



Lip Movements Generation at a Glance

Lele Chen^(✉) , Zhiheng Li , Ross K. Maddox , Zhiyao Duan ,
and Chenliang Xu 

University of Rochester, Rochester, USA
lchen63@cs.rochester.edu

Abstract. In this paper, we consider the task: given an arbitrary audio speech and one lip image of arbitrary target identity, generate synthesized lip movements of the target identity saying the speech. To perform well, a model needs to not only consider the retention of target identity, photo-realistic of synthesized images, consistency and smoothness of lip images in a sequence, but more importantly, learn the correlations between audio speech and lip movements. To solve the collective problems, we devise a network to synthesize lip movements and propose a novel correlation loss to synchronize lip changes and speech changes. Our full model utilizes four losses for a comprehensive consideration; it is trained end-to-end and is robust to lip shapes, view angles and different facial characteristics. Thoughtful experiments on three datasets ranging from lab-recorded to lips in-the-wild show that our model significantly outperforms other state-of-the-art methods extended to this task.

Keywords: Lip movements generation · Audio visual correlation

1 Introduction

Cross-modality generation has become an important and emerging topic of computer vision and its broader AI communities, where examples are beyond the most prominent image/video-to-text [10, 19] and can be found in video-to-sound [23], text-to-image [25], and even sound-to-image [4]. This paper considers a task: given an arbitrary audio speech and one lip image of arbitrary target identity, generate synthesized lip movements of the target identity saying the speech. Notice that the speech does not have to be spoken by the target identity, and neither the speech nor the image of target identity is required to be appeared in the training set (see Fig. 1). Solving this task is crucial to many applications, e.g., enhancing speech comprehension while preserving privacy or assistive devices for hearing impaired people.

Lip movements generation has been traditionally solved as a sub-problem in synthesizing a talking face from speech audio of a target identity [3, 12, 13, 29]. For example, Bo et al. [12] restitch the lower half of the face via a bi-directional LSTM to re-dub a target video from a different audio source. Their model selects

L. Chen and Z. Li—Equal contribution.




Fig. 1. The model takes an audio speech of the women and one lip image of the target identity, a male celebrity in this case, and synthesizes a video of the man’s lip saying the same speech. The synthesized lip movements need to correspond to the speech audio and also maintain the target identity, video smoothness and sharpness

a target mouth region from a dictionary of saved target frames. More recently, Suwajanakorn et al. [29] generate synthesized taking face of President Obama with accurate lip synchronization, given his speech audio. They first use an LSTM model trained on many hours of his weekly address footage to generate mouth landmarks, then retrieve mapped texture and apply complicated post-processing to sharpen the generated video. However, one common problem for these many methods is that they retrieve rather than generating images and thus, require a sizable amount of video frames of the target identity to choose from, whereas our method generates lip movements from a single image of the target identity, i.e., at a glance.

The only work we are aware of that addresses the same task as ours is Chung et al. [6]. They propose an image generator network with skip-connections, and optimize the reconstruction loss between synthesized images and real images. Each time, their model generates one image from 0.35-second audio. Although their video generated image-by-image and enhanced by post-processing looks fine, they have essentially bypassed the harder questions concerning the consistency and smoothness of images in a sequence, as well as the temporal correlations of audio speech and lip movements in a video.

To overcome the above limitations, we propose a novel method that takes speech audio and a lip image of the target identity as input, and generates multiple lip images (16 frames) in a video depicting the corresponding lip movements (see Fig. 1). Observing that speech is highly correlated with lip movements even across identities, a concept grounds lip reading [1, 7], the core of our paper is to explore the best modeling of such correlations in building and training a lip movement generator network. To achieve this goal, we devise a method to fuse time-series audio embedding and identity image embedding in generating multiple lip images, and propose a novel audio-visual correlation loss to synchronize lip changes and speech changes in a video. Our final model utilizes a combination of four losses including the proposed audio-visual correlation loss, a novel three-stream adversarial learning loss to guide a discriminator to judge both image quality and motion quality, a feature-space loss to minimize perceptual-level differences, and a reconstruction loss to minimize pixel-level differences, for

a comprehensive consideration of lip movements generation. The whole system is trained in an end-to-end fashion and is robust to lip shapes, view angles, and different facial characteristics (e.g., beard v.s. no beard). Our code is available at https://github.com/lelechen63/3d_gan. Check out more results at https://youtu.be/7IX_sIL5v0c.

We evaluate our model along with its variants on three datasets: The GRID audiovisual sentence corpus (GRID) [8], Linguistic Data Consortium (LDC) [26] and Lip Reading in the Wild (LRW) [7]. To measure the quantitative accuracy of lip movements, we propose a novel metric that evaluates the detected landmark distance of synthesized lips to ground-truth lips. In addition, we use a cohort of three metrics, Peak Signal to Noise Ratio (PSNR), Structure Similarity Index Measure (SSIM) [32], and perceptual-based no-reference objective image sharpness metric (CPBD) [21], to measure the quality of synthesized lip images, e.g., image sharpness. We compare our model with Chung et al. [6] and an extended version of the state-of-the-art video Generative Adversarial Network (GAN) model [30] to our task. Experimental results show that our model outperforms them significantly on all three datasets (see “Full model” in Table 3). Furthermore, we also show real-world novel examples of synthesized lip movements of celebrities, who are not in our dataset.

Our paper marks three contributions. First, to the best of our knowledge, we are the first to consider the correlations among speech and lip movements in generating multiple lip images at a glance. Second, we explore various models and loss functions in building and training a lip movement generator network. Third, we quantify the evaluation metrics and our final model achieves significant improvement over state-of-the-art methods extended to this task on three datasets ranging from lab-recorded to lips in-the-wild.

2 Related Work

We have briefly surveyed work in lip movement generation in the Introduction section. Here, we discuss related work of each techniques used in our model.

A related but different task to ours is lip reading, where it also tackles the cross-modality generation problem. [1, 7] use the correlation between lip movement and the sentences/words to interpret the audio information from the visual information. Rasiwasia et al. [24] use Canonical Correlation Analysis (CCA) [16] subspace learning to learn two intermediate feature spaces for two modalities where they do correlation on the projected features. Cutler and Davis [9] use Time Delay Neural Network [31] (TDNN) to extract temporal invariant audio features and visual features. These works have inspired us to model correlations between speech audio and lip movements in generating videos.

Audio variations and lip movements are not always synchronized in the production of human speech; lips often move before the audio signal is produced [2]. Such delay between audio and visual needs to be considered when designing a model. Suwajanakorn et al. [29] apply a time-delayed RNN without outputting value in the first few RNN cells. Therefore, the output is shifted accordingly to

the delayed steps. However, such delay is empirically fixed by hand and thus, it is hard to determine the amount of delay for videos in-the-wild. We follow [31] to extract features with a large receptive field along temporal dimension, but use a convolutional network instead of TDNN that leads to a simpler design.

Adversarial training [14] is recently introduced as a novel and effective way to train generative models. Researchers find that by conditioning the model on additional information, it is possible to direct the data generation process [5, 20, 22]. Furthermore, GAN has shown its ability to bridge the gap between different modalities and produce useful joint representations. We also use GAN loss in our training but we show that combining it with other losses leads to better results.

3 Lip-Movement Generator Network

The overall data flow of our lip-movement generator network is depicted in Fig. 2. **In this paper, we omit channel dimension of all tensors for simple illustration.** Recall that the input to our network are a speech audio and one single image of the target identity, and the output of our network are synthesized lip images of the target identity saying that audio. The synthesized lip movements need to correspond to the speech audio, maintain the target identity, ensure the video smoothness, and be photo-realistic.




Fig. 2. Full model illustration. Audio encoder and identity encoder extracts and fuses audio and visual embeddings. Audio-Identity fusion network fuses features from two modalities. Decoder expands fused feature to synthesized video. Correlation Networks are in charge of strengthening the audio-visual mapping. Three-Stream discriminator is responsible for distinguishing generated video and real video

3.1 Audio-Identity Fusion and Generation

First, we encode the two-stream input information. For audio stream, the raw audio waveform, denoted as S_{raw} , is first transformed into log-mel spectrogram (see detail in Sect. 5.1), denoted as S_{lms} , then encoded by an audio encoder network into audio features $f_s \in \mathbb{R}^{T \times F}$, where T and F denote the number of time frames and frequency channels. For visual stream, an input identity image, denoted as p_r , is encoded by an identity encoder network. The network outputs

image features $f_p \in \mathbb{R}^{H \times W}$, where H and W denote the height, width of the output image features.

We fuse audio features f_s and visual features f_p together, whose output, the synthesized video feature f_v , will be expanded by several residual blocks and 3D deconvolution operations to generate synthesized video \hat{v} . In order to make sure the synthesized clip is based on the target person and also captures the time-variation of speech, we investigate an effective way to fuse f_s and f_p to get f_v for generating a video. Here, the challenge is that the feature maps exist in different modalities, e.g., audio, visual, and audio-visual, and reside in different feature spaces, e.g., time-frequency, space, and space-time.

Our fusion method is based on duplication and concatenation. This process is depicted in Fig. 3. For each audio feature, we duplicate that feature along frequency dimension in each time step, i.e., from the size of $T \times F$ to the size of $T \times F \times F$. Image feature, which can be viewed as a template for video representation, is copied T times, i.e., from $H \times W$ to a new size $T \times H \times W$. We set $H = W = F$ in this method. Then, two kinds of duplicated features are concatenated along channel dimension.




Fig. 3. Audio-Identity fusion. Transfer audio time-frequency features and image spatial features to video spatial-temporal features

3.2 Audio-Visual Derivative Correlation Loss

We believe that the acoustic information of audio speech is correlated with lip movements even across identities because of their shared high-level representation. Besides, we also regard that variation along temporal axis between two modalities are more likely to be correlated. In other words, compared with acoustic feature and visual feature of lip shape themselves, the variation of audio feature (e.g. the voice raising to a higher pitch) and variation of visual feature (e.g. mouth opening) have a higher likelihood to be correlated. Therefore, we propose a method to optimize the correlations of the two modalities in their feature spaces. We use f'_s in size of $(T - 1) \times F$, the derivative of audio feature f_s (with size of $T \times F$) between consecutive frames in temporal dimension, to represent the changes in speech. It goes through an audio derivative encoder network ϕ_s , and thus, we have audio derivative feature $\phi_s(f'_s)$. Similarly, we use $\mathcal{F}(v)$ to represent optical flows of each consecutive frames in a video v , where \mathcal{F} is an optical flow estimation algorithm. It goes through an optical flow encoder network ϕ_v , and thus, we have $\phi_v(\mathcal{F}(v))$ to depict the visual variations of lip

movements in the feature space. We use cosine similarity loss to maximize the correlation between audio derivative feature and visual derivative feature:

$$\ell_{corr} = 1 - \frac{\phi_s(f'_s) \cdot \phi_v(\mathcal{F}(v))}{\|\phi_s(f'_s)\|_2 \cdot \|\phi_v(\mathcal{F}(v))\|_2}. \quad (1)$$

Here, the optical flow algorithm applied to the synthesized frames needs to be differentiable for back-propagation [27]. In our implementation, we add a small number ($\epsilon = 10^{-8}$) to the denominator to avoid division by zero. In order to avoid trivial solution when ϕ_s and ϕ_v are learned to predict constant outputs $\phi_s(f'_x)$ and $\phi_v(\mathcal{F}(v))$ which are perfectly correlated and the ℓ_{corr} will go to 0, we combine other losses during the training process (see Eq. 2).

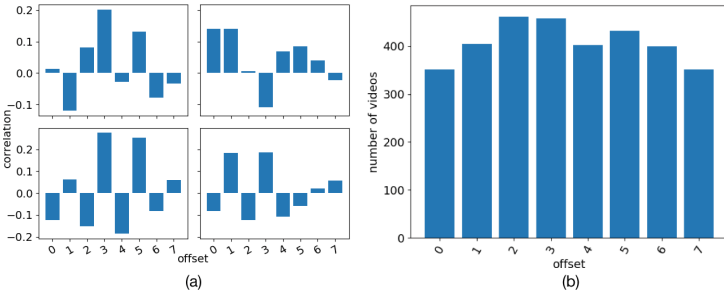


Fig. 4. (a): Correlation coefficients with different offsets of four example videos. (b): Number of videos of different offsets with which the video has the maximum correlation coefficient. X-axes of both (a) and (b) stands for time steps of flow field shifted forward

Correlation Networks. The audio and visual information are not perfectly aligned in time. Usually, lip shape forms earlier than sound. For instance, when we say word ‘bed’, upper and lower lips meet before speaking the word [2]. If such delay problem exists, aforementioned correlation loss, assuming audio-visual information are perfectly aligned, may not work. We verify the delayed correspondence problem between audio and visual information by designing a case study on 3260 videos randomly sampled from the GRID dataset. The solution for the delayed correspondence problem is given in the next paragraph. In the case study, for each 75-frame video v , we calculate the mean values of each 74 derivatives of audio s_{lms} and mean values of each 74 optical flow fields $\phi_v(\mathcal{F}_v)$. With respect to each video, we shift mean values of optical flows forward along time at different offsets (0 to 7 in our case study) and calculate Pearson correlation coefficients of those two parts. Results of four videos, calculated by aforementioned procedures, are shown in Fig. 4(a). Finally, we count the number of videos in different offsets at which the video has the largest correlation coefficient, as shown in Fig. 4(b). Figure 4 shows that different videos prefer different offsets to output the maximum correlation coefficient, which indicates that fixing a constant offset of all audio-visual inputs would not solve the problem of correlation with inconsistent delays among all videos in a dataset.

To mitigate such delayed correlation problem, we design correlation networks (as shown in Fig. 2) containing an audio derivative encoder ϕ_s and an optical flow encoder ϕ_v to extract features used for calculating the correlation loss in Eq. 1. These networks reduce the feature size but retain the temporal length simultaneously. The sizes of the two outputs are matched for calculating the correlation loss. We use 3D CNNs to implement these networks, which are also helpful to mitigate the fixed offset problem happens in previous works [29]. Both ϕ_s and ϕ_v output features with large receptive fields (9 for $\phi_s(f'_s)$ and 13 for $\phi_v(\mathcal{F}_v)$), which consider the audio-visual correlation in large temporal dimension. Compared with time-delayed RNN proposed in [29], CNN can learn delay from the dataset rather than set it as a hyper-parameter. Besides, CNN architecture benefits from its weight sharing property leading to a simpler and smaller design than TDNN [31].

4 Full Model and Training

Without loss of generality, we use pairs of lip movement video and speech audio $\{(v^j, s^j)\}$ in the training stage, where v^j represents the j th video in our dataset and s^j represents the corresponding speech audio. We omit the superscript j when it is not necessary for the discussion of one sample. We use p_r to denote one lip image of the target speaker, which can provide the initial texture information. During training, we train over (v, s) in the training set and sample p_r to be one frame randomly selected from the raw video where v^j is sampled from to ensure that v and p_r contain the same identity. Therefore, the system is robust to the lip shape of the identity p_r . The objective of training is to generate a realistic video \hat{v} that resembles v . For testing, the speech s and identity image p_r can be any speech and any lip image (even out of the dataset we used in training). Next, we present the full model in the context of training.

Our full model (see Fig. 2) is end-to-end trainable and is optimized according to the following objective function:

$$\mathcal{L} = \ell_{corr} + \lambda_1 \ell_{pix} + \lambda_2 \ell_{perc} + \lambda_3 \ell_{gen} , \quad (2)$$

where λ_1 , λ_2 and λ_3 are coefficients of different loss terms. We set them as 0.5, 1.0, 1.0 respectively in this paper. The intuitions behind the four losses are as follows:

- ℓ_{corr} : Correlation loss, illustrated in Sect. 3.2, is introduced to ensure the correlation between audio and visual information.
- ℓ_{pix} : Pixel-level reconstruction loss, defined as $\ell_{pix}(\hat{v}, v) = \|v - \hat{v}\|$, which aims to make the model sensitive to speaker’s appearance, i.e., retain the identity texture. However, we find that using it alone will reduce the sharpness of the synthesized video frames.
- ℓ_{perc} : Perceptual loss, which is originally proposed by [17] as a method used in image style transfer and super-resolution. It utilizes high-level features to compare generated images and ground-truth images, resulting in better

sharpness of the synthesized image. We adapt this perceptual loss and detail it in Sect. 4.1.

- ℓ_{gen} : Adversarial loss allows our model to generate overall realistic looking images and is defined as: $\ell_{gen} = -\log D([s^j, \hat{v}^j])$, where D is a discriminator network. We describe the detail of our proposed stream-stream GAN discriminator in Sect. 4.2.

4.1 Autoencoder and Perceptual Loss

In order to avoid over-smoothed phenomenon of synthesized video frames \hat{v} , we adapt perceptual loss proposed by Johnson et al. [17], which reflects perceptual-level similarity of images. The perceptual loss is defined as:

$$\ell_{perc}(\hat{v}, v) = \|\varphi(v) - \varphi(\hat{v})\|_2^2, \quad (3)$$

where φ is a feature extraction network. We train an autoencoder to reconstruct video clips. To let the network be more sensitive to structure features, we apply six residual blocks after the convolution layers. We train the autoencoder from scratch, then fix the weights and use its encoder part as φ to calculate perceptual loss for training the full model.

4.2 Three-Stream GAN Discriminator

The GAN discriminator in [30] for synthesizing video considers the motion changes implicitly by 3D convolution. In order to generate sharp and smooth changing video frames, we propose a three-stream discriminator network (see Fig. 5) to distinguish the synthesized video (\hat{v}^j) from real video (v^j) that not only considers motion explicitly and but also conditions on the input speech signal. The input to the discriminator is a video clip with the corresponding audio. We have the following three streams. For audio stream (also used in our generator), we first convert the raw audio to log-mel spectrogram, then use four convolutional layers followed by a fully-connected layer to get a 1D vector. We duplicate it to match features from other streams. For video stream, we use four 3D CNN layers to extract video features. In addition, we include an optical flow stream that attends to motion changes explicitly. We fine-tune the FlowNet [11], which is pre-trained on FlyingChairs dataset, to extract optical flows, then apply four 3D CNN layers to extract features.

Finally, we concatenate the three-stream features in channel dimension and let them go through two convolutional layers to output the discriminator probability. We adapt mismatch strategy [25] to make sure that our discriminator

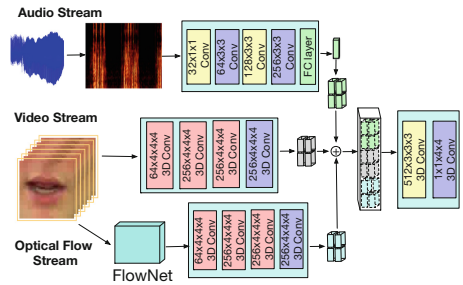


Fig. 5. Three-stream GAN discriminator illustration

Table 1. Dataset information. Validation set: known speakers but unseen sentences. Testing set: unseen speakers and unseen sentences

Dataset	GRID	LRW	LDC
Train	211k (37.5 h)	841k (159.8 h)	36k (6.4 h)
Val.	23 (4.2 h)	N/A	4k (0.7 h)
Test	7k (1.3 h)	40k (7.8 h)	6.6k (1.2 h)

is also sensitive to mismatched audio and visual information. Therefore, the discriminator loss is defined as:

$$\begin{aligned}
 \ell_{dis} = & -\log D([s^j, v^j]) \\
 & -\lambda_p \log(1 - D([s^j, \hat{v}])) \\
 & -\lambda_u \log(1 - D([s^j, v^k])), \quad k \neq j,
 \end{aligned} \tag{4}$$

where v^k represents a mismatch real video. We set both λ_p and λ_u 0.5 in our experiment. The performance of the optical flow stream is discussed in Sect. 5.3.

5 Experiments

In this section, we first introduce datasets and experimental settings, and our adapted evaluation metrics. Then, we show ablation study and comparison to the state of the art. Finally, we demonstrate real-world novel examples.

5.1 Datasets and Settings

We present our experiments on GRID [8], LRW [7] and LDC [26] datasets (see Table 1). There are 33 different speakers in GRID. Each speaker has 1000 short videos. The LRW dataset consists of 500 different words spoken by hundreds of different speakers. There are 14 speakers in the LDC dataset in which each speaker reads 238 different words and 166 different sentences. Videos in GRID and LDC are lab-recorded while videos in LRW are collected from news. Our data is composed by two parts: audio and image frames. The network can output different numbers of frames. In this work, we only consider generating 16 frames of mouth region. As the videos are sampled at 25fps, the time span of the synthesized image frames is 0.64s. We use sliding window approach (window size: 16 frames, overlap: 8 frames) to obtain training and testing video samples over raw videos.

Audio: We extract audio from the video file with a sampling rate of 41.1 kHz. Each input audio is 0.64s long (0.04×16). To encode audio, we first transform the raw audio waveform into the time-frequency domain by calculating the Log-amplitude Mel-frequency Spectrum (LMS). When we calculate the LMS, the number of samples between successive frames, the length of the FFT window,

and the number of Mel bands are 512, 1024 and 128, respectively. This operation will convert a 0.64-s raw audio to a 64×128 time-frequency representation.

Images: First, we extract all image frames from videos. Then, we extract lip landmarks [18] and crop the image around the lip. Landmarks are only used for cropping and evaluation. We resize all of the cropped images to 64×64 . So, each 0.64-s audio corresponds to a $16 \times 3 \times 64 \times 64$ RGB image sequence.

We adopt Adam optimizer during training and fixed learning rates of 2×10^{-4} with weight decay of 4×10^{-4} , and coefficients β_1 and β_2 are 0.5 and 0.999, respectively. We initialize all network layers according to the method described in [15]. All models are trained and tested on a single NVIDIA GTX 1080Ti. During testing, generating one single frame costs 0.015 s.

5.2 Evaluation Metrics

To evaluate the quality of the synthesized video frames, we compute Peak Signal to Noise Ratio (PSNR) and Structure Similarity Index Measure (SSIM) [32]. To evaluate the sharpness of the generated image frames, we compute the perceptual-based no-reference objective image sharpness metric (CPBD) [21].

As far as we know, no quantitative metric has been used to evaluate the accuracy of generated lip movements video. Therefore, to evaluate whether the synthesized video \hat{v} corresponds to accurate lip movements based on the input audio, a new metric is proposed by calculating the Landmark Distance (LMD). We use Dlib [18], a HOG-based facial landmarks detector, which is also widely used in lip-movement generation task and other related works [28, 29], to detect lip landmarks on \hat{v} and v , and mark them as LF and LR , respectively. To eliminate the geometric difference, we calibrate the two mean points of lip landmarks in LF and LR . Then, we calculate the Euclidean distance between each corresponding pairs of landmarks on LF and LR , and finally normalized them with temporal length and number of landmark points. LMD is defined as:

$$LMD = \frac{1}{T} \times \frac{1}{P} \sum_{t=1}^T \sum_{p=1}^P \|LR_{t,p} - LF_{t,p}\|_2, \quad (5)$$

where T denotes the temporal length of video and P denotes the total number of landmark points on each image (20 points).

5.3 Ablation Study

We conduct ablation experiments to study the contributions of the three components in our full model separately: correlation loss, three-stream GAN discriminator and perceptual loss. The ablation study is conducted on GRID dataset. Results are shown in Table 2. Different implementations are discussed in below as well. The following ablation studies are trained and tested on the GRID dataset.

Perceptual Loss and Reconstruction Loss. Generally, we find that perceptual loss can help our model generate more accurate lip movements with higher

Table 2. Ablation results on GRID dataset. The full model (method (e)) uses all four losses as described in Sect. 4. For LMD, the lower the better. SSIM, PSNR and CPBD, the higher the better. We bold top-2 leading scores for each metric

Methods	(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)	(i)
l_{pix}	✓		✓	✓	✓	✓	✓	✓	✓
l_{perc}	✓	✓		✓	✓	✓	✓	✓	✓
l_{corr}	✓	✓	✓		✓		✓		✓
l_{corr} (Non-Derivative Corr.)						✓			
l_{gen} (Three-Stream)		✓	✓	✓	✓	✓			
l_{gen} (Two-Stream)							✓	✓	
l_{gen} (Three-Str. Frame-Diff.)									✓
Metrics									
LMD	1.24	1.31	1.38	1.31	1.18	1.96	1.39	1.42	1.40
SSIM	0.77	0.71	0.72	0.70	0.73	0.52	0.68	0.59	0.66
PSNR	29.36	29.79	29.66	29.80	29.98	28.6	29.59	29.46	29.51
CPBD	0.01	0.18	0.17	0.21	0.18	0.22	0.19	0.18	0.21

image quality, and improve image sharpness at the same time (see method (c) v.s. method (e) in Table 2). If we compare method (b) with method (e), we can find that pixel-wise reconstruction loss can improve the LMD, SSIM and PSNR while decreasing the CPBD.

Correlation Models. When correlation loss is removed from final objective function Eq. 2, results are worse than final objective in LMD, SSIM and PSNR, demonstrating the importance of correlation loss in generating more accurate lip movement (see method (d) v.s. method (e) or method (g) v.s. method (h)).

Besides, we investigate a model variant, *Non-Derivative Correlation* (see method (f) in Table 2), for analyzing the necessity of applying derivative features to ϕ_s and ϕ_v . Instead of using the derivative of audio features and the optical flow, this variant just uses audio features f_s and video frames v directly as inputs. Neither the derivative nor the optical flow is calculated here. Other settings (e.g., network structure and loss functions) are identical with the full model (denoted as method (e) in Table 2). The comparison between method (e) and method (f) in Table 2 shows that derivative correlation model outperforms the *Non-Derivative Correlation* model in metrics such as SSIM, PSNR and LMD. With respect to *Non-Derivative Correlation* model, landmark distance is even worse than model without correlation loss (method (d)). The experimental result proves our assumption that it is the derivatives of audio and visual information rather than the direct features that are correlated. Furthermore, since *Non-Derivative Correlation* model fails to learn the derivative feature implicitly (i.e. convolutional layers fails to transform feature to their derivatives), using the derivatives of audio and visual features to do correlation as a strong expert prior knowledge is necessary.

GAN Discriminator. We find that ℓ_{gen} improves the CPBD result a lot (see method (a) and method (e) in Table 2), demonstrating that discriminator can improve the frame sharpness. Furthermore, we use two model variants to study the effectiveness of proposed three-stream GAN discriminator. ℓ_{gen} (Two-Stream) only contains audio stream and video stream. ℓ_{gen} (Three-Str. Frame-Diff.) replaces the optical flow with frame-wise difference, i.e., $L1$ distance between adjacent frames, as the third stream to capture motion changes. First, compared with the two-stream variant, our full model with proposed three-stream discriminator gives better result (see method (e) v.s. method (g)), which indicates the effectiveness of explicitly modeling motion changes among the frames. Second, compared with the three-stream frame-difference variant, the full model generates more realistic (higher CPBD) and accurate lip movements (lower LMD) (see method (e) and (i)), which indicates that optical flow is a better representation than frame-wise difference for modeling motion changes.

5.4 Comparison to State-of-the-Art

In this section, we compare our full model with two state-of-the-art methods [6, 30]. We extend [30] to a conditional GAN structure, which receives the same target image information and audio information as our models. There are a few changes made for ensuring a fair comparison with [6]: we did not pretrain the identity encoder; we changed two convolution layers to fit the image size (64×64); we omitted the De-blurring post-processing as we aim to compare directly the generative models themselves.

The quantitative results are shown in Table 3. We test our models on three different datasets. The results show that our proposed models outperform state-of-the-art models in most of the metrics. In terms of LMD and PSNR, our full model shows better performance than methods that use discriminator [30] or reconstruction loss [6]. Model proposed by Chung et al., based on reconstruction loss, generates blurred images, which makes them look unrealistic. We can see this phenomenon in the CPBD column. The LRW dataset consists of people talking in the wild so resolution is much smaller in terms of lip region. We need to scale up the ground truth to 64×64 , which leads to a lower resolution and CPBD. We suspect this is the reason why we achieve a better CPBD than ground truth in LRW dataset.




Fig. 6. Generated videos of our model on three testing datasets compared with state-of-the-art methods. In the testing set, none of the speakers were in the training set

Table 3. Results on three datasets. Models mentioned in this table are trained from scratch (no pre-training included) and be tested on each dataset a time. We bold each leading score

Method	GRID				LDC				LRW			
	LMD	SSIM	PSNR	CPBD	LMD	SSIM	PSNR	CPBD	LMD	SSIM	PSNR	CPBD
G. T.	0	N/A	N/A	0.141	0	N/A	N/A	0.211	0	N/A	N/A	0.068
Vondrick [30]	2.38	0.60	28.45	0.129	2.34	0.75	27.96	0.160	3.28	0.34	28.03	0.082
Chung [6]	1.35	0.74	29.36	0.016	2.13	0.50	28.22	0.010	2.25	0.46	28.06	0.083
Full model	1.18	0.73	29.89	0.175	1.82	0.57	28.87	0.172	1.92	0.53	28.65	0.075

The qualitative results compared with other methods are shown in Fig. 6. Our model generates sharper video frames on all three datasets, which has also been supported by the CPBD results, even if input identity images are in low resolution. We show additional results of our method in Fig. 7. Our model can generate realistic lip movement videos that are robust to view angles, lip shapes and facial characteristics in most of the times. However, sometimes our model fails to preserve the skin color (see the last two examples in Fig. 7), which, we suspect, is due to the imbalanced data distribution in LRW dataset. Furthermore, the model has difficulties in capturing the amount of lip deformations of each person, which is an intrinsic problem when learning from a single image.




Fig. 7. Randomly selected outputs of the full model on the LRW testing set. The lip shape in videos not only synchronize well with the ground truth, but maintain identity information, such as (beard v.s. no beard)

5.5 Real-World Novel Examples

For generating videos given unpaired identity image and audio in the real-world, i.e., source identity of provided audio is different from the target identity and out of the datasets, our model can still perform well. Results are shown in Fig. 8, in which three identity images of celebrities are selected outside of the datasets the model trained on and the input audio is selected in GRID dataset. For our model trained on LRW, both identity images and audio are unseen. For our model trained on GRID, we leave the source identity out of the training.




Fig. 8. The figure shows the generated images based on three identity images outside of dataset, which is also not paired with the input audio from GRID dataset. Two full models trained on GRID and LRW datasets are used here for a comparison

The videos generated by our model show promising qualitative performance. Both lip regions of Musk and Sandburg are rotated by some degrees. We can see that the rotation phenomenon in the generated video frames as well. Besides, our model can also retain beards in our generated clip when identity (target person) has beards as well. However, we observe that model trained on GRID dataset fails to reserve the identity information. Because of the fact that LRW dataset has much more identities than GRID dataset (hundreds v.s. 33), the model trained on LRW has better generalization ability.

6 Conclusion and Future Work

In this paper, we study the task: given an arbitrary audio speech and one lip image of arbitrary target identity, generate synthesized lip movements of the target identity saying the speech. To perform well in this task, it requires a model to not only consider the retention of the target identity, photo-realistic of synthesized images, consistency and smoothness of images in a video, but also learn the correlations between the speech audio and lip movements. We achieve this by proposing a new generator network, a novel audio-visual correlation loss and a full model that considers four complementary losses. We show significant

improvements on three datasets compared to two state-of-the-art methods. There are several future directions. First, non-fixed length lip movements generation is needed for a more practical purpose. Second, it is valuable to extend our method to one generating full face in an end-to-end paradigm.

Acknowledgement. This work was supported in part by NSF BIGDATA 1741472, NIH grant R00 DC014288 and the University of Rochester AR/VR Pilot Award. We gratefully acknowledge the gift donations of Markable, Inc., Tencent and the support of NVIDIA with the donation of GPUs used for this research. This article solely reflects the opinions and conclusions of its authors and not the funding agents.

References

1. Assael, Y.M., Shillingford, B., Whiteson, S., de Freitas, N.: LipNet: end-to-end sentence-level lipreading. arXiv preprint [arXiv:1611.01599](https://arxiv.org/abs/1611.01599) (2017)
2. Chandrasekaran, C., Trubanova, A., Stillitano, S., Caplier, A., Ghazanfar, A.A.: The natural statistics of audiovisual speech. *PLOS Comput. Biol.* **5**(7) (2009)
3. Charles, J., Magee, D., Hogg, D.: Virtual immortality: reanimating characters from TV shows. In: Hua, G., Jégou, H. (eds.) *ECCV 2016*. LNCS, vol. 9915, pp. 879–886. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-49409-8_71
4. Chen, L., Srivastava, S., Duan, Z., Xu, C.: Deep cross-modal audio-visual generation. In: *Proceedings of Multimedia Thematic Workshops*. ACM (2017)
5. Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., Abbeel, P.: InfoGAN: interpretable representation learning by information maximizing generative adversarial nets. In: *Proceedings of NIPS*. Curran Associates, Inc. (2016)
6. Chung, J.S., Jamaludin, A., Zisserman, A.: You said that? In: *Proceedings of BMVC*. Springer (2017)
7. Chung, J.S., Zisserman, A.: Lip reading in the wild. In: Lai, S.-H., Lepetit, V., Nishino, K., Sato, Y. (eds.) *ACCV 2016*. LNCS, vol. 10112, pp. 87–103. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-54184-6_6
8. Cooke, M., Barker, J., Cunningham, S., Shao, X.: An audio-visual corpus for speech perception and automatic speech recognition. *J. Acoust. Soc. Am.* **120**(5), 2421–2424 (2006)
9. Cutler, R., Davis, L.S.: Look who’s talking: speaker detection using video and audio correlation. In: *Proceedings of ICME*. IEEE (2000)
10. Das, P., Xu, C., Doell, R., Corso, J.J.: A thousand frames in just a few words: lingual description of videos through latent topics and sparse object stitching. In: *Proceedings of CVPR*. IEEE (2013)
11. Dosovitskiy, A., et al.: FlowNet: learning optical flow with convolutional networks. In: *Proceedings of ICCV*. IEEE (2015)
12. Fan, B., Wang, L., Soong, F.K., Xie, L.: Photo-real talking head with deep bidirectional LSTM. In: *ICASSP*. IEEE (2015)
13. Garrido, P., et al.: VDub: modifying face video of actors for plausible visual alignment to a dubbed audio track. *Comput. Graph. Forum* **34**(2), 193–204 (2015)
14. Goodfellow, I.J., et al.: Generative adversarial nets. In: *Proceedings of NIPS*. Curran Associates, Inc. (2014)
15. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: surpassing human-level performance on ImageNet classification. In: *Proceedings of CVPR*. IEEE (2015)

16. Hotelling, H.: Relations between two sets of variates. In: Kotz, S., Johnson, N.L. (eds.) *Breakthroughs in Statistics*, pp. 162–190. Springer, New York (1992). https://doi.org/10.1007/978-1-4612-4380-9_14
17. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016*. LNCS, vol. 9906, pp. 694–711. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46475-6_43
18. King, D.E.: Dlib-ml: a machine learning toolkit. *JMLR* **10**, 1755–1758 (2009)
19. Kulkarni, G., et al.: Baby talk: understanding and generating simple image descriptions. In: *Proceedings of CVPR*. IEEE (2011)
20. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint [arXiv:1411.1784](https://arxiv.org/abs/1411.1784) (2014)
21. Narvekar, N.D., Karam, L.J.: A no-reference image blur metric based on the cumulative probability of blur detection (CPBD). *IEEE TIP* **20**(9), 2678–2683 (2011)
22. Odena, A., Olah, C., Shlens, J.: Conditional image synthesis with auxiliary classifier GANs. In: *Proceedings of ICML*. PMLR (2017)
23. Owens, A., Isola, P., McDermott, J., Torralba, A., Adelson, E.H., Freeman, W.T.: Visually indicated sounds. In: *Proceedings of CVPR*. IEEE (2016)
24. Rasiwasia, N., et al.: A new approach to cross-modal multimedia retrieval. In: *Proceedings of Multimedia*. ACM (2010)
25. Reed, S.E., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., Lee, H.: Generative adversarial text to image synthesis. In: *ICML*. PMLR (2016)
26. Richie, S., Warburton, C., Carter, M.: Audiovisual database of spoken American English. Linguistic Data Consortium (2009)
27. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning representations by back-propagating errors. *Nature* **323**(6088), 533–536 (1986)
28. Son Chung, J., Senior, A., Vinyals, O., Zisserman, A.: Lip reading sentences in the wild. In: *Proceedings of CVPR*. IEEE (2017)
29. Suwajanakorn, S., Seitz, S.M., Kemelmacher-Shlizerman, I.: Synthesizing Obama: learning lip sync from audio. *ACM Trans. Graph.* **36**(4), 95:1–95:13 (2017)
30. Vondrick, C., Pirsiavash, H., Torralba, A.: Generating videos with scene dynamics. In: *Proceedings of NIPS*. Curran Associates, Inc. (2016)
31. Waibel, A.H., Hanazawa, T., Hinton, G.E., Shikano, K., Lang, K.J.: Phoneme recognition using time-delay neural networks. *IEEE Trans. Acoust. Speech Signal Process.* **37**(3), 328–339 (1989)
32. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE TIP* **13**(4), 600–612 (2004)