra rule-based improvements on the unreal lighting and abrupt mask boundaries.

Beyond rule-based methods, Geng et al. (2020) proposed an NN model, Identity Aware Mask GAN (IAMGAN), to synthesize masked faces. It consists of a CycleGAN-like generation module and an Identity Preservation (IP) module. The IP module has a mask region predictor and an identity classifier. The former predicts mask regions and removes them from both masked (output) and unmasked (input) images. The latter classifies identities using the rest parts and penalizes identity change between input and output based on information from multiple classifier levels.

GAN Models for I2I. Image to image translation is a problem of mapping images from one domain into another, with the intrinsic source content preserved and the extrinsic target style transferred (Pang et al., 2021). It has been used in a broad set of scenarios such as image synthesis (Regmi and Borji, 2018), image segmentation (Guo et al., 2020), style transfer (Zhu et al., 2017), and more. GAN models are adapted to the I2I tasks by replacing the generator’s standard random input with source image embeddings and keeping the discriminator’s function of distinguishing synthesized and authentic images.

The earliest versions of GANs for image translation, such as pix2pix (Isola et al., 2016), require aligned image pairs as the training data. However, getting a large number of strictly aligned image pairs is hard or even impossible in many situations. Zhu et al. (2017) proposed CycleGAN using a cyclic loss, making it possible to train an I2I on unpaired data. The previously mentioned IAMGAN (Geng et al., 2020) is an enhanced CycleGAN aiming specifically to generate masks. Based similarly on CycleGAN, Tang et al. (2021) trained the generator to produce attention masks together with the generated image contents, which guided the generator to translate individual objects without changing the background. They call their general-purpose model “AttentionGAN”.

Our work adapts AttentionGAN (Tang et al., 2021). The key innovation is that our source images are not full-face but rule-based mask images, i.e., CMFD (Cabani et al., 2021). We use a rule-based method as the first step and the adapted AttentionGAN as the second step. We show that: 1. warping

¹Occlusions here include various types like body parts, scarves, etc. Medical masks constitute only a tiny portion.

²<https://github.com/X-zhangyang/Real-World-Masked-Face-Dataset/issues/16>



Figure 1: Examples from dataset B.

mask images onto full faces provides style guidance and ground truth attention for better GAN model results; 2. the GAN model renders the rule-based results into more realistic details. With the latest advances in rule-based methods such as Wang et al. (2021), our two-step proposal may still improve details.

3 DATA

Like other unpaired I2I models, both training and testing of AttentionGAN require two sets of data, A and B, A being the source and B being the destination.

Dataset B. We manually cropped web images and extracted MAFA bounding boxes for dataset B. From MAFA, we first extracted 8938 single faces with “simple” masks, full occlusions, sizes of at least 60×60 pixels, and front-facing orientations. Finding that annotations for MAFA are not accurate, we followed it by hand-picking a subset that strictly matches the criteria mentioned above. At the same time, we added “faces without pitch/roll with light-colored medical masks” as additional criteria. In this way, we have a controlled dataset for a more accessible proof of our concept that a superimposed mask image helps the subsequent I2I step, and they together achieve better results than the superimposing alone. With 1597 final images from MAFA, we supplemented the small subset with 98 cropped faces from open-source photos on <https://unsplash.com/>³, resulting in a total of 1695 images for training set B. Example images from B are shown in Figure 1

Dataset A. We use down-sampled CMFD as set A, which uses uniformly blue medical masks with rare occasions of misplacement. The downsampling makes sure the size of set A matches that of set B.

4 METHOD

AttentionGAN is designed based on CycleGAN and shares with CycleGAN the co-training of translation models in two opposite directions. We use most of the default parameters in the AttentionGAN repository⁴, with batch size 4, learning rate 0.0002, Adam momentum 0.5, and weights initialized from Gaussian distribution $w \sim \mathcal{N}(0.0, 0.02)$.

³Images credit to Jana Shnipelson et al. on Unsplash

⁴<https://github.com/Ha0Tang/AttentionGAN>

Our input images are all resized to 256×256 with no cropping before being fed into the models. In each translation direction, our generator starts with a 3-pixel Reflection Padding followed by three convolutional layers with instance normalizations and nine ResNet blocks. In the convolutional layers, we increase the number of channels at each layer, and in the ResNet blocks, we keep the number of channels unchanged. Still in the generator, content tensors and attention masks are generated after the nine ResNet blocks in two separate pipelines, each of which consists of three Transposed Convolutional (TC) layers with two instance normalizations, where the number of channels increases layer-wise. The content tensors are activated by tanh and attention masks by softmax.

As shown in Figure 2, the generator produces corresponding attention masks for both output content tensors and the input image tensor. The tensor values are filtered by their attention masks and summed up to produce the final generator result.

On the other hand, the discriminator is a three-layer PatchGAN with kernel size 4×4 , judging the input image’s realness based on whether each 4×4 patch looks real. While retaining most of AttentionGAN’s structures, we used multiple ways to adapt them to our training needs.

4.1 Adding Non-Mask Change Loss

As mentioned, the generator in AttentionGAN outputs a set of attention masks, which learn to find the most distinguishing parts between sets A and B unsupervised. However, we obtained sets A and B from different sources, causing heterogeneity beyond the facial mask differences. Examples include 1. faces in set B are more zoomed-in, in most extreme cases with foreheads partly cut out; and 2. Set B includes more outdoor scarves and hoods, which often occlude the lower parts of the masks and the foreheads. These additional sources of heterogeneity caused the generator to produce inaccurate attention. On the other hand, the inputs for our GAN model are generated by warping mask images onto full faces, so the mask regions can be determined by comparing the full-face images with rule-based mask images pixel by pixel. Utilizing this pixel-by-pixel comparison to improve the misled attention became one of our improvement directions.

To achieve this, we created an extra “Non-Mask Change (NMC)” loss that calculates the L1 distances between the rule-based mask images and the realistic mask images for all pixels outside the mask regions. We minimize the sum of this L1 distance and other losses in AttentionGAN. The non-mask region for each rule-based mask image is calculated pre-training as a 256×256 boolean tensor, stored with the same

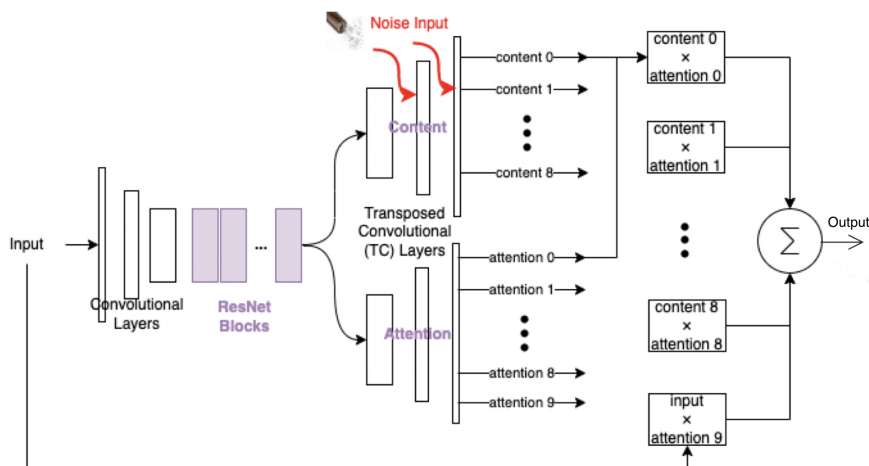


Figure 2: Generator architecture in AttentionGAN. Our noise inputs are depicted in red.

file name as the rule-based mask image but with a different extension name. During training, these tensor files are paired with their rule-based mask images and the generated realistic mask images for the NMC loss calculation. Note that we only calculate the loss at the training stage, so we do not have to get these tensor files for any generalized model test or usage once we finish training. Our model learns to automatically generate the attention mask instead of relying on ground truth attention masks beyond training.

4.2 Adding Noise

Inspired by StyleGAN (Karras et al., 2018), which takes random noise input to multiple generator layers for result diversity, we modified AttentionGAN to include zero-mean Gaussian noise input of a similar style. We first attempted noise input to the first content-generating TC layer following the nine ResNet blocks. We tested the noise-tuned model structure on parameters trained without noise and found that the generated images remained the same no matter how we amplified the noise. We concluded that the first TC layer assimilates noise and does not map noise input to identifiable features in the final output. Then, with one TC layer noise-tuned at a time, we tested different amplitudes of noise input on the other two content-generating TC layers, using the same model mentioned above. We found that noise added to either of the other two TC layers makes a difference to the model output, producing random noise pixels as shown in Figure 3.

Based on these tests, we trained our final model using a modified structure: the generator takes zero-mean Gaussian noise inputs to the last two content-generating layers, with standard deviations being 1 and 0.2 separately for the second and the last content-



Figure 3: Feed noise to the last two content-generating layers without training.

generating layers. We denoted our noise input location in Figure 2.




Figure 4: Uniform-colored masks generated at epoch 279.

Section 5.1 shows that models at two different epochs generated diverse colors after adding the noise input. In contrast, before adding the noise, our model output is in uniform color for any single checkpoint, as shown in Figure 4. Therefore, we conclude that the noise input effectively results in increased diversity.

Besides diversity, the noise input, together with the NMC loss, also reduced distortions to non-mask areas and stabilized training. Before the improvements, with the training process producing one sample per epoch, the samples did include faces mostly preserved from the input for consecutive epochs. However, drastically distorted and completely re-drawn faces, as in Figure 5, also last for epochs often, interweaving with the preserved faces. In contrast, all epoch samples stably retained input faces after applying the improvements, with only local changes fluctuating, proving the increased training stability. The remaining non-mask changes still exist, with examples shown in Figure 8, but both the extent of changes and the stableness of training were improved.



Figure 5: Models drastically distort inputs at two different epochs. Outputs like images 2 and 4 intermittently appear as the training progresses.

4.3 Transfer Learning and the Training Timeline

We used transfer learning from our trial-and-error experiments. We gradually improved datasets and methods in our training timeline, but we did not discard previous checkpoints. Instead, we believe that a checkpoint from training epochs with less ideal model settings and datasets is better than a random start.

Our initial parameters for the single-face mask generation task are from a different research topic: a multi-face image translator trained with the same AttentionGAN model. These initial weights were gained by 60 epochs of training on Face Detection Data Set and Benchmark (Jain and Learned-Miller, 2010) and MAFA data, with hyperparameters in the original AttentionGAN code set to $\lambda_A = 5$, $\lambda_B = 5$, and $\lambda_{identity} = 0.2$. From each of Fddb and MAFA, about 25,000 images were selected as training sets A and B separately. We will not elaborate on the rationale behind our choices since it belongs to other research. The information here is only for the reproduction of our result.

Within the single-face mask generation task itself, our first experiment used all the 9,517 face images extracted from the MAFA training set, which were annotated as 1) fully occluded (with occlusion degree equal to or higher than three), 2) of “simple” mask type, and 3) at least 60×60 in size. From the analysis of this first experiment, we found our training set A, i.e., the CMFD data, is much less diverse in orientations of faces than set B. To prevent set B’s additional orientation variations from confusing the model, we limited the image extraction to only facial orientations of “front,” “front left,” and “front right,” getting 8,938 images for set B. We downsampled set A to match the number. Our continued training used these 8,938-image datasets, with starting weights copied from the model at epoch 60 in our first experiment.

Then we noticed that faces in MAFA also involve pitch and roll rotations, which, unlike yaw rotations, are not annotated. On the other hand, CMFD mostly restricted pitch and roll rotations. Also, masks annotated as “simple” in set B are not always simple medical masks similar to CMFD. Cloth or gauze veils are also annotated as simple types. We believe limiting the source and destination datasets to have similar

variations in the aspects above may help the model focus on the target modifications, i.e., the masks. Therefore, out of the 8,938 images, we manually selected 1,597, which are limited in pitch and roll rotations with only light-colored medical masks.

Adopting 98 additional real-world masked faces from online for set B, we finalized our datasets with 1,695 images in each set as described in Section 3. After this, training and improvements, including in Sections 4.1 and 4.2, are all based on 1,695 real-world images and 1,695 CMFD images. The whole training timeline on the single face task is shown in Table 1.

The training model was updated along the way as we designed new improvements, but training epochs on old models were utilized and stacked together. We might have tried more clear-cut experiments if we had time, but this methodology has accelerated training and alleviated that our final datasets are small.

5 DISCUSSIONS

Output from training epochs showed that our model slowly converged after applying noise input and NMC loss. We carefully watched this trend and visually picked two better-performing checkpoints, checkpoints 313 and 476, from the later epochs for our test.

5.1 Improvements on Top of CMFD

Test results in Figure 6 show that, compared to CMFD inputs, epoch 313 provides a diversity of mask colors that match dataset B’s color distribution. It also shows better details than CMFD on various aspects:

- Fabric folds and resulting irregular mask region boundaries;
- Straps or their connecting points with the masks;
- More realistic lighting matching cheek curvatures;
- Visual effects of masks lifted by the nose bridges;
- More natural transitions from masks to faces.

In epoch 476, as shown in Figure 7, all the diversity and details mentioned here are rendered in even more powerful ways. Images such as the top-left and the bottom-right ones in Figure 7 even learned to partly put other occlusions, such as hand or veil, in front of the mask occlusion.

However, epoch 476 also produced more noise and distortions than epoch 313. We think this is due to overfitting the small training set. Red and white colors in the hair and forehead (Figure 8a) are likely caused by faces wearing not only masks but also Christmas hats appearing in our target training set repeatedly.

Table 1: Transfer learning from trial-and-error experiments. All epochs use Learning Rate 0.0002 and Lambda_identity 0.5.

Epochs	Λ_A / Λ_B	Noise Input	Non-Mask Change Loss	Dataset Size	Training Data Selection Restrictions Added to B
1~60	10/10	No	No	9,517	Simple mask, fully occluded
61~90	10/10	No	No	8,938	Front facing
91~140	8/8	No	No	1,695	No pitch/roll, light-colored medical mask
141~298	8/8	No	Yes	1,695	None
299~510	8/8	Yes	Yes	1,695	None



Figure 6: Results produced by the model at epoch 313. Input and output images are paired side by side.



Figure 7: Results produced by the model at epoch 476. Input and output images are paired side by side.

Patterns on the bottom parts of masks (Figure 8b) may be caused by a large portion of images with patterned scarves occluding the masks.

5.2 Comparison with IAMGAN

Both Geng et al. (2020) and our research used CycleGAN-based methods to turn full faces into

masked ones. Our differences include:

- IAMGAN uses a multi-layer identity loss, while our NMC loss is pixel-level only. They differ because IAMGAN aims to keep the person's identity after adding a mask, while we want to keep the images exactly the same except the mask region to facilitate both recognition and detection tasks.



Figure 8: Noisy output in epoch 476 test result. (a) Red and white blocks in hair. (b) Patterns on the mask.

- IAMGAN always predicts the mask regions, while we utilize ground truth mask regions during training and only predict it during testing.
- IAMGAN works on more diverse data, while we have pioneer work on constrained datasets.
- IAMGAN turns full faces directly into masked ones, while we require a pre-step and turn fake masks into more realistic ones after the pre-step.

Performance scores such as Frechet Inception Distance (FID) (Heusel et al., 2017) and Kernel Inception Distance (KID) (Bińkowski et al., 2018) are usually used to compare the fidelity of different synthesized datasets, but we lack a real-world masked face dataset as the baseline. Datasets highly similar to either IAMGAN’s or our training data would unfaithfully push one party’s score high. Therefore, we ran our model on some examples shown in the IAMGAN paper (Geng et al., 2020) and demonstrated a qualitative comparison in Figure 9.

The two models showed similar abilities to retain non-mask regions. Benefiting from the guidance of the superimposed fake masks, our model got more accurate nose bridge positions and occasionally more details such as fabric folds and connecting points between masks and straps. However, IAMGAN offered good fabric and lighting details in many cases, too, and it offered higher diversity in mask colors.

5.3 Potential Improvements

We believe a more thorough work on datasets would greatly benefit the results in the future. It would be best to simultaneously achieve the mutual similarity between sets A and B, the size, and the diversity.

Making the two datasets, i.e., A and B, more similar to each other, with masks being the only source of heterogeneity, is an alternative approach to improving attention learning and reducing distortions to non-mask face areas, complementing the extra loss func-




Figure 9: Comparing our model with IAMGAN.

tion in Section 4.1. As mentioned in Section 3, we have already limited our data to only a subset from MAFA and Unsplash.com. This step is exactly based on the consideration of limiting irrelevant heterogeneity. However, it resulted in a small dataset. While the steps described in Section 4.3 using larger datasets to get our initial weights before training on the small datasets helped alleviate the dataset size problem, getting a more considerable amount of quality data may provide further improvements.

Future work on datasets may also emphasize increasing mask color/type diversity, but the type diversity should be attempted together with increased model abilities in learning diverse mask shapes. The cycle loss in our model is better at dealing with point-to-point mapping with little shape-changing, so it may be insufficient for the shape diversity. Besides, we only need the A-to-B model, not the opposite direction. Therefore, the single-sided domain mapping proposed by Benaim and Wolf (2017), with its distance constraints substituting the cycle loss, could be one direction we consider together with an increased mask type diversity. If we retain the two-directional training architecture, the distance constraints may even be used together with the cycle loss.

Besides improving the data, the non-mask penalty loss itself can be improved in two different ways. First, besides calculating the loss based on the improper content change, we may compare the ground truth attention masks directly with the generator-produced attention masks, taking the differences between the two as an extra loss. Second, instead of using a binary tensor indicating whether each pixel is supposed to be changed or not, we may set a finer-weighted penalty that punishes pixel changes farther away from the mask more than those closer to the mask. Such a weighted penalty would allow more room for the model to create realistic details in the transition regions, for example, mask straps and fabric

fold. These improvements to the non-mask penalty will further increase the learning stability and reduce improper changes outside the mask.

6 CONCLUSIONS

We aimed at turning full face detection/recognition datasets into masked face datasets, supplementing the limited training data for masked face tasks. For this purpose, we proposed a two-step data augmentation method, utilizing Cabani et al. (2021)'s algorithm to warp mask images onto faces as a pre-step to an AttentionGAN-like model that generates more realistically masked faces. We applied multiple improvements to the GAN model training and verified their effectiveness through experimental results. Analyses of our final results showed that the two-step method provided noticeable improvements compared to using a rule-based method alone. Even with the latest advances of the rule-based method by Wang et al. (2021), we still expect an extra I2I step to render the rule-based results with more details, such as irregular region boundaries caused by fabric folds and straps. Our results are also comparable with state-of-the-art NN-only mask generation methods such as IAMGAN, with complementary details. For example, we produced lighting changes and mask stripes or their connecting points missing in IAMGAN results.

While our current model and the generated images can be used in masked face detection or recognition tasks, we have limitations, including patterned noise caused by overfitting small datasets, the remaining face distortions, and the lacking of diversity in mask color and type. Based on discussions about these limitations, we pointed out several directions to generate even better supplemental training data in the future.

REFERENCES

- Anwar, A. and Raychowdhury, A. (2020). Masked Face Recognition for Secure Authentication. *CoRR*, abs/2008.11104.
- Benaïm, S. and Wolf, L. (2017). One-Sided Unsupervised Domain Mapping. *CoRR*, abs/1706.00826.
- Bińkowski, M., Sutherland, D. J., Arbel, M., and Gretton, A. (2018). Demystifying MMD GANs. *arXiv preprint arXiv:1801.01401*.
- Cabani, A., Hammoudi, K., Benhabiles, H., and Melkemi, M. (2021). MaskedFace-Net – A Dataset of Correctly/Incorrectly Masked Face Images in the Context of COVID-19. *Smart Health*, 19:100144.
- Ge, S., Li, J., Ye, Q., and Luo, Z. (2017). Detecting Masked Faces in the Wild with LLE-CNNs. In *CVPR*, pages 426–434.
- Geng, M., Peng, P., Huang, Y., and Tian, Y. (2020). Masked Face Recognition with Generative Data Augmentation and Domain Constrained Ranking. In *ACM-MM*, pages 2246–2254.
- Guo, X., Wang, Z., Yang, Q., Lv, W., Liu, X., Wu, Q., and Huang, J. (2020). GAN-Based Virtual-to-Real Image Translation for Urban Scene Semantic Segmentation. *Neurocomputing*, 394:127–135.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Klambauer, G., and Hochreiter, S. (2017). GANs Trained by a Two Time-Scale Update Rule Converge to a Nash Equilibrium. *CoRR*, abs/1706.08500.
- Isola, P., Zhu, J., Zhou, T., and Efros, A. A. (2016). Image-to-Image Translation with Conditional Adversarial Networks. *CoRR*, abs/1611.07004.
- Jain, V. and Learned-Miller, E. (2010). Fddb: A Benchmark for Face Detection in Unconstrained Settings. Technical report, UMass Amherst.
- Karras, T., Laine, S., and Aila, T. (2018). A Style-Based Generator Architecture for Generative Adversarial Networks. *CoRR*, abs/1812.04948.
- Masi, I., Wu, Y., Hassner, T., and Natarajan, P. (2018). Deep Face Recognition: A Survey. In *SIBGRAPI*, pages 471–478.
- Montero, D., Nieto, M., Leskovský, P., and Aginako, N. (2021). Boosting Masked Face Recognition with Multi-Task ArcFace. *CoRR*, abs/2104.09874.
- Pang, Y., Lin, J., Qin, T., and Chen, Z. (2021). Image-to-Image Translation: Methods and Applications. *CoRR*, abs/2101.08629.
- Regmi, K. and Borji, A. (2018). Cross-View Image Synthesis using Conditional GANs. *CoRR*, abs/1803.03396.
- Sagonas, C., Antonakos, E., Tzimiropoulos, G., Zafeiriou, S., and Pantic, M. (2016). 300 Faces In-The-Wild Challenge: Database and Results. *Image Vis. Comput.*, 47:3–18.
- Singh, S., Ahuja, U., Kumar, M., Kumar, K., and Sachdeva, M. (2021). Face Mask Detection Using YOLOv3 and Faster R-CNN Models: COVID-19 Environment. *Multimed. Tools. Appl.*, 80(13):19753–19768.
- Tang, H., Liu, H., Xu, D., Torr, P. H., and Sebe, N. (2021). AttentionGAN: Unpaired Image-to-Image Translation using Attention-Guided Generative Adversarial Networks. *TNNLS*.
- Wang, C., Fang, H., Zhong, Y., and Deng, W. (2021). MLFW: A Database for Face Recognition on Masked Faces. *CoRR*, abs/2109.05804.
- Wang, X., Wang, K., and Lian, S. (2020a). A Survey on Face Data Augmentation for the Training of Deep Neural Networks. *Neural. Comput. Appl.*, 32(19):15503–15531.
- Wang, Z., Wang, G., Huang, B., Xiong, Z., Hong, Q., Wu, H., Yi, P., Jiang, K., Wang, N., Pei, Y., et al. (2020b). Masked Face Recognition Dataset and Application. *arXiv preprint arXiv:2003.09093*.
- Yang, W. and Jiachun, Z. (2018). Real-Time Face Detection Based on YOLO. In *ICKI2018*, pages 221–224.
- Zhu, J., Park, T., Isola, P., and Efros, A. A. (2017). Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. *CoRR*, abs/1703.10593.