

APB2FACE: AUDIO-GUIDED FACE REENACTMENT WITH AUXILIARY POSE AND BLINK SIGNALS

Jiangning Zhang¹, Liang Liu¹, Zhucun Xue², Yong Liu¹

¹Zhejiang University, Hangzhou, China

²Wuhan University, Wuhan, China

ABSTRACT

Audio-guided face reenactment aims at generating photorealistic faces using audio information while maintaining the same facial movement as when speaking to a real person. However, existing methods can not generate vivid face images or only reenact low-resolution faces, which limits the application value. To solve those problems, we propose a novel deep neural network named *APB2Face*, which consists of *GeometryPredictor* and *FaceReenactor* modules. *GeometryPredictor* uses extra head pose and blink state signals as well as audio to predict the latent landmark geometry information, while *FaceReenactor* inputs the face landmark image to reenact the photorealistic face. A new dataset *AnnVI* collected from YouTube is presented to support the approach, and experimental results indicate the superiority of our method than state-of-the-arts, whether in authenticity or controllability.

Index Terms— face reenactment, audio, generative adversarial neural network, information fusion

1. INTRODUCTION

Audio and image are two commonly used signal transmission modes by humans, which both send overlapped conversation contents and rich emotional expression. When only hearing a familiar voice of a human, the subtle facial movements and the emotions of the people who is speaking could be imagined, especially for the mouth area. It is significant to study audio-to-face task that aims to generate the face image using audio as input, and this technique can be applied in many situations, such as virtual announcer, film-making, game, etc.

Recently, many works focus on solving the audio-to-face task and have made some achievements. The methods mainly fall into two categories: parameter prediction and pixel-level generation. The former usually predict parameters of the pre-defined model, which can control the face model in a post-processing way. The images generated in this way are of high quality, but it is time-consuming and does not work in real-time. Pixel-level generation methods generally introduce the encoder to get a latent vector and then reenact face using the generative adversarial neural network (GAN) [1], which require no post-processing and can run in an end-to-end manner.

Most recently, Tae-Hyun *et al.* [2] propose the S2F model to predict the frontal face from the audio clip, and Amanda *et al.* [3] successfully reenact the special person by Wav2Pix model. However, nearly all recent works can not generate high-quality faces only using audio as input and ignore other attributes such as head pose and eye blink, which results in stiff facial expressions. So in order to generate more photorealistic images, audio-independent control signals over head pose and eye blink are required.

To solve aforementioned problems, we propose a novel *APB2Face* structure as shown in Figure 1. *GeometryPredictor* module inputs audio, pose, and blink signals to regress latent *landmark* that contains the facial geometry information (Thus we can use landmark image as input rather than a vector), and then *FaceReenactor* reenacts the target image using the generated landmark. Such a two-stage design can not only generate photorealistic faces but also fuse pose and blink signals into audio. Our main contributions are as follows:

- i) A novel *APB2Face* structure is proposed to generate photorealistic faces from audio, pose, and blink signals.
- ii) A new dataset *AnnVI* is created, which contains 6 announcers with synchronized audio, image, pose, landmark, and blink information.
- iii) Abundant experiments show that our approach can generate high-quality faces.

2. RELATED WORKS

Generative Adversarial Networks. GAN method generally consists of two adversarial models: a generative model G and a discriminative model D . The generator G has to capture the data distribution and generate the realistic image, while the discriminator D is used to discriminate generated fake images from G and real images, in which way they can enhance each other's performance by playing a min-max game. Specifically, the generator G with parameter θ_g maps input z from a prior distribution $p_z(z)$ to data space as $G(z; \theta_g)$, and the discriminator D with parameter θ_d outputs a probability value representing how much the input is from fake data space $G(z; \theta_g)$ or real data space x . The value function $V(G)$ and $V(D)$ are as follows:

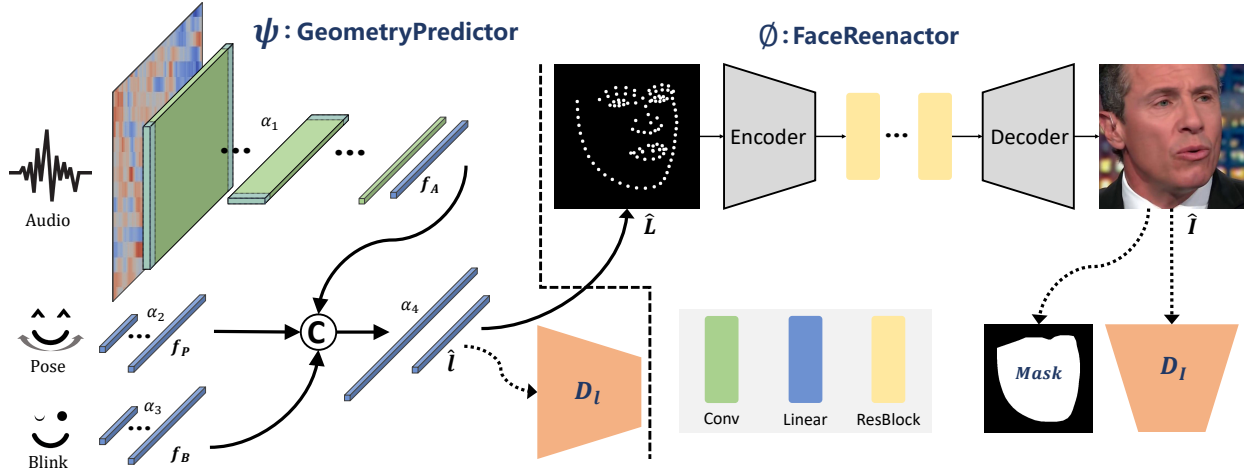


Fig. 1. Overview of the proposed APE2Face network, which consists of a geometry predictor ψ and a face reenactor ϕ . The ψ first encodes audio, head pose, and eye blink information as f_A , f_P , and f_B , and then fuses them to regress the latent landmark geometry \hat{L} . The discriminator D_I is used to help boost performance and $\alpha_i (i = 1, 2, 3, 4)$ indicate network parameters. The ϕ leverages the transformed landmark image \hat{L} as input to reenact target face image \hat{I} , whose characteristics should match the input audio, pose, and blink information. The discriminator D_I helps to generate photorealistic images in the training phase.

$$\min_G V(G) = \mathbb{E}_{z \sim p_z(z)} [\log(D(G(z))) - 1]. \quad (1)$$

$$\min_D V(D) = \mathbb{E}_{z \sim p_z(z)} [\log(D(G(z)))] + \mathbb{E}_{x \sim p_{data}(x)} [\log(D(x) - 1)], \quad (2)$$

Since Ian *et al.* first proposed generative adversarial network [1] method in image generation task, many excellent jobs have come up [4, 5, 6, 7, 8]. Considering the model size and the generation effect, we modified Pix2Pix model with a landmark input as our *FaceReenactor*.

Face Reenactment via Audio. Some recent works reenact face by predicting parameters of the predefined face model [9, 10, 11]. Tian *et al.* [10] directly use audio to regress blendshape parameters and then drive the predefined model. Cudeiro *et al.* [11] design a VOCA model to drive the static 3D template. But they are not real-time for requiring post-processing, so another major approach is to directly generate pixel-level face images [12, 13, 3, 14]. The state-of-the-art method [3] successfully reenacts faces of the special person by Wav2Pix model, which introduces a speech-conditioned GAN architecture. However, the quality of the generated images is not particularly good by this method, and it can not control generated attributes such as head pose and eye blink.

3. ANNVI DATASET

Consider that the current approaches usually use only audio information without other signals such as head pose and eye blink, we propose a new *AnnVI* dataset with additional head pose and eye blink annotations which can complement latent

head information that the audio cannot control. The dataset contains six announcers (three man and three women) and 23790 frames totally with corresponding audio clip, head pose, eye blink, and landmark information.

Image processing: To crop out the sequence of faces from videos, we first extract video frames from original videos that are collected from YouTube at 1080P resolution, and then landmark geometry and head pose information of each frame are detected by [15]. Notice that we have manually checked the detection results of each image to ensure the reliability of the labeled information. After acquiring landmarks, we crop each face using its 1.4x minimum outer square and resize them to 256×256 resolution.

Audio processing: In this paper, all audios are firstly re-sampled to 44.1kHz and then we extract mel-frequency cepstral coefficients (MFCCs) [16] features using the same method as [9].

Pose and blink: Head pose information is detected by [15] that contains yaw, pitch, and roll angles. In *AnnVI* dataset, the value ranges of three components are $-0.354 \sim 0.196$, $-0.367 \sim 0.379$, and $-0.502 \sim 0.509$ respectively (radian). As for eye blink, we define it as the height of the eye divided by the width, where the calculation scale is performed on the normalized 256×256 resolution.

4. METHOD

Directly using a vector to reenact face generally limits the resolution and authenticity of the generated images. So as shown in Figure 1, we propose a novel APB2Face structure, which consists of *GeometryPredictor* and *FaceReenactor*, to generate photorealistic pixel-level face from the audio input, along

with additional head pose and eye blink control signals. **Geometry predictor:** The geometry predictor module contains three paths to extract various input information: audio, head pose, and eye blink. For audio, we first use 5 convolutional layers to extract individual feature information and then use 5 convolutional layers to fuse features in the time dimension to get the feature \mathbf{f}_A . For head pose and eye blink, 4 and 3 linear layers are used to extract features \mathbf{f}_P and \mathbf{f}_B . After that, we concatenate them and use 2 consecutive linear layers to predict finally landmark geometry. The process can be denoted as:

$$\begin{aligned} \hat{\mathbf{I}} &= \psi_{\alpha_4}([\mathbf{f}_A, \mathbf{f}_P, \mathbf{f}_B]) \\ &= \psi_{\alpha_4}([\psi_{\alpha_1}(Audio), \psi_{\alpha_2}(Pose), \psi_{\alpha_3}(Blink)]), \end{aligned} \quad (3)$$

Note that the predictor is designed in a lightweight way, which can run in real-time: 260 FPS in CPU (with batch size 1, using i7-8700K @ 3.70GHz) and 823 FPS in GPU (with batch size 1, using 2080Ti).

During the training phase, an additional discriminator D_I (containing 7 linear layers) is applied to provide stronger supervision beside L1 loss, which is designed to judge real or fake of generated landmarks. So the overall loss function $\mathcal{L}_{predictor}$ is defined as:

$$\mathcal{L}_{predictor} = \lambda_1 \mathcal{L}_{L_1} + \lambda_2 \mathcal{L}_{D_I}, \quad (4)$$

Where \mathcal{L}_{L_1} and \mathcal{L}_{D_I} indicate L1 loss and discriminator D_I loss respectively, and $\lambda_1 = 100, \lambda_2 = 0.1$ in all experiments. In this way, we can reduce the L1 loss of average generated landmarks from 1.500 ± 0.064 pixels to 0.666 ± 0.034 , which greatly boosts the performance.

Face reenactor: Given the predicted landmark $\hat{\mathbf{I}}$, we first plot it to the binary image according to its point coordinates, and then the resized landmark image $\hat{\mathbf{L}}$ (with 256*256 resolution) is fed into the *FaceReenactor*. Each reenactor is associated with a special identity, which means that the input landmark image provides geometric information while the network itself contains the appearance information of the person. The process can be denoted as:

$$\hat{\mathbf{I}} = \phi(\hat{\mathbf{L}}), \quad (5)$$

Note that the reenactor can generate photorealistic high-resolution faces (256*256 resolution) and run in real-time: 75 FPS in GPU under the same setting with the predictor. During the training stage, in addition to L1 and GAN losses [5], we propose another mask loss function $\mathcal{L}_{L_{mask}}$. In detail, we first fill a black image with white in the smallest external polygon of the landmark and then expand the white area with the expansion algorithm to cover the face. Then L1 loss is applied to the white mask area, because we want to increase the weight of the face area that can help to generate the more photorealistic face. The overall loss function $\mathcal{L}_{reenactor}$ is defined as:

$$\mathcal{L}_{reenactor} = \lambda_1 \mathcal{L}_{L_1} + \lambda_2 \mathcal{L}_{L_{mask}} + \lambda_3 \mathcal{L}_{D_I}, \quad (6)$$

Where \mathcal{L}_{L_1} , $\mathcal{L}_{L_{mask}}$, and \mathcal{L}_{D_I} indicate L1 loss, mask loss, and discriminator D_I loss [5] respectively, and $\lambda_1 = 100, \lambda_2 = 100, \lambda_3 = 1$ in all experiments.

5. EXPERIMENTS

In this section, we elaborate on the training details and show some experimental results to prove the effectiveness of the proposed method.

Implementation details. For the geometry predictor, we use Adam [17] optimizer with $\beta_1 = 0.99, \beta_2 = 0.999$, and the learning rate is set to $3e^{-4}$. We train the module for 1000 epochs with batch size 32. For the discriminator D_I , we design 7 linear layers and the training setting is the same as the predictor. For the face reenactor, we use Adam optimizer with $\beta_1 = 0.5, \beta_2 = 0.999$, and the learning rate is set to $2e^{-4}$. We train the module for 100 epochs with batch size 16 in a single GPU. For the discriminator D_I , we use PatchGAN proposed in [5] and the training setting is the same as the reenactor.

Qualitative results. To visually demonstrate the effectiveness and superiority of our approach, we conduct some qualitative experiments as shown in Figure 2. The top half shows high-quality generation results of six persons, in which the input information of each identify is from himself/herself. To further evaluate the generalization capability of our model, input signals (audio, pose, and blink information) from $P1-P6$ are used to generate the identity $P1$ (Because of the limited space, we only show the results of $P1$), as shown in the bottom half in Figure 2. The generated faces match original faces in facial movements (especially mouth area), which indicates our approach can use audios from other persons to drive the given person.

Quantitative results. As shown in the first two rows of the Table 1, we choose generally used SSIM [18] and FID [19] metrics to quantitatively evaluate our approach, where results indicate that our model can generate photorealistic face images. The bottom four rows are evaluation results on our new defined *Detection Rate* (DR), *Average Landmark Error* (ALE), *Average Pose Error* (APE), and *Average Blink Error* (ABE) metrics, which are calculated on 1000 generated images for each identity. The results show that our method can reenact faces with an extremely high detection rate and low errors in landmark, pose, and blink information.

Comparison with the state-of-the-art. We further conduct a comparison experiment with the state-of-the-art *Wav2Pix* in the same person (We use the video with the same identity as the original paper in this experiment, for the original author not providing the original data). As shown in Figure 3, the result shows the superiority of our model no matter in vision (more photorealistic image generation) or metric evaluations (higher SSIM score and lower FID score).

Decoupling experiments. We also conduct a decoupling experiment to test the disentanglement of the audio, pose, and blink input information. As shown in Figure 4, the first four rows are generation results with only yaw, pitch, roll, and



Fig. 2. Face reenactment results of our proposed method. The top half is generated from individual audio input with pose and blink signals, while the bottom half is generated from the information of other persons. P1 to P6 represents different identities.

Table 1. Quantitative evaluations of our proposed approach on six persons (P1-P6) in AnnVI dataset. The first two rows are SSIM and FID metric evaluation results, while the bottom four rows are evaluation results on our defined DR, ALE, APE, and ABE metrics. The upward arrow indicates that the larger the value, the better the performance, and vice versa.

Metric	P1 (6400)	P2 (6110)	P3 (3450)	P4 (3500)	P5 (1230)	P6 (3100)	Avg
SSIM \uparrow	0.764	0.843	0.879	0.786	0.761	0.758	0.799
FID \downarrow	4.600	6.514	5.743	14.960	17.408	21.944	11.862
DR(%) \uparrow	98.9	98.8	99.1	98.8	98.7	98.5	98.8
ALE \downarrow	1.237	1.119	1.303	1.587	1.464	1.861	1.429
APE \downarrow	0.0168	0.0143	0.0152	0.0250	0.0221	0.0233	0.0195
ABE \downarrow	0.0356	0.0370	0.0371	0.0502	0.0443	0.0437	0.0413




Fig. 3. Comparison with the state-of-the-art method. The results of the first row are generated from Wav2Pix [3], while the second row is ours.

blink information changes respectively, and results indicate that our approach can well disentangle each component of the pose and blink signals, which gives us more room to control the generated image properties. Results of the last row are generated with only audio signals (pose and blink are set to 0) like other methods [2, 3], but the generated image quality of our method is visually better.

6. CONCLUSIONS

A novel *APB2Face* structure is proposed in this paper, which can use multi-information as control signals to reenact pho-



Fig. 4. Decoupling experiments of the pose and the blink information. Images of the first four rows are under only one control variable, while the last row contains all variables.

to-realistic faces. Specifically, we first design a *GeometryPredictor* to regress the landmark geometry information from the audio, pose, and blink signals, and then we use *FaceReenactor* to reenact target faces that match the input information. To verify the effectiveness of our approach and ensure the high quality of the results, we propose a new dataset named AnnVI which contains six announcers.

In the future, we will further extend our dataset as a stronger benchmark for the audio-to-face task. In the meantime, a more powerful structure can be used to boost *FaceReenactor* performance, so as to generate more photorealistic images.

Acknowledgements

This work is supported by the National Natural Science Foundation of China under Grant 61836015, 61771193.

7. REFERENCES

- [1] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, “Generative adversarial nets,” in *NeurIPS*, 2014, pp. 2672–2680.
- [2] Tae-Hyun Oh, Tali Dekel, Changil Kim, Inbar Mosseri, William T Freeman, Michael Rubinstein, and Wojciech Matusik, “Speech2face: Learning the face behind a voice,” in *CVPR*, 2019, pp. 7539–7548.
- [3] Amanda Duarte, Francisco Roldan, Miquel Tubau, Janna Escur, Santiago Pascual, Amaia Salvador, Eva Mohedano, Kevin McGuinness, Jordi Torres, and Xavier Giro-i Nieto, “Wav2pix: speech-conditioned face generation using generative adversarial networks,” in *ICASSP*, 2019, vol. 3.
- [4] Mehdi Mirza and Simon Osindero, “Conditional generative adversarial nets,” *arXiv preprint arXiv:1411.1784*, 2014.
- [5] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros, “Image-to-image translation with conditional adversarial networks,” in *CVPR*, 2017, pp. 1125–1134.
- [6] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *ICCV*, 2017, pp. 2223–2232.
- [7] Tero Karras, Samuli Laine, and Timo Aila, “A style-based generator architecture for generative adversarial networks,” in *CVPR*, 2019, pp. 4401–4410.
- [8] Xun Huang and Serge Belongie, “Arbitrary style transfer in real-time with adaptive instance normalization,” in *ICCV*, 2017, pp. 1501–1510.
- [9] Tero Karras, Timo Aila, Samuli Laine, Antti Herva, and Jaakko Lehtinen, “Audio-driven facial animation by joint end-to-end learning of pose and emotion,” *TOG*, vol. 36, no. 4, pp. 94, 2017.
- [10] Guanzhong Tian, Yi Yuan, and Yong Liu, “Audio2face: Generating speech/face animation from single audio with attention-based bidirectional lstm networks,” in *ICMEW. IEEE*, 2019, pp. 366–371.
- [11] Daniel Cudeiro, Timo Bolkart, Cassidy Laidlaw, Anurag Ranjan, and Michael J Black, “Capture, learning, and synthesis of 3d speaking styles,” in *CVPR*, 2019, pp. 10101–10111.
- [12] Supasorn Suwajanakorn, Steven M Seitz, and Ira Kemelmacher-Shlizerman, “Synthesizing obama: learning lip sync from audio,” *TOG*, vol. 36, no. 4, pp. 95, 2017.
- [13] Najmeh Sadoughi and Carlos Busso, “Speech-driven expressive talking lips with conditional sequential generative adversarial networks,” *IEEE Transactions on Affective Computing*, 2019.
- [14] Olivia Wiles, A Sophia Koepke, and Andrew Zisserman, “X2face: A network for controlling face generation using images, audio, and pose codes,” in *ECCV*, 2018, pp. 670–686.
- [15] Face++, ,” <https://www.faceplusplus.com/attributes/>, 2019, Accessed September 22, 2019.
- [16] Lindasalwa Muda, Mumtaj Begam, and Irraivan Elamvazuthi, “Voice recognition algorithms using mel frequency cepstral coefficient (mfcc) and dynamic time warping (dtw) techniques,” *arXiv preprint arXiv:1003.4083*, 2010.
- [17] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [18] Zhou Wang, Alan C Bovik, Hamid R Sheikh, Eero P Simoncelli, et al., “Image quality assessment: from error visibility to structural similarity,” *IEEE transactions on image processing*, vol. 13, no. 4, pp. 600–612, 2004.
- [19] 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 *Neurips*, 2017.