A 3D MODEL-BASED APPROACH FOR FITTING MASKS TO FACES IN THE WILD

Je Hyeong Hong[*1], *Hanjo Kim*[*2], *Minsoo Kim*[1,3], *Gi Pyo Nam*[1],
Junghyun Cho[1], *Hyeong-Seok Ko*[2], *Ig-Jae Kim*[1,3]

[1]Korea Institute of Science and Technology   [2]Seoul National University
[3]University of Science and Technology

ABSTRACT

Face recognition now requires a large number of labelled masked face images in the era of this unprecedented COVID-19 pandemic. Unfortunately, the rapid spread of the virus has left us little time to prepare for such dataset in the wild. To circumvent this issue, we present a 3D model-based approach called *WearMask3D* for augmenting face images of various poses to the masked face counterparts. Our method proceeds by first fitting a 3D morphable model on the input image, second overlaying the mask surface onto the face model and warping the respective mask texture, and last projecting the 3D mask back to 2D. The mask texture is adapted based on the brightness and resolution of the input image. By working in 3D, our method can produce more natural masked faces of diverse poses from a single mask texture. To compare precisely between different augmentation approaches, we have constructed a dataset comprising masked and unmasked faces with labels called *MFW-mini*. Experimental results demonstrate WearMask3D[1] produces more realistic masked faces, and utilizing these images for training leads to state-of-the-art recognition accuracy for masked faces.

Index Terms— covid, mask, face, augmentation, 3dmm

1. INTRODUCTION

Many countries are imposing a strict mask-wearing policy in public places with aims to prevent transmission of the novel coronavirus through respiratory droplets. Consequently, face recognition in unconstrained environments has become a rather more difficult task both for humans and machines, each having to identify people through occlusions from their hairstyles, eyes and eyebrows only. While it is hoped we will one day return to our lifestyle before the COVID-19 pandemic, improving face recognition models to work with masked faces is a current challenge we are directly facing.

The state-of-the-art face recognition methods [1, 2, 3] are mostly, if not all, data-driven, which have been made possible by the availability of large-scale datasets of labelled face images in the wild, e.g. VGGFace [4, 5], CelebA [6] and CASIA-Webface [7] just to name a few. It is similarly anticipated a recognition model for masked faces will require a large number of labelled masked face images. In addition, the identities of masked faces would need to overlap with those of non-masked faces for the models to learn the correct representation of the individuals. Unfortunately, wearing a mask has only become a common global practice in recent months, and thus it still requires some time for such large-scale masked dataset with labels to be constructed and publicly released.

Meanwhile, to avoid lack of data from being the bottleneck of research in masked face recognition, a practical short-term solution has been to augment publicly available face datasets to create pseudo-real masked face images. Ud din et al. [10] randomly placed frontal mask images to frontal faces using Adobe Photoshop, but this involves a manual process. Cabani et al. [11] used 2D facial landmarks to fit 2D masks to faces, but the fitting is natural for near-frontal images only. Anwar and Raychowdhury [12] improved this by utilizing 3 textures (front, left and right) for each type of mask to handle side poses, but this requires additional efforts for each newly added mask and the fitting result is still inaccurate for extreme poses (see Fig. 5). Automatically fitting masks more naturally to faces of extreme poses serves as our main goal.
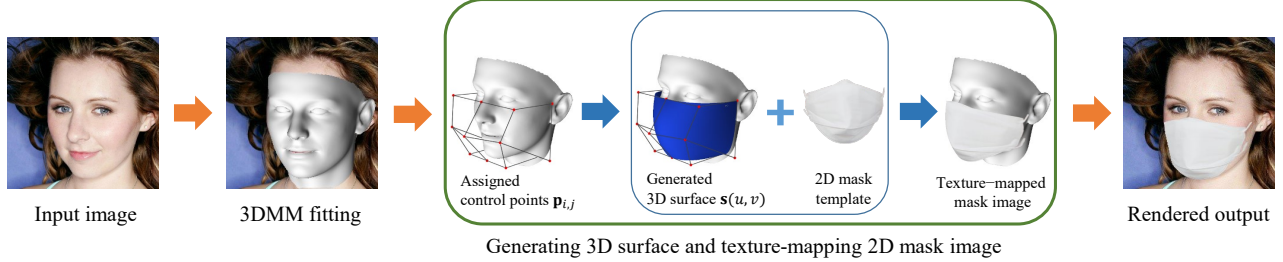
In this work, we propose a 3D model-based approach called *WearMask3D* for masking face images in the wild. Our method utilizes the well-known 3D morphable model (3DMM) [13, 8] to estimate the face surface, from which the mask surface is constructed and the corresponding mask texture warped. The output is formed by back-projecting the 3D mask to the input image then adjusting the mask brightness and resolution. The whole pipeline can be processed efficiently as 3DMM coefficients can be regressed using a neural network [14, 9] without iterative optimization.

The main contributions of this work are as follows:

+ *WearMask3D*, an efficient mask-fitting pipeline based on the 3D morphable model and head pose estimation to produce natural masked faces across a wide range of poses,

+ *Masked Faces in the Wild (MFW) mini*, a dataset of labeled masked and non-masked faces larger than previously reported and allowing various types of verification pairs, and

+ comparisons of WearMask3D against baseline and state-of-the-art augmentation method (MaskTheFace).

∗ Je Hyeong Hong and Hanjo Kim contributed equally to this work.
[1] https://github.com/jhh37/wearmask3d

Generating 3D surface and texture-mapping 2D mask image

| Input image | 3DMM fitting | Assigned control points $\mathbf{p}_{i,j}$ | Generated 3D surface $\mathbf{s}(u,v)$ | 2D mask template | Texture−mapped mask image | Rendered output |

Fig. 1. An overview of our *WearMask3D* pipeline. We first fit a 3D morphable model (Basel Face Model '09 [8]) using a pretrained off-the-shelf neural 3DMM model regressor called 3DDFA [9]. Then, we compute 15 control points at predefined locations with respect to the 3D model. From these control points, we fit a mask surface based on non-uniform rational B-spline (NURBS) and warp the mask texture onto it. The 3D warped mask is then projected back to 2D for image rendering.

2. PROPOSED MASK AUGMENTATION METHOD

We now illustrate each stage of our pipeline from Fig. 1.

2.1. Fitting a 3D morphable model (3DMM)

The origin of 3DMM dates back to the seminal work of Blanz and Vetter [13], in which a 3D face comprising N 3D points, $\mathbf{S} \in \mathbb{R}^{3 \times N}$, is expressed as a weighted sum of low-rank basis shapes. These shapes are grouped into those based on the person's identity and others on the person's expression. i.e.

$$\mathbf{S}(\boldsymbol{\alpha}_{\mathrm{id}}, \boldsymbol{\alpha}_{\mathrm{exp}}) = \bar{\mathbf{S}} + \sum_{j=1}^{N_{\mathrm{id}}} \alpha_{\mathrm{id}}^j \mathbf{A}_{\mathrm{id}}^j + \sum_{k=1}^{N_{\mathrm{exp}}} \alpha_{\mathrm{exp}}^k \mathbf{A}_{\mathrm{exp}}^k, \quad (1)$$

where $\bar{\mathbf{S}} \in \mathbb{R}^{3 \times N}$ is the mean shape, $\mathbf{A}_{\mathrm{id}}^j \in \mathbb{R}^{3 \times N}$ is the j-th identity basis, $\mathbf{A}_{\mathrm{exp}}^k \in \mathbb{R}^{3 \times N}$ is the k-th expression basis, and α_{id}^j and α_{exp}^k are the scalar weights for the respective bases.

One often uses a pretrained basis model ($\{\mathbf{A}_{\mathrm{id}}^j\}, \{\mathbf{A}_{\mathrm{exp}}^k\}$), and we also use Gerig et al.'s Basel face model 2009 [15].

Assuming the face is relatively flat and far away from the camera, each 3DMM point can be projected to 2D using the weak-perspective camera model [16] as follows:

$$V(\boldsymbol{\alpha}_{\mathrm{id}}, \boldsymbol{\alpha}_{\mathrm{exp}}, \mathbf{p}) = f \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \mathbf{R}(\mathbf{q})\mathbf{S} + \mathbf{t}\mathbf{1}^\top, \quad (2)$$

where $f \in \mathbb{R}$ is the focal length, $\mathbf{R} \in SO(3)$ is the rotation matrix, $\mathbf{t} \in \mathbb{R}^2$ is translation, $\mathbf{q} \in \mathbb{R}^4$ are the quaternions representing $\mathbf{R}$ and $\mathbf{p}$ denotes the camera variables $[\mathbf{f}, \mathbf{q}, \mathbf{t}]$.

We employ Zhu et al.'s end-to-end pretrained network called 3DDFA [9] to jointly estimate $\mathbf{p}$, $\boldsymbol{\alpha}_{\mathrm{id}}$ and $\boldsymbol{\alpha}_{\mathrm{exp}}$ from the input image. In this case, the number of identity bases (N_{id}) and expression bases (N_{exp}) are 40 and 10 respectively, with the number of dense 3DMM points N set to 53,490.

2.2. Generating 3D mask surface

After fitting the 3DMM, we generate a 3D mask surface on which the 2D mask template is mapped. For this purpose, We use the 3D non-uniform rational B-spline (NURBS) surface, which is defined for $0 \le u \le 1$ and $0 \le v \le 1$ as the tensor product of a horizontal and vertical NURBS curves:

$$\mathbf{s}(u,v) = \sum_{i=1}^{K} \sum_{j=1}^{L} \Phi_{i,j}(u,v)\mathbf{c}_{i,j}, \quad (3)$$

where $\Phi_{i,j}(u,v)$ is the rational basis function defined as

$$\Phi_{i,j}(u,v) := \frac{B_{i,n}(u)B_{j,m}(v)w_{i,j}}{\sum_{p=1}^{K} \sum_{q=1}^{L} B_{p,n}(u)B_{q,m}(v)w_{p,q}}, \quad (4)$$

$\mathbf{c}_{i,j} \in \mathbb{R}^3$ is the (i,j)-th control point, K and L are the number of control points for the horizontal and vertical NURBS curves respectively, $B : \mathbb{R} \to \mathbb{R}$ represents the B-spline basis function, n and m are the polynomial degrees of the NURBS curves, and $w_{i,j}$ is the weight of the control point (i,j).

In our implementation based on NURBS-Python [17], we assign 15 control points (K=3 and L=5) for the NURBS surface, second order polynomials for both curves (m=2, n=2) and equal unit weights across all control points. The location of each control point is automatically calculated based on the positions of several 3D facial landmarks such as tip of the nose, cheeks, bottom of the lips in the 3DMM. Nevertheless, we did go through some trial-and-errors to find the optimal position of each control point for realistic mask-fitting. The control points are set differently for different types of masks.

2.3. Mapping the mask texture

We have obtained several mask textures by gathering high-quality redistributable images of real masked faces in frontal view, cropping the mask region and resizing the image. The deleted region is left transparent as shown in Fig. 2.

From the frontal (opposite-the-face) viewpoint, we assign $(u,v) = (0,0)$ to the bottom-left vertex of the 3D mask surface and $(u,v) = (1,1)$ to the top right vertex. Then, the uv-coordinates of approximately 1000 mask surface vertices are obtained through interpolation.

We then arbitrarily select one mask template from Fig. 2 (with different probabilities) and map its texture to the 3D mask surface. We apply bilinear interpolation to fill in the texture between adjacent surface points. Since the 3D mask surface smoothly covers the lower half of the face, the mapped mask surface looks realistic in 3D as shown in Fig. 1.
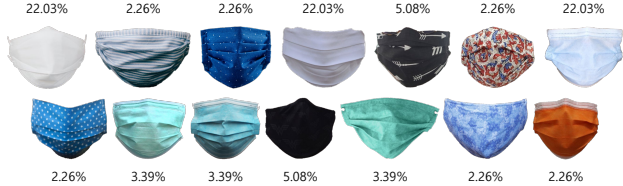
Fig. 2. A list of utilized mask templates with probabilities

22.03% 2.26% 2.26% 22.03% 5.08% 2.26% 22.03%

2.26% 3.39% 3.39% 5.08% 3.39% 2.26% 2.26%

(a) Lower brightness (b) Downsampled mask

Fig. 3. Effects of controlling mask brightness and resolution

2.4. Rendering and post-processing

Finally, we render the texture-mapped 3D mask surface on the 2D input image using the camera pose obtained using 3DDFA in Sec. 2.1. So long as the 3DMM is accurately fitted on the input image, the rendered mask also fits the face naturally. We then refine the mask rendering for quality improvement.

Brightness control   Just rendering the mask with equal brightness across all images sometimes produces visually unsatisfactory results largely due to significant variations of illumination across images. To mitigate this issue, we estimate the brightness of the input image by averaging the intensities (obtained by converting RGB to greyscale and dividing by 255) of the entire pixels. If the measured brightness is lower than our pre-determined threshold of 0.4, we reduce brightness of the mask ρ as $0.7 + 0.75 * \min(0.4, \rho)$, where the numbers are deduced from empirical trials.

Resolution adjustment   We also adjust the mask resolution depending on the sharpness of the input image. Variance of the Laplacian has been known as an effective concept to estimate the image sharpness [18]. We apply a conventional 3×3 Laplacian kernel to each pixel of the input image and calculate its variance σ^2. If $\sigma^2 < 400$, we regard the input as a low-quality image and reduce its resolution by the factor of $\min(10, 400/\sigma^2)$, where the values are again empirically set.

3. MASKED FACES IN THE WILD WITH LABELS

So far, the largest public dataset of real labelled masked faces has been MFR2 [12], comprising 171 masked and 98 non-masked faces from 53 identities with 848 test pairs for verification. We found this to be small for precise comparisons, and decided to construct a larger dataset for testing.

We gathered 3,000 images of 300 individuals from the internet, forming a class-balanced dataset by collecting 5 masked and 5 non-masked images per identity. This is to incorporate various matching situations, e.g. between masked and non-masked faces as well as masked vs masked. We

Fig. 4. Images from our class-balanced *MFW-mini* dataset.

Table 1. Fréchet inception distances (FIDs) achieved by different mask augmentation methods on MFR2 [12] and MFW-mini (lower is better). The tested feature dimensions are set to be below the number of images available in each dataset.

Augmentation method	MFR2	MFW-mini (Sec. 3)		
	64-d	64-d	192-d	768-d
No aug. (non-masked)	0.665	0.556	3.114	0.954
MaskTheFace [12]	0.468	0.410	2.011	0.304
WearMask3D	**0.398**	**0.153**	0.856	**0.189**
w/o brightness adj.	0.404	0.154	**0.854**	**0.189**
w/o resolution adj.	0.399	0.203	1.171	0.200

cropped and aligned images using the bounding box and 5 facial landmarks extracted from RetinaFace [19]. About 3% of images were incorrectly cropped or aligned, and these were re-cropped by manually extracting the 5 facial landmarks. The dataset is named *Masked Faces in the Wild (MFW) mini*, and some of the constituent images are shown in Fig. 4.

While the number of images is still relatively small compared to that of LFW [20], this allows 13,500 genuine pairs to be formed (compared to ≈400 for MFR2) including 3,000 genuine pairs of masked images, allowing the associated verification accuracies to be more reliable. We plan to release either the actual dataset or the URLs of the utilized images.

4. EXPERIMENTAL RESULTS

Our approach is compared against baseline (no augmentation) and state-of-the-art mask augmentation method (MaskTheFace [12]) on two aspects, namely the quality of masked images by calculating the Fréchet inception distance (FID) [21] and the usefulness of mask augmentation for training by measuring the accuracy improvement in masked face recognition.

4.1. Image quality analysis

Fig. 5 shows representative mask augmentation results across full range of horizontal poses (0–90°). This shows visually our method yields more natural mask fitting for profile faces.

We also compared the FIDs of different approaches on MFR2 and MFW-mini. FID is the Fréchet distance between the distribution of the Inception v3 features from real samples and those from fake ones (lower is better). While FID is frequently used to compare the quality of fake images produced

Fig. 5. A visual comparison of the masked faces. The top row contains original images from VGGFace2 [5], while the middle and bottom rows show outputs from MaskTheFace [12] and ours respectively. The horizontal pose varies from 0° to 90°.

Table 2. Verification accuracies (%@EER) achieved by different mask augmentation schemes. N+N denotes pairs of non-masked images, M+N pairs of masked and non-masked images, M+M pairs of masked images, and Mix. mixed pairs. "Bottom-out" overwrites 0s to the bottom 45% of the image.

Augmentation Method	LFW -	MFR2 Mix.	MFW-mini (Sec. 3)		
			N+N	M+N	M+M
No augment.	**97.3**	90.1	**96.7**	87.3	89.7
Bottom-out	96.3	93.9	96.1	89.5	91.1
MaskTheFace	96.6	94.6	95.9	90.3	92.1
WearMask3D	96.6	**95.8**	95.9	**90.9**	**92.6**

by a deep generator, we would like to stress that FID itself does not enforce any constraint on the nature of fake images, i.e. a fake image can be produced by a 3D rendering process.

On each dataset, we fitted a mask to each non-masked face using MaskedTheFace and our method. Then, the FID between the real masked faces and the mask-fitted images was calculated for each method. For the baseline comparison, we also provide the FID between the real masked images and non-masked images.

While the FID is usually computed on 2,048-d Inception v3 features, we had to use lower dimensional features (64, 192 and 768) as the number of non-masked images is smaller than 2,048 for MFR2 (98) and MFW-mini (1,500). Table 1 shows WearMask3D yields lower FIDs across all tested Inception v3 features, indicating more realistic mask augmentation.

4.2. Masked-face recognition accuracy

We compared the verification accuracies by training a network with mask augmented images from different methods.

As in [12], we created a class-balanced random subset of VGGFace2 [5] for training, with 42 images for each of the 8,631 identities. For each augmentation method, we fitted masks to 362,502 faces and used them with the original non-masked subset for training over 25 epochs. The "no augmentation" setting was trained with the non-masked subset only but for 50 epochs to compensate for halved number of training images. Across all settings, we used the ResNet-50 backbone, and applied basic softmax loss for training.

We set the initial learning rate to 10^{-1}, and slowed the learning rate by factor of 10 at specified epochs (16, 28, 40

for "no augmentation" and 8, 14, 20 for others). We set momentum to 0.9 and weight decay to 5e-4 as in [3]. We used batch size of 256 for "no augmentation" and 512 (256 non-masked + 256 masked) for others, distributed over 8 Nvidia RTX 2080 Ti GPUs.

For evaluation, we utilized authors-provided 848 test pairs for MFR2, and created 6,000 unique test pairs (50% genuine) for each MFW-mini setting. For detailed analysis, we varied types of tested face pairs, namely non-masked (N+N), masked and non-masked (M+N), masked (M+M). In each pair, we extracted 2,048-d feature vectors and calculated the angle between them. For each method on each setting, we deduced the acceptance threshold ϕ at which the equal error rate holds (FAR=FRR), and reported the sum of true positives and true negatives over the total number of test pairs.

Table 2 shows WearMask3D achieves state-of-the-art recognition accuracy when test pairs include real masked faces. On the other hand, the network with no augmentation performs best for non-masked face pairs, implying better network training is required to cover all verification cases.

5. CONCLUSION

In this work, we have proposed *WearMask3D*, an automatic 3D model-based approach for fitting masks to face images in the wild. We have shown that incorporating the 3D morphable model allows more naturally mask fitted faces even with extreme poses. We have also demonstrated the recognition accuracy of masked faces can be improved by training with masked faces generated from our approach. As a by-product, we have created *MFW-mini*, a dataset of 300 identities each with 5 masked and 5 non-masked faces in the wild that can be used for testing. We have released the relevant code aims to facilitate future research.

6. REFERENCES

[1] Florian Schroff, Dmitry Kalenichenko, and James Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 1

[2] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song, "Sphereface: Deep hypersphere embedding for face recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. 1

[3] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 1, 4

[4] Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman, "Deep face recognition," in *Proceedings of the British Machine Vision Conference (BMVC)*. September 2015, pp. 41.1–41.12, BMVA Press. 1

[5] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "Vggface2: A dataset for recognising faces across pose and age," in *2018 13th IEEE International Conference on Automatic Face Gesture Recognition (FG 2018)*, 2018, pp. 67–74. 1, 4

[6] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang, "Deep learning face attributes in the wild," in *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015. 1

[7] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z. Li, "Learning face representation from scratch," 2014. 1

[8] P. Paysan, R. Knothe, B. Amberg, S. Romdhani, and T. Vetter, "A 3d face model for pose and illumination invariant face recognition," in *Proceedings of the 6th IEEE International Conference on Advanced Video and Signal based Surveillance (AVSS) for Security, Safety and Monitoring in Smart Environments*, Genova, Italy, 2009, IEEE. 1, 2

[9] X. Zhu, X. Liu, Z. Lei, and S. Z. Li, "Face alignment in full pose range: A 3d total solution," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 1, pp. 78–92, 2019. 1, 2

[10] N. Ud Din, K. Javed, S. Bae, and J. Yi, "A novel gan-based network for unmasking of masked face," *IEEE Access*, vol. 8, pp. 44276–44287, 2020. 1

[11] Adnane Cabani, Karim Hammoudi, Halim Benhabiles, and Mahmoud Melkemi, "Maskedface-net – a dataset of correctly/incorrectly masked face images in the context of covid-19," *Smart Health*, vol. 19, pp. 100144, 2021. 1

[12] Aqeel Anwar and Arijit Raychowdhury, "Masked face recognition for secure authentication," 2020. 1, 3, 4

[13] Volker Blanz and Thomas Vetter, "A morphable model for the synthesis of 3d faces," in *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, USA, 1999, SIGGRAPH '99, pp. 187—194, ACM Press/Addison-Wesley Publishing Co. 1, 2

[14] X. Yin, X. Yu, K. Sohn, X. Liu, and M. Chandraker, "Towards large-pose face frontalization in the wild," in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 4010–4019. 1

[15] T. Gerig, A. Morel-Forster, C. Blumer, B. Egger, M. Luthi, S. Schoenborn, and T. Vetter, "Morphable face models - an open framework," in *2018 13th IEEE International Conference on Automatic Face Gesture Recognition (FG 2018)*, 2018, pp. 75–82. 2

[16] Richard Hartley and Andrew Zisserman, *Multiple View Geometry in Computer Vision*, Cambridge University Press, USA, 2 edition, 2003. 2

[17] Onur Rauf Bingol and Adarsh Krishnamurthy, "NURBS-Python: An open-source object-oriented NURBS modeling framework in Python," *SoftwareX*, vol. 9, pp. 85–94, 2019. 2

[18] J. L. Pech-Pacheco, G. Cristobal, J. Chamorro-Martinez, and J. Fernandez-Valdivia, "Diatom autofocusing in brightfield microscopy: a comparative study," in *Proceedings 15th International Conference on Pattern Recognition. ICPR-2000*, 2000, vol. 3, pp. 314–317 vol.3. 3

[19] Jiankang Deng, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou, "Retinaface: Single-shot multi-level face localisation in the wild," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 3

[20] Gary B. Huang, Manu Ramesh, Tamara Berg, and Erik Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," Tech. Rep. 07-49, University of Massachusetts, Amherst, October 2007. 3

[21] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter, "Gans trained by a two time-scale update rule converge to a local nash equilibrium," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, NIPS'17, pp. 6629—-6640. 3