

# Extreme-scale Talking-Face Video Upsampling with Audio-Visual Priors

Sindhu B Hegde\*

sindhu.hegde@research.iiit.ac.in

International Institute of Information Technology  
Hyderabad, India

Vinay P Namboodiri

vpn22@bath.ac.uk  
University of Bath  
United Kingdom



Figure 1: We solve the problem of upsampling extremely low-resolution (LR) talking-face videos to generate high-resolution (HR) outputs. Our approach exploits LR frames ( $8 \times 8$  pixels), corresponding audio signal and a single HR target identity image to synthesize realistic, high-quality talking-face videos ( $256 \times 256$  pixels). Please check our project page for video results.

## ABSTRACT

In this paper, we explore an interesting question of what can be obtained from an  $8 \times 8$  pixel video sequence. Surprisingly, it turns out to be quite a lot. We show that when we process this  $8 \times 8$  video with the right set of audio and image priors, we can obtain a full-length,  $256 \times 256$  video. We achieve this  $32\times$  scaling of an extremely low-resolution input using our novel audio-visual upsampling network. The audio prior helps to recover the elemental facial details and precise lip shapes and a single high-resolution target identity image prior provides us with rich appearance details. Our approach is an end-to-end multi-stage framework. The first stage produces a coarse intermediate output video that can be then used to animate single target identity image and generate realistic, accurate and high-quality outputs. Our approach is simple and performs exceedingly well (an  $8\times$  improvement in FID score) compared to previous super-resolution methods. We also extend our model to talking-face video compression, and show that we obtain a  $3.5\times$  improvement in terms of bits/pixel over the previous state-of-the-art. The results from our network are thoroughly analyzed through extensive ablation experiments (in the paper and supplementary material). We also provide the demo video along with code and models on our website<sup>1</sup>.

\*Both authors contributed equally to this research.

<sup>1</sup><http://cvit.iiit.ac.in/research/projects/cvit-projects/talking-face-video-upsampling>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '22, October 10–14, 2022, Lisboa, Portugal

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9203-7/22/10...\$15.00

<https://doi.org/10.1145/3503161.3548080>

Rudrabha Mukhopadhyay\*

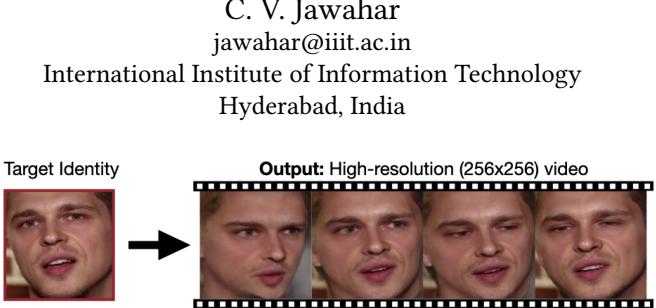
radrabha.m@research.iiit.ac.in

International Institute of Information Technology  
Hyderabad, India

C. V. Jawahar

jawahar@iiit.ac.in

International Institute of Information Technology  
Hyderabad, India



## CCS CONCEPTS

- Computing methodologies → Reconstruction; Neural networks.

## KEYWORDS

video upsampling, video super-resolution, video compression, talking-face videos, audio-visual learning

### ACM Reference Format:

Sindhu B Hegde, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and C. V. Jawahar. 2022. Extreme-scale Talking-Face Video Upsampling with Audio-Visual Priors. In *Proceedings of the 30th ACM International Conference on Multimedia (MM '22), October 10–14, 2022, Lisboa, Portugal*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3503161.3548080>

## 1 INTRODUCTION

Over the years, we have always been fascinated with questions that can push the limits of computer vision. For instance, can we recognize actions [15] or objects [48] in videos/images of about  $30 \times 30$  pixels? Or how small a face can we detect in an image? It turns out, the face can be as tiny as 3 pixels tall [25]! Recovering extremely feeble signals has also led to remarkable achievements, such as the imaging of the black hole. In this vein, we explore whether we can upsample talking-face videos with resolutions as low as  $8 \times 8$  pixels. Clearly, this is an exceptionally challenging task. Surprisingly, we reveal that one can obtain realistic high-resolution (HR) talking-faces ( $256 \times 256$  pixels) when provided with the right set of additional information. We utilize a single target identity image and the accompanying audio to upsample  $8 \times 8$  video to a full  $256 \times 256$  dimensional video (a fascinating  $32\times$  scale-factor) that far exceeds previous methods. In today's digitally connected world, where talking-face videos are among the most common forms of video content, our multi-modal system can have numerous potential

applications. Some of them include: (i) video conferencing in low-bandwidth situations, (ii) recovering low-quality archival footage of public talks and speeches and (iii) enhancing videos captured from a distance with a high camera zoom.

**Challenges:** Although our task has promising applications, synthesizing HR videos from extremely low-resolution (LR) inputs (e.g.,  $8 \times 8$  pixels) is a very challenging task. Essential face attributes such as identity, age and gender are almost entirely lost at such low resolutions and cannot be directly recovered (see LR video in Figure 1). Apart from reconstructing these elementary identity details, the network must also learn to predict the original head poses and accurate lip shapes. Given the arduous nature of the task, it is evident that the model will struggle to achieve the desired quality if it relies solely on the LR input. Thus, in our work, we argue that considering appropriate priors is quintessential to obtain high-quality results. Specifically, we assess the importance of two kinds of priors: (i) audio signal, which can help to recover elemental facial attributes and can significantly improve lip shape generation; (ii) a single HR target identity image which can aid in restoring fine-grained details such as skin texture, color, hair, teeth and surrounding background. The target identity image can be any sample frame, either from the same video or any other image with similar characteristics in terms of the face identity, pose, clothing and background. To understand our task better, we will now explore how our work is connected to some of the existing problems in literature.

**Super-Resolution (SR) Perspective:** The SR literature till date has focused on super-resolving inputs (either faces, generic images or videos) [7, 18, 19, 31] where sufficient information is already available (e.g.,  $256 \times 256$  pixels input). None of these methods can handle extreme scale-factors like  $32\times$ . When the input resolution is very low (like  $8 \times 8$ ), we observe that most of the current SR works [6, 7, 18, 19, 31, 36, 43] generate sub-optimal results. The essential visual attributes such as the identity, face texture and lip shape do not accurately match the original face. This is natural because the network is forced to speculate these details without adequate priors. Thus, the existing methods: (i) do not aim to preserve the specific identity details and (ii) do not explicitly deal with talking-face videos where specialized temporal information like lip synchronisation must be maintained throughout the video.

In our work, we propose to address these limitations by generating high-quality talking-face videos with accurate lip-sync. It is important to note that although we aim to generate HR talking-faces from LR inputs, the task we are attempting is very different compared to the typical SR problem. The use of a single HD target identity image and synthesizing from extremely LR talking-faces ( $8 \times 8$  pixels) sets us considerably apart from the traditional SR.

**Compression Perspective:** Our task enables applications such as low-bandwidth video calling; thus making it closely related to the task of talking-face video compression. However, unlike the existing works like “os-synth” [50] where 3D face keypoints are transmitted, we take a unique path in our work. We propose to transmit the LR frames to extract the face structure, motion and pose information. Transmitting keypoints have several limitations: (i) Keypoints can only be extracted if we have the HR video beforehand. While the availability of actual video might be a possibility in

video conferencing, this poses severe constraints in various other applications where the actual HD video is not present; (ii) Keypoints do not encode adequate head pose information, thus requiring additional specialized head pose estimation models; (iii) Keypoints do not cover other information like background, lighting, accessories like glasses and beards that could be present in talking-faces. Thus, in our work, we demonstrate the advantages of using LR frames and achieve better compression ratio compared to the standard codecs, while also not compromising on the desired quality.

**Talking-Face Animation Perspective:** Our task also shares similarities with audio-driven talking-face generation (A2TF) [26, 28, 42, 51, 58] and face re-enactment (FR) [44, 57] tasks. A2TF works aim to generate videos of a target identity conditioned on the audio signal (input: single target identity image + audio). FR methods ingest a single target identity image and an HD video of a different identity with an aim to animate the target image according to the motion of the driving video (input: single target identity image + HD video as pose prior). Although we agree that in terms of the problem space, our work resembles A2TF/FR, we want to point out that our focus is very different - upsampling extremely low-resolution videos. There are key differences as noted in Table 1. We leverage positive aspects from these dimensions to solve an entirely new task - generating HD videos from extremely LR inputs while preserving the exact same facial features, e.g., pose and identity. Nevertheless, we also include a comparison with A2TF and FR works in Section 4.3.

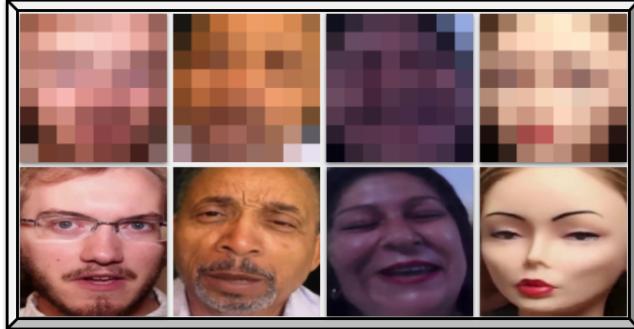
**Table 1: Key differences between audio-driven talking-face generation (A2TF), face re-enactment (FR) and proposed task.**

	Pose prior used	Lips are in-sync with audio?	Matches GT frames?	Can be used for SR?
A2TF	None	✓	✗ (changes pose)	✗
FR	HD video	✓	✓ (same id-recons.)	✗
<b>Ours</b>	LR frames	✓	✓	✓

**Overview of this Work:** In this work, we propose a talking-face video upsampling framework, where the core idea is to utilize adequate priors to generate high-quality ( $256 \times 256$ ) videos from extremely low-resolution inputs. A gallery showing our synthesized frames is displayed in Figure 2. We conduct extensive experiments and comparisons with state-of-the-art methods for the tasks of super-resolution and compression. To the best of our knowledge, we are the first to synthesize high-quality talking-faces at scale-factors of  $32\times$  from an input as small as  $8 \times 8$  pixels. We also show how our network can be utilized for low-bandwidth video conferencing, along with a demo video on our project page.

To summarize, the major contributions of our work are:

- We present a novel audio-visual network that can upsample very low-resolution talking-face videos at scale-factors previously unseen in video SR literature ( $32\times$ ).
- Our approach to make use of adequate priors: (i) audio signal and (ii) a single target identity, achieves significant improvements over the existing works for SR and compression tasks.
- Our system has strong real-world use-cases, owing to its ability to preserve true identity information. We specifically demonstrate the superior quality of our results for low-bandwidth video conferencing and achieve a significant reduction in bandwidth compared to H.264 standard.



**Figure 2: Input: LR frames (top); Output: HR frames (bottom).** Our approach effectively synthesizes a diverse set of HR talking-faces, handling different poses, gender, age and race. Our method also works for “human-like” synthetic faces even though it was trained only on real faces.

## 2 RELATED WORK

**Image & Video Super-Resolution:** Single image super-resolution (SISR) has progressed tremendously with the advent of deep learning and CNNs [13]. A flurry of works [18, 33, 35, 45, 47, 51, 53] have followed this initial method, thereby improving the performance by many folds. To account for specialized facial attributes which are not present in generic images, face super-resolution (FSR) methods [5, 6, 34] with specific loss functions like facial landmarks [36] and facial heatmaps [31] were proposed. However, these FSR approaches focus on scale-factors up to 8 $\times$ , which is in stark contrast to our attempt of 32 $\times$  SR.

Recently, PULSE [37] proposed a StyleGAN-based approach to super-resolve static faces at high scale-factors of 64 $\times$ . However, the model generates imaginary faces of people who do not exist, posing severe constraints on applications where the person’s identity needs to be matched/recovered correctly. Conversely, in our work, we aim to generate HR faces of the specific (real) identity.

To capture the temporal aspect which is not present in the static faces, video SR approaches came into picture. From early recurrent architectures [20] to more recent advancements [7, 19, 43], impressive results have been achieved for scale-factors up to 4 $\times$ . However, these generic video SR methods produce blurred results with many artifacts when the input resolution is very low ( $8 \times 8$ ) and the scale-factors are very high ( $> 8\mathcal{X}$ ). Moreover, “talking-face videos” have their own set of additional challenges, which are neither tackled in video SR nor in static FSR works. Thus, in our work, we overcome these limitations and design a novel network that specifically deals with talking-face videos.

**Talking-Face Video Generation:** Audio-driven talking-face generation (A2TF) is an active research area. Various methods have been proposed [11, 17, 42, 54, 58] to accurately morph the lip movements of input target identity to be in-sync with the corresponding speech. Face re-enactment (FR) works [32, 44, 46] have also shown impressive performance in transferring head motions and expressions based on guiding videos. While our task shares some similarities with A2TF/FR works, we differ in the fact that we are given very sparse information in the form of very low-resolution input. This is not the case in talking-face generation models where HR frames (or extracted landmarks) are used for conditional generation.

Few works [32, 55] use a 3D model as an intermediate step to recover high-quality videos. The major benefit of adapting a 3D model is the superior quality of the output generations. But, it comes with an additional overhead - the computational complexity, making such heavy models very impractical for mobile hardware deployments. In contrast, our method is (i) simple, (ii) generalised, since it can be applied for any in-the-wild identity (unlike some of the 3D models that require speaker-specific training) and (iii) does not require specialized large-scale 3D datasets to train the models.

**Data Compression using Deep Learning:** Deep learning has lead to profound improvements in image and video compression techniques. Starting from initial auto-encoder based methods [14, 52] to flow-based approaches [3, 24], there have been multiple efforts to obtain a compact image/video representation. To enable video calls at a reduced bandwidth, specific methods like [39, 41] have been designed. In “SRVC” [29] authors used video SR as a tool for compression, but at a scale-factor of merely 2 $\times$ . Recently, “os-synth” [50] demonstrated impressive results by transmitting a sequence of learned 3D facial keypoints. However, as discussed previously, transmitting keypoints has its own set of limitations. Thus, unlike the existing compression methods, we take a path of extreme-scale (32 $\times$ ) SR for the first time and transmit very low-resolution ( $8 \times 8$ ) videos. We show that our proposed approach of utilizing the audio and the visual modalities can enable video-conferencing in bandwidth-limited regions, while also achieving a better compression ratio over the existing works.

## 3 LEARNING TO UPSAMPLE LR VIDEOS

We start the discussion by highlighting the critical elements of our approach. We then present our framework with a detailed description of the modules involved.

### 3.1 Critical Elements of our Approach

**Audio Prior:** As discussed previously, when the input resolution is a meager  $8 \times 8$  pixel video, the ambiguity and the loss of information is so paramount that the person’s original identity characteristics are barely discernible. In such situations, we show that audio can aid in the recovery of dominant facial traits of the person because the audio and the face share multiple common features [30, 38, 40] like gender, age and ethnicity. We exploit the audio signal not only to disambiguate the LR input, but also to greatly improve the lip shape generation. Although precise lip shape is not a crucial necessity for static face SR, it is a very important aspect of video SR where the generated lip movements should sync with the given speech. In our work, we explore the natural correlation between speech and lips [1, 10, 21, 42] to generate accurate lip movements.

**Visual Prior:** To be able to generate faces that replicate the actual identity, it is important to preserve sharp details like face texture, lip colour, hair, teeth and skin tone. Most of the current works hallucinate these details, leading to significant variations in fine-grained information. Such models are thus unusable for real-world use-cases, where videos of a specific identity needs to be generated. We argue that considering the adequate prior information is of utmost importance to: (i) generate a video of a specific identity and (ii) reconstruct the high-quality facial details. To achieve this, we provide our network with a single HR image of the target identity,

which helps to transfer the identity-specific sharp details to the synthesized video. In most applications, a single HR identity image is easily accessible. For example, during video conferencing, the first frame can be transmitted in the original resolution.

### 3.2 Our Approach

An overview of our proposed framework is depicted in Figure 3. The goal of our work is to generate a sequence of HR frames,  $F_{hr} = (hr_1, hr_2, \dots, hr_n)$  from the LR input,  $F_{lr} = (lr_1, lr_2, \dots, lr_n)$  that accurately match the ground-truth frames  $F_{gt} = (gt_1, gt_2, \dots, gt_n)$ . We consider the corresponding audio signal,  $A$  and a single HR target identity,  $F_{id}$  as prior information. We detail the different components of our model and the network architecture below.

**Backbone Network:** We pre-train a backbone network to generate a “driving video” for our face animation network. It extracts the relevant features of the identity, face structure, pose and motion information. We observe that pre-training the network before training the face animation network helps to improve the model’s performance because the basic face details and motion information captured is essential to animate the HR face (details about the animation network is explained below). We empirically demonstrate the importance of this in the ablation study (in supplementary).

**Visual Encoder:** We extract the visual features from the LR frames  $F_{lr}$  using a visual encoder, which comprises a series of 3D convolution layers with residual connections. Our visual encoder resembles multiple previous models [2, 10, 21] designed specifically for processing talking-face videos. The input to the visual encoder is a contiguous window of  $T$  LR frames  $l : (N, T, 3, 8, 8)$  where  $N$  refers to the batch size and  $T$  refers to the window frames (here  $T=5$ ). The encoder processes these input frames and generates the visual embeddings,  $f_v : (N, T, 512, 8, 8)$ .

**Audio Encoder:** We consider the corresponding audio segment  $A$  and extract the melspectrogram representation using a window length of 25ms with a hop length of 10ms sampled at 16kHz. The melspectrograms ( $T', 80$ ) are given to the audio encoder, which is a stack of residual 1D convolutions with appropriate strides to match the visual time-steps  $T$ . The generated features ( $N, T, 512$ ) are then upsampled using transpose convolution layers to obtain the audio embeddings,  $f_a : (N, T, 512, 8, 8)$ .

**Identity Structure & Motion Predictor:** We concatenate the learned visual and audio embeddings in the latent space (along the channel dimension) to obtain  $f_{cat} = (N, T, 1024, 8, 8)$ . Inspired from Deep Back-Projection Network (DBPN) [18], we use iterative upsampling and downsampling layers in our module, where the primary idea is to effectively capture the mutual relationship between the LR and HR frames. We consider the fused representation  $f_{cat}$  and stack the time steps along the batch dimension to obtain  $f_{cat} : (N*T, 1024, 8, 8)$ . As the visual and the audio encoders have already captured the temporal information, the stacking strategy improves the convergence speed and also gives us the desired performance. The output of this block is a sequence of frames ( $F_{int}$ ) of resolution  $256 \times 256$  pixels, which encapsulates the elemental facial details, face structure, pose and motion information. The network is trained to minimize the  $L_1$  reconstruction loss:

$$L_{rec} = \frac{1}{N} \sum_{i=1}^N \|F_{int} - F_{gt}\|_1 \quad (1)$$

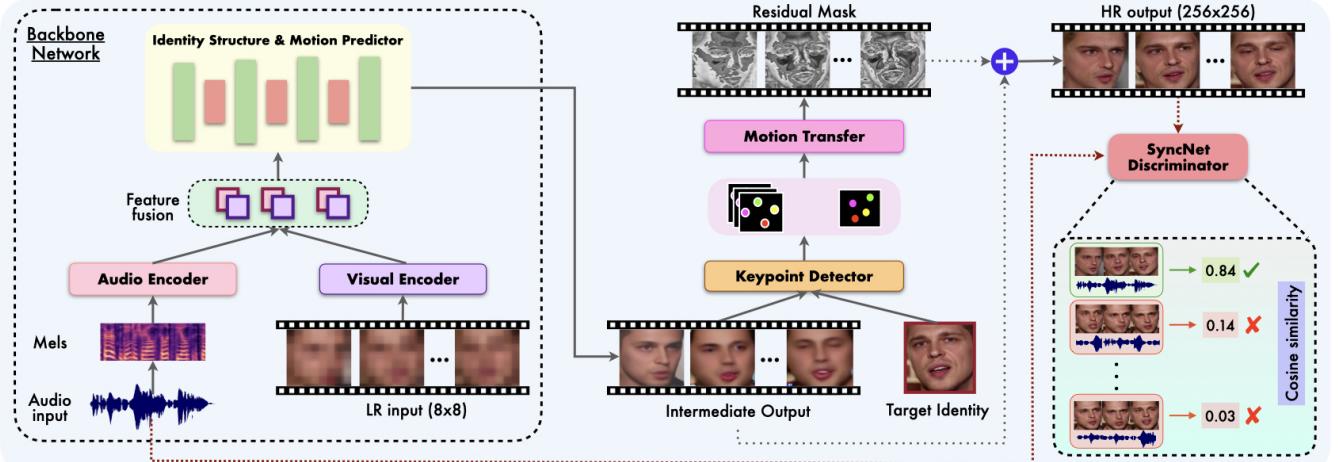
**Face Animation Network:** To synthesize the HR videos of the target identity by preserving all the details, we consider a single HR image of the target identity as our input. This is in-line with the audio-driven talking-face generation works [42, 44, 50, 57] where a single target identity is considered to replicate the identity-specific details. The target identity image helps to capture the fine-grained features like face texture, skin tone, hair and lip color, which are otherwise not recovered in the existing FSR works [6, 18, 36]. These details are crucial, especially when our input is a mere  $8 \times 8$  pixel video and make our model applicable in cases where we must match the actual identity to the maximum extent. We adopt one of the popular methods, FOMM [44] to animate the target identity based on the driving video obtained from our backbone network.

**Overview of First-Order Motion Model (FOMM):** FOMM ingests a target identity image (to extract the appearance) and a driving video (to extract the pose and motion). The learned latent representation of motion in the driving video is combined with the target identity to synthesize the output video. During training, the model observes the target-driving image pairs and predicts a dense motion field which is later encoded using a keypoint detection network. The target image is then rendered according to the learned trajectories in the driving video. We refer the reader to [44] for more details about FOMM. Note that FOMM is designed to work even when the target image and the driving video are of different identities. However, this feature is not necessary in our work where the aim is to preserve the target identity, so we make appropriate modifications, as described below.

**Adapting FOMM to our Task:** For our task at hand, the goal is to animate the HR target identity image  $F_{id}$  in accordance with the motion of the driving video  $F_{int}$ . As described previously, our backbone network reconstructs the basic identity attributes like face structure, age and gender. We thus generate a residual mask as the output of the animation network and add the intermediate outputs  $F_{int}$  (see Figure 3) to obtain a realistic HR talking-face video  $F_{hr}$  of the target identity as the final output. We fine-tune the entire network (including the backbone network) end-to-end, by optimizing the FOMM loss  $L_{fomm}$  and our task-specific losses,  $L_{rec}$ ,  $L_{region}$  and  $L_{sync}$ . The FOMM loss  $L_{fomm}$  consists of: (i) a VGG-19 based perceptual loss at multiple resolutions and (ii) an equivariance constraint to enforce the model to predict consistent keypoints to known geometric transformations. We describe other task-specific losses that we use in our network below.

**Enforcing Local Correspondence:** In our experiments, we observed that the model at times generates frames where the facial regions like eyes, eyebrows and lips are slightly off-position. This occurs if the target identity ingested by our animation network is significantly different from the driving video. Hence, to further improve the generation quality, we add a face landmark-based region loss, which penalizes the network for generating incorrect regions. We compute the face landmarks [4] for both GT  $F_{gt}$  and generated frames  $F_{hr}$  and extract the following  $R$  face regions: lips, nose, eyes and eyebrows. We then add a patch-based local loss by minimizing the  $L_2$  distance for all these  $R$  regions (here  $R=4$ ) to ensure that the predicted regions are as close as possible to the actual regions.

$$L_{region} = \frac{1}{N} \sum_{i=1}^N \sum_{r=1}^R \|F_{hr} - F_{gt}\|_2 \quad (2)$$



**Figure 3:** We propose a novel audio-visual talking-face upsampling framework. Our approach recovers the basic facial attributes like the identity structure, pose and motion using a backbone network. We then consider a single HR target identity image to capture the fine-grained details. Our end-to-end trainable animation network ingests the intermediate outputs (from backbone network) and the target identity and reconstructs realistic, high-quality videos that accurately match the actual identity.

Enforcing Accurate Lip-sync: Prior works [42] involving speech and lip movements have shown that using a lip-sync discriminator can greatly benefit in enforcing strong audio-visual correspondences. This can also be observed in our task, as learning to synthesize HR frames from very low-resolution input might lead to generating lips that are out-of-sync with the audio segment. Thus, we pre-train a lip-sync discriminator, “SyncNet” adapted from [10], trained to maximize the cosine similarity between the lip-speech pair ( $F_{gt}, A$ ) when they are in-sync (and minimise the similarity if they are out-of-sync). Once trained, we use this network as a frozen discriminator to penalize the generated frames  $F_{hr}$  if they do not match the corresponding audio segment. In our end-to-end network, we minimize the sync loss:

$$L_{sync} = -\frac{1}{N} \sum_{i=1}^N \log\left(\frac{f_{hr} \cdot a}{\max(\|f_{hr}\|_2 \cdot \|a\|_2, \epsilon)}\right) \quad (3)$$

### 3.3 Training Settings and Datasets

The final loss function is the combination of the above losses:

$$L_{HR} = \lambda_{rec} L_{rec} + L_{fomm} + \lambda_{region} L_{region} + \lambda_{sync} L_{sync} \quad (4)$$

In our experiments, we set  $\lambda_{rec} = 50$ ,  $\lambda_{region} = 100$  and  $\lambda_{sync} = 0.05$ . We provide the details regarding pre-processing and training settings in supplementary file on our project page.

**Datasets:** We train our model using AVSpeech [16] and VoxCeleb2 [8] datasets; both containing talking-face videos spanning a wide variety of identities, languages and poses. For AVSpeech data, we extract the face tracks using an off-the-shelf face detector [56]. We curate a set of 50 hours for training and  $\sim 3$  hours from the official test split for testing and verified it for accurate lip-sync using SyncNet [9]. We also benchmark our model on VoxCeleb2 data which comprises face tracks with a fair amount of background. Owing to computational limitations, we randomly sample a subset of 100 hours for training and use the full official test split for testing. Note that there are no overlaps between the identities used in training

and testing sets in both datasets. The code, models and file-lists are released on our website for reproducibility and future research.

## 4 EXPERIMENTS

### 4.1 Extreme-scale Super-Resolution

**Baselines:** The state-of-the-art works in video SR literature super-resolve up to a scale-factor of  $4\times$ . We thus re-train the existing state-of-the-art video SR method, TecoGAN [7] at a scale-factor of  $32\times$  on the same training dataset as ours. We extend the existing face SR approach, SPARNet [6] to work for video SR by appropriately modifying the architecture (ingest a window of 5 frames) and train using the same settings as mentioned above. These methods originally do not consider a HR target identity as input; thus, it would be unfair to compare them without the identity information. Hence we provide a HR target identity image to these models in a manner typically used in talking-face generation methods: concatenating it channel-wise with the input.

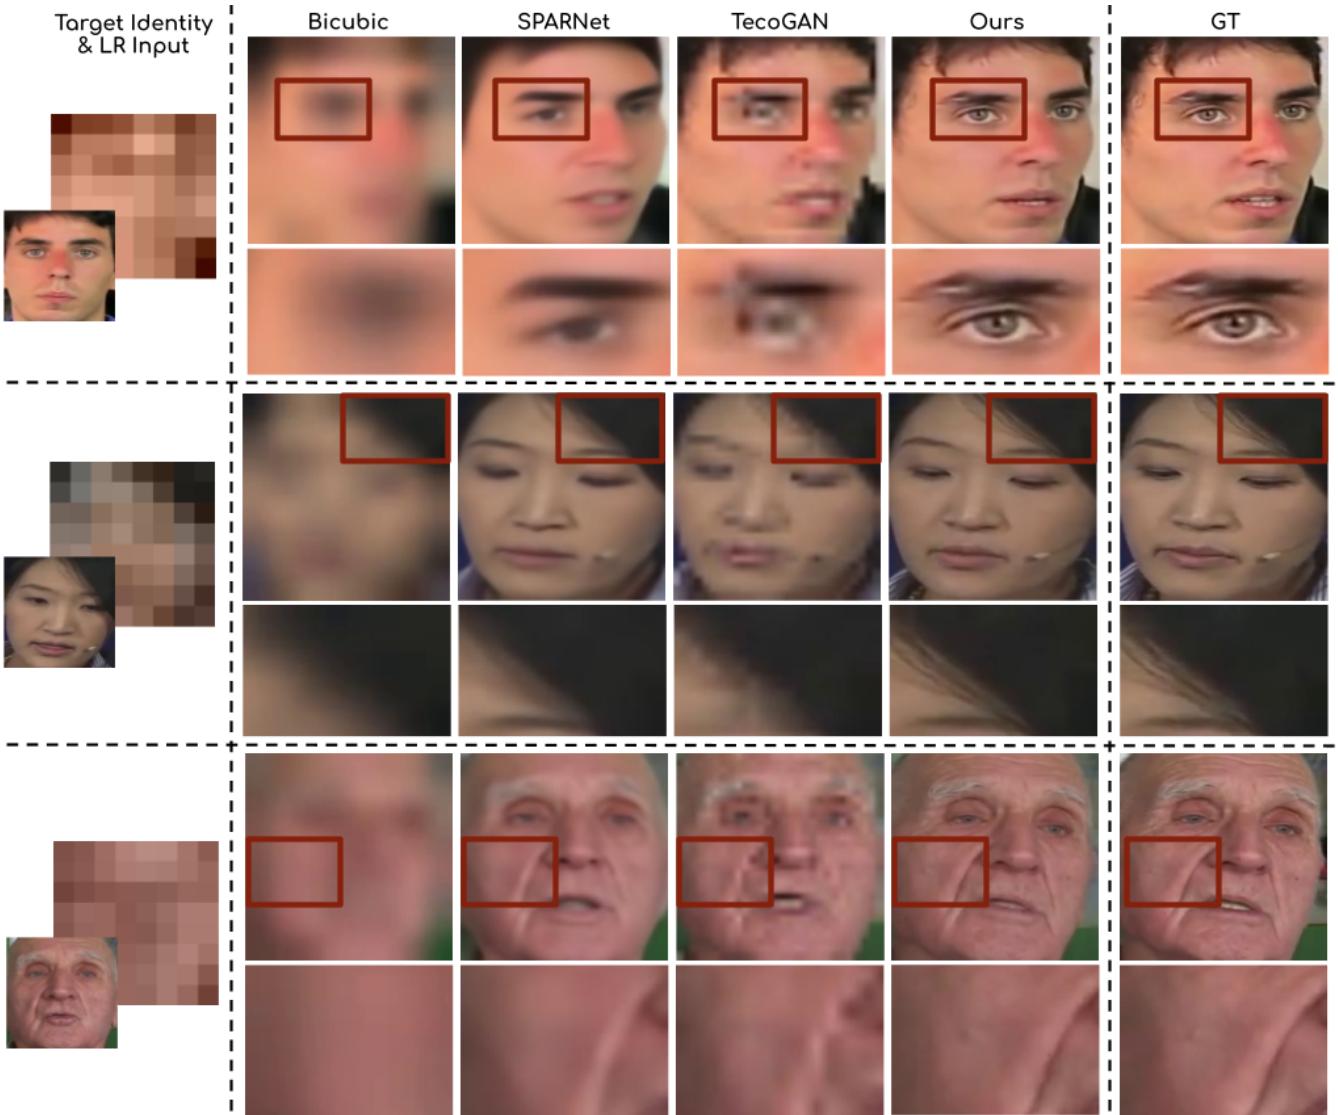
**Metrics:** We evaluate our SR model on: (i) PSNR, (ii) SSIM, (iii) Fréchet Inception Distance (FID) [23], (iv) Landmark Distance (LMD) [4] and (v) Lip-Sync Error Distance (LSE-D) [42]. More details about the metrics can be found in the supplementary file.

**Results:** We compare our results with existing SR approaches at extreme scale-factor of  $32\times$  in Table 2. As we can see from the table, our method outperforms the existing works by a significant margin on both AVSpeech [16] and VoxCeleb2 [8] datasets. None of the current techniques match the ground-truth identity (measured using PSNR) and perceptual quality (measured using FID) of our generations. The LSE-D metric indicates that our method achieves accurate lip-synchronization with audio, thus validating our claim that the audio signal enables us to generate far more accurate lip shapes than the competing methods. Our method also surpasses the existing approaches in preserving the overall face structure (measured using SSIM and LMD).

Figures 4 and 5 show the qualitative comparisons. We can clearly observe that our models generate results with far fewer artifacts

**Table 2: Quantitative comparison for 32 $\times$  SR on AVSpeech [16] and VoxCeleb2 [8] datasets. Our method outperforms the baselines by a significant margin across all metrics. Note that the baselines have also been trained with a single identity image.**

Dataset	AVSpeech [16]					VoxCeleb2 [8]				
	PSNR $\uparrow$	SSIM $\uparrow$	FID $\downarrow$	LMD $\downarrow$	LSE-D $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	FID $\downarrow$	LMD $\downarrow$	LSE-D $\downarrow$
Bicubic	22.33	0.60	102.41	0.246	14.18	22.16	0.60	105.14	0.255	17.83
SPARNet [6]	23.17	0.68	92.14	0.201	12.87	22.98	0.67	83.01	0.228	14.07
TecoGAN [7]	19.26	0.62	84.73	0.213	13.01	16.91	0.54	82.19	0.234	14.12
<b>Ours</b>	<b>25.06</b>	<b>0.73</b>	<b>11.54</b>	<b>0.162</b>	<b>12.43</b>	<b>24.95</b>	<b>0.71</b>	<b>14.10</b>	<b>0.196</b>	<b>13.91</b>



**Figure 4: Qualitative comparisons on AVSpeech dataset [16]. Our method captures the rich identity-specific attributes like eyeballs, hair strands, face texture and lip shape, far better compared to the existing approaches.**

and captures rich, fine-grained details. Although all the comparison methods consider the HR target identity as input, they do not match the quality of our generations. This shows that our overall network design is highly effective in making use of the available target identity image. In the examples, we can also see the diverse range of our models' generative capabilities: eyeballs with precise eye color (Fig. 4: row 1), microphone (Fig. 4: row 2), hair strands (Fig. 4:

row 2), lip shape (Fig. 4: rows 2,3 and Fig. 5: rows 1,2), face texture such as wrinkles (Fig. 4: row 3), beard (Fig. 5: rows 1,2). More visual examples can be found on our project page.

**Ablation study:** We validate the design choices of our network by analyzing the importance of audio signal, landmark-based region loss, the use of different target identity images and several other

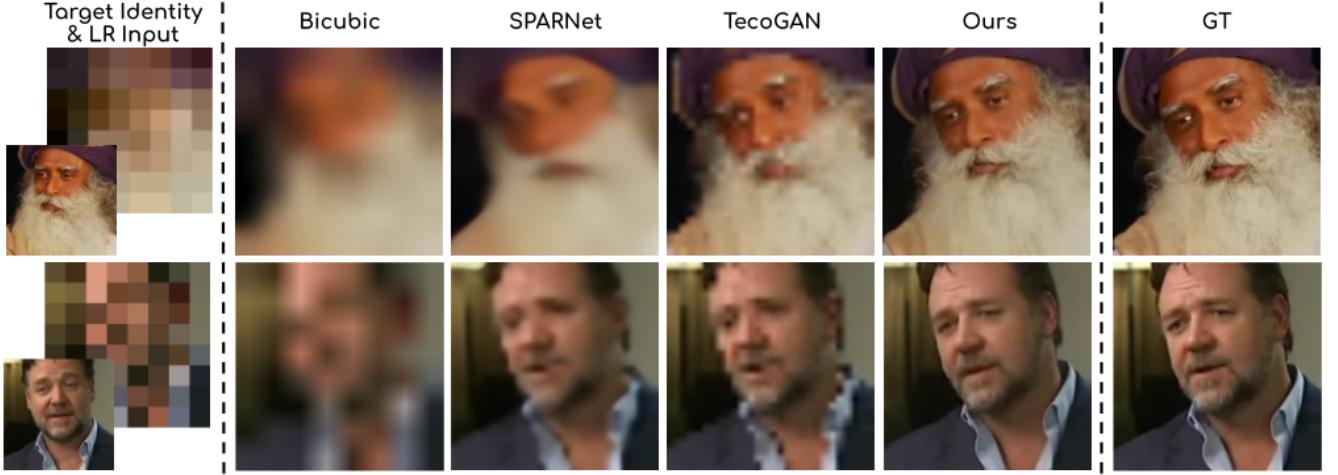


Figure 5: Qualitative comparisons on the models trained on VoxCeleb2 dataset [8]. Our method surpasses the existing baselines in generating the outputs that accurately match the ground-truth identity.

Table 3: Quantitative comparison for talking-face video compression on VoxCeleb2 dataset [8]. We achieve the best trade-off in terms of quality versus compression ratio. Our method achieves the lowest FID (indicates very high perceptual quality) and a very low/comparable BPP.

Method	BPP↓	PSNR↑	SSIM↑	FID↓
H.264 (CRF=23) (min. compression)	0.109	32.96	0.79	9.75
H.264 (CRF=36)	0.027	19.24	0.67	30.12
H.266	0.0076	23.27	0.70	58.32
fs-vid2vid [49]	n/a	20.36	0.71	85.76
os-synth [50]	0.016	24.37	<b>0.80</b>	69.13
<b>Ours</b>	0.023	<b>24.95</b>	0.71	<b>14.10</b>
<b>Ours (Frame-Interpolation)</b>	<b>0.0046</b>	23.72	0.68	14.51

additional experiments, along with human evaluations in our supplementary. We also compare the performance of different models at multiple scale-factors like 4×, 8×, 16× and 32× in supplementary.

## 4.2 Talking-Face Video Compression

One of the major applications of our system is in compressing talking-face videos to reduce the bandwidth in video conferencing applications. We can transmit the LR frames ( $8 \times 8$  pixels) along with the audio signal on the sender's side and the receiver can reconstruct the high-quality video ( $256 \times 256$  pixels) using a single HR target identity image. A sample video calling demo is illustrated in Figure 6. We assume that a single target identity image can be sent at the beginning (e.g., 1<sup>st</sup> frame) and hence does not consume additional bandwidth. Note that this is very different from the standard codecs, where full resolution I-frames are transmitted at regular intervals. Also, since the audio signal is always accompanied in a video call, we do not consider this as an extra overhead.

**Baselines:** We benchmark our model's capability using the existing talking-face video compression methods: few-shot vid2vid (fs-vid2vid) [49], one-shot free-view synthesis (os-synth) [50] and the standard codecs: (i) H.264 (with CRF of 23 and 36) and (ii) H.266 (implemented using vvenc: <https://github.com/fraunhoferhhi/vvenc>).

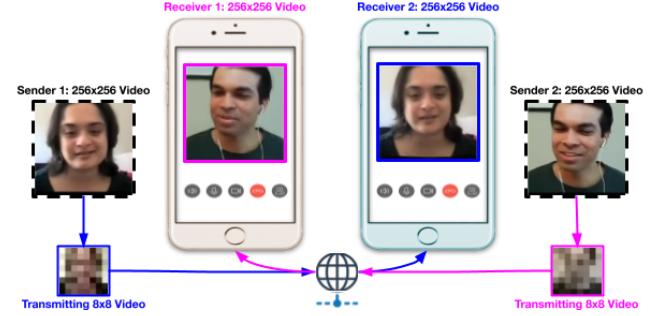


Figure 6: Illustration of low-bandwidth video calling enabled by our system. Note that HR frames shown at both the sender's end are taken from an actual video call recording (credits: <https://www.youtube.com/watch?v=lQJD8RAq3IY>).

Since we train and evaluate on the same dataset (VoxCeleb2 [8]), we directly take the scores reported in os-synth [50] for comparison.

**Metrics:** We compare the compression factor using the standard bits-per-pixels (BPP) metric and measure the reconstruction quality using PSNR, SSIM and FID metrics.

**Results:** Table 3 shows the comparison of our approach with the competing methods. We calculate the average BPP across all test videos for H.264 and H.266 codecs. The number of bits required by the current state-of-the-art "os-synth" [50] to represent a  $256 \times 256$  image is 1056 (20 keypoints:  $(20 \times 6 + 12) \times 8$ ). Our approach requires 1536 bits ( $8 \times 8 \times 3 \times 8$ ), with a BPP of 0.023. The BPP obtained using our method beats H.264 and is comparable to os-synth [50]. In terms of perceptual quality, we can observe that there is a trade-off in compression factor v/s quality for H.264, (i) CRF of 23: Higher quality, with very less compression and (ii) CRF of 36: Better compression, but poor reconstruction. While H.266 achieves comparable PSNR and SSIM measures with a BPP of 0.0076, it still lags behind perceptually as shown by the FID metric. In contrast, our method is able to generate higher quality videos (remarkably low FID), with a better/comparable compression factor.

**Computation Comparison:** Our network has ~ 20% fewer parameters (143M) while being 2× faster (50FPS) than state-of-the-art



**Figure 7: Failure cases: (a) Drastic color changes, (b) Extreme variations of facial expressions and (c) Rapid head motions.**

compression method (os-synth) on an NVIDIA 2080Ti GPU. Previous talking-face generation works focus on obtaining plausible faces, whereas our aim is to develop the first architecture to reconstruct the specific-identity by preserving most of the actual identity details; future works can build upon this to achieve optimal design. Further, model optimization techniques can be applied to the current design to make it suitable for mobile hardware deployment.

**Frame-Interpolation Network:** To further reduce the bandwidth consumption, we develop a baseline frame-interpolation network with an aim to upsample 5FPS videos to 25FPS. While such networks have been studied in literature [27], most of them restrict the amount of upsampling, mainly due to the computational resources involved in training the bulky networks (often consisting of 3D CNNs). In contrast, we take the advantage of LR videos and design a model to increase the temporal resolution of frames (5FPS to 25FPS) for the first time. Our encoder-decoder based model processes 5FPS LR frames (5, 3, 8, 8) and upsamples it 5× to generate 25FPS LR frames as output (25, 3, 8, 8). We refer the reader to our supplementary for more details. The lower resolution of both the input and output allows us to train a very light 3D CNN model with only 0.2M parameters. This strategy effectively permits us to transmit 5FPS LR videos. On receiving these frames, frame-interpolation network initially upsamples them to 25FPS LR videos, which can subsequently be ingested by our spatial talking-face video upsampling network to render the final HR video. As shown in Table 3, our frame-interpolation network achieves a further reduction in BPP, since only 1 in 5 LR frames needs to be transferred. We thus obtain  $\sim 6\times$  and  $\sim 25\times$  reduction in bandwidth compared to os-synth [50] and H.264 codec (with CRF 23) respectively, without a significant compromise in the generated quality.

### 4.3 Audio-driven Talking-Face Generation (A2TF) and Face Re-enactment (FR)

**Baselines:** For A2TF, we compare with Wav2Lip [42] and MakeITalk [58]. Talking-face videos are generated using the audio segment and the first frame of the original HD video as inputs (same strategy as our model). For FR, the original models take the actual HD frames as the driving video input, however, we do not have access to the actual HD video in our model. Thus, we upsample the LR input using the existing video SR method, TecoGAN [7] (trained on VoxCeleb2 dataset) and use it as our input driving video, along with the first frame of the original HD video as target identity input.

**Metrics:** Along with the standard metrics, we specifically evaluate the ability of models to match the original identity using head pose estimation metrics [12, 22]. We convert the rotation matrix<sup>2</sup> to Euler angles and report Mean Absolute Error (MAE) of these angles.

<sup>2</sup><https://learnopencv.com/head-pose-estimation-using-opencv-and-dlib/>

**Results:** Table 4 shows the performance of different A2TF and FR methods in comparison to our approach. A2TF models generate accurate lip-sync (LSE-D), but they fail to match the pose of the original identity (as noted by Yaw, Pitch, Roll and MAE metrics). FR methods also face similar issues and do not match the identity, as they do not have the actual HD video as driving frames. Our model is clearly the most applicable for this task, is consistent across all metrics and produces noticeably better results. Most importantly, our model preserves the identity to the maximum extent.

**Table 4: Quantitative comparison with A2TF [42, 58] and FR [44, 57] methods on VoxCeleb2 test set.**

Method	PSNR↑	SSIM↑	FID↓	LMD↓	LSE-D↓	Yaw↓	Pitch↓	Roll↓	MAE↓
Wav2Lip [42]	14.18	0.32	<b>8.15</b>	4.320	9.19	24.68	38.31	28.94	30.64
MakeITalk [58]	18.88	0.49	31.19	2.012	11.91	26.29	40.13	31.42	32.61
FOMM [44]	20.14	0.56	21.18	0.864	14.03	19.14	30.57	22.76	24.35
PC-AVS [57]	15.68	0.37	33.38	1.063	<b>8.42</b>	22.27	31.89	25.80	26.65
<b>Ours</b>	<b>24.95</b>	<b>0.71</b>	14.10	<b>0.196</b>	13.91	<b>13.55</b>	<b>21.01</b>	<b>15.48</b>	<b>16.68</b>

## 5 LIMITATIONS AND FUTURE DIRECTIONS

Although our method generates realistic results for a wide variety of inputs, there are certain situations as shown in Figure 7 where our model results in sub-optimal generations. For example, if the color contrast changes drastically as the video progresses, the model fails to capture these details (Figure 7 (a)). Significant variations in facial expressions is another case where our model struggles to replicate the details (Figure 7 (b)). Explicitly handling the expressions is an interesting direction that can be investigated in the future, which we currently do not handle in our work. Finally, in the case of sudden/rapid changes in view, camera angle, or head movements, our model attempts to generate smooth transitions (Figure 7 (c)). However, we found our method to be stable over a large variety of inputs and anticipate that our idea of utilizing extremely LR frames will be a basis for other domains and applications.

## 6 CONCLUSION

In this work, we present a novel framework for extreme-scale talking-face video super-resolution and compression. We show that by considering appropriate priors (audio signal and a single target identity image), we are able to generate realistic, high-quality talking-face videos from (very) low-resolution frames. Our method handles various inputs, including but not limited to people of different ages, gender and ethnicity. Most importantly, our framework is the first of its kind to produce photo-realistic lip-synced talking-face videos while also matching the actual identity. By dramatically reducing the bandwidth requirements, our approach can be utilized as a tool for a seamless video-conferencing experience. We believe our core idea of exploiting very low-resolution videos along with adequate priors will be an important step towards the future of super-resolution and low-bandwidth video-conferencing.

## REFERENCES

- [1] T. Afouras, J. S. Chung, and A. Zisserman. 2018. The Conversation: Deep Audio-Visual Speech Enhancement. In *INTERSPEECH*.
- [2] Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman. 2020. Self-Supervised Learning of Audio-Visual Objects from Video. In *European Conference on Computer Vision*.
- [3] Eirikur Agustsson, David Minnen, Nick Johnston, Johannes Balle, Sung Jin Hwang, and George Toderici. 2020. Scale-Space Flow for End-to-End Optimized Video Compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [4] Adrian Bulat and Georgios Tzimiropoulos. 2017. How far are we from solving the 2D & 3D Face Alignment problem? (and a dataset of 230,000 3D facial landmarks). In *International Conference on Computer Vision*.
- [5] Adrian Bulat, Jing Yang, and Georgios Tzimiropoulos. 2018. To learn image super-resolution, use a GAN to learn how to do image degradation first. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 185–200.
- [6] Chaofeng Chen, Dihong Gong, Hao Wang, Zhifeng Li, and Kwan-Yee K. Wong. 2020. Learning Spatial Attention for Face Super-Resolution. *IEEE Transactions on Image Processing (TIP)*.
- [7] Mengyu Chu, You Xie, Jonas Mayer, Laura Leal-Taixé, and Nils Thuerey. 2020. Learning Temporal Coherence via Self-Supervision for GAN-Based Video Generation. *ACM Trans. Graph.* 39 (July 2020), 13 pages. <https://doi.org/10.1145/3386569.3392457>
- [8] J. S. Chung, A. Nagrani, and A. Zisserman. 2018. VoxCeleb2: Deep Speaker Recognition. In *INTERSPEECH*.
- [9] J. S. Chung and A. Zisserman. 2016. Out of time: automated lip sync in the wild. In *Workshop on Multi-view Lip-reading, ACCV*.
- [10] S. Chung, J. S. Chung, and H. Kang. 2019. Perfect Match: Improved Cross-modal Embeddings for Audio-visual Synchronisation. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 3965–3969. <https://doi.org/10.1109/ICASSP.2019.8682524>
- [11] Filipe Antonio de Barros Reis, Paula Dornhofer Paro Costa, and José Mario de Martino. 2020. Deeply Emotional Talking Head: A Generative Adversarial Network Approach to Expressive Speech Synthesis with Emotion Control. In *ACM SIGGRAPH 2020 Posters* (Virtual Event, USA) (*SIGGRAPH '20*). Association for Computing Machinery, New York, NY, USA, Article 33, 2 pages. <https://doi.org/10.1145/3388770.3407417>
- [12] Naina Dhingra. 2022. LwPosr: Lightweight Efficient Fine Grained Head Pose Estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 1495–1505.
- [13] Chao Dong, Chen Change Loy, Kaiming He, and Xiaou Tang. 2016. Image Super-Resolution Using Deep Convolutional Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38, 2 (2016), 295–307. <https://doi.org/10.1109/TPAMI.2015.2439281>
- [14] Thierry Dumas, Aline Roumy, and Christine Guillemot. 2018. Autoencoder Based Image Compression: Can the Learning be Quantization Independent? *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (2018), 1188–1192.
- [15] Efros, Berg, Mori, and Malik. 2003. Recognizing action at a distance. In *Proceedings Ninth IEEE International Conference on Computer Vision*. 726–733 vol.2. <https://doi.org/10.1109/ICCV.2003.1238420>
- [16] Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T. Freeman, and Michael Rubinstein. 2018. Looking to Listen at the Cocktail Party: A Speaker-Independent Audio-Visual Model for Speech Separation. *ACM Trans. Graph.* 37, 4, Article 112 (2018), 11 pages. <https://doi.org/10.1145/3197517.3201357>
- [17] Ohad Fried, Ayush Tewari, Michael Zollhöfer, Adam Finkelstein, Eli Shechtman, Dan B Goldman, Kyle Genova, Zeyu Jin, Christian Theobalt, and Maneesh Agrawala. 2019. Text-Based Editing of Talking-Head Video. *ACM Trans. Graph.* 38, Article 68 (2019), 14 pages. <https://doi.org/10.1145/3306346.3323028>
- [18] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. 2018. Deep Back-Projection Networks for Super-Resolution. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1664–1673. <https://doi.org/10.1109/CVPR.2018.00179>
- [19] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. 2019. Recurrent Back-Projection Network for Video Super-Resolution. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [20] Patrick Héas, Angélique Dréneau, and Cédric Herzet. 2015. Estimation of Super-Resolved Video Dynamics. *CoRR* abs/1506.00473 (2015). arXiv:1506.00473
- [21] Sindhu B. Hegde, K.R. Prajwal, Rudrabha Mukhopadhyay, Vinay P. Namboodiri, and C.V. Jawahar. 2021. Visual Speech Enhancement Without a Real Visual Stream. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. 1926–1935.
- [22] Thorsten Hempel, Ahmed A Abdelaal, and Ayoub Al-Hamadi. 2022. 6D Rotation Representation For Unconstrained Head Pose Estimation. *arXiv preprint arXiv:2202.12555* (2022).
- [23] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2018. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. *arXiv:1706.08500 [cs.LG]*
- [24] Jonathan Ho, Evan Lohn, and P. Abbeel. 2019. Compression with Flows via Local Bits-Back Coding. In *NeurIPS*.
- [25] Peiyun Hu and Deva Ramanan. 2017. Finding Tiny Faces. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [26] Amir Jamaludin, Joon Son Chung, and Andrew Zisserman. 2019. You said that? : Synthesising talking faces from audio. *International Journal of Computer Vision* (2019).
- [27] Huaizu Jiang, Deqing Sun, V. Jampani, Ming-Hsuan Yang, Erik G. Learned-Miller, and Jan Kautz. 2018. Super SloMo: High Quality Estimation of Multiple Intermediate Frames for Video Interpolation. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 9000–9008.
- [28] Prajwal K R, Rudrabha Mukhopadhyay, Jerin Philip, Abhishek Jha, Vinay Namboodiri, and C V Jawahar. 2019. Towards Automatic Face-to-Face Translation. In *Proceedings of the 27th ACM International Conference on Multimedia* (Nice, France) (*MM '19*). ACM, 9 pages. <https://doi.org/10.1145/3343031.3351066>
- [29] Mehrdad Khani, Vibhaalakshmi Sivaraman, and Mohammad Alizadeh. 2021. Efficient Video Compression via Content-Adaptive Super-Resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 4521–4530.
- [30] Changil Kim, Hijung Valentina Shin, Tae-Hyun Oh, Alexandre Kaspar, Mohamed Elgharib, and Wojciech Matusik. 2018. On Learning Associations of Faces and Voices. In *Proceedings of Asian Conference on Computer Vision (ACCV)*.
- [31] Deokyun Kim, Minseon Kim, Gihyun Kwon, and Dae-Shik Kim. 2019. Progressive Face Super-Resolution via Attention to Facial Landmark. In *Proceedings of the 30th British Machine Vision Conference (BMVC)*.
- [32] H. Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Nießner, P. Pérez, Christian Richardt, M. Zollhöfer, and C. Theobalt. 2018. Deep video portraits. *ACM Transactions on Graphics (TOG)* 37 (2018), 1 – 14.
- [33] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. 2015. Accurate Image Super-Resolution Using Very Deep Convolutional Networks. *CoRR* abs/1511.04587 (2015). arXiv:1511.04587 <http://arxiv.org/abs/1511.04587>
- [34] Jonghyun Kim, Gen Li, Inyong Yun, Cheolkon Jung, and Joongkyu Kim. 2021. Edge and identity preserving network for face super-resolution. *Neurocomputing* 446 (2021), 11–22. <https://doi.org/10.1016/j.neucom.2021.03.048>
- [35] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. 2017. Enhanced Deep Residual Networks for Single Image Super-Resolution. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 1132–1140. <https://doi.org/10.1109/CVPRW.2017.151>
- [36] Cheng Ma, Zhenyu Jiang, Yongming Rao, Jiwen Lu, and J. Zhou. 2020. Deep Face Super-Resolution With Iterative Collaboration Between Attentive Recovery and Landmark Estimation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020), 5568–5577.
- [37] Sachit Menon, Alex Damian, McCourt Hu, Nikhil Ravi, and Cynthia Rudin. 2020. PULSE: Self-Supervised Photo Upsampling via Latent Space Exploration of Generative Models. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [38] Arsha Nagrani, Samuel Albanie, and Andrew Zisserman. 2018. Seeing Voices and Hearing Faces: Cross-Modal Biometric Matching. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2018), 8427–8436.
- [39] Arun Shankar Narayanan and Tan Kok Kiong. 2013. Video conferencing solution for medical applications in low-bandwidth networks. In *2013 Australasian Telecommunication Networks and Applications Conference (ATNAC)*. 195–200. <https://doi.org/10.1109/ATNAC.2013.6705380>
- [40] Tae-Hyun Oh, Tali Dekel, Changil Kim, Inbar Mosseri, William T. Freeman, Michael Rubinstein, and Wojciech Matusik. 2019. Speech2Face: Learning the Face Behind a Voice. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 7531–7540. <https://doi.org/10.1109/CVPR.2019.00772>
- [41] Maxime Oquab, Pierre Stock, Oran Gafni, Daniel Haziza, Tao Xu, Peizhao Zhang, Onur Çelebi, Yana Hasson, Patrick Labatut, Bobo Bose-Kolanu, Thibault Peyronel, and Camille Couprie. 2021. Low Bandwidth Video-Chat Compression using Deep Generative Models. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)* (2021), 2388–2397.
- [42] K R Prajwal, Rudrabha Mukhopadhyay, Vinay P. Namboodiri, and C.V. Jawahar. 2020. A Lip Sync Expert Is All You Need for Speech to Lip Generation In the Wild. In *Proceedings of the 28th ACM International Conference on Multimedia* (*MM '20*), 484–492. <https://doi.org/10.1145/3394171.3413532>
- [43] Mehdi S. M. Sajjadi, Raviteja Vemulapalli, and Matthew Brown. 2018. Frame-Recurrent Video Super-Resolution. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [44] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. 2019. First Order Motion Model for Image Animation. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- [45] Jiaxiang Tang, Xiaokang Chen, and Gang Zeng. 2021. *Joint Implicit Image Function for Guided Depth Super-Resolution*. Association for Computing Machinery, New York, NY, USA, 4390–4399. <https://doi.org/10.1145/3474085.3475584>
- [46] Justus Thies, Michael Zollhöfer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. 2018. Face2Face: Real-Time Face Capture and Reenactment of RGB Videos. *Commun. ACM* 62, 1 (Dec. 2018), 96–104. <https://doi.org/10.1145/3197517.3201357>

- 3292039
- [47] Radu Timofte, Shuhang Gu, Jiqing Wu, and Luc Van Gool. 2018. NTIRE 2018 Challenge on Single Image Super-Resolution: Methods and Results. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
  - [48] Antonio Torralba, Kevin P. Murphy, and William T. Freeman. 2007. Sharing Visual Features for Multiclass and Multiview Object Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29, 5 (2007), 854–869. <https://doi.org/10.1109/TPAMI.2007.1055>
  - [49] Ting-Chun Wang, Ming-Yu Liu, Andrew Tao, Guilin Liu, Jan Kautz, and Bryan Catanzaro. 2019. Few-shot Video-to-Video Synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*.
  - [50] Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. 2021. One-Shot Free-View Neural Talking-Head Synthesis for Video Conferencing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
  - [51] Haozhe Wu, Jia Jia, Haoyu Wang, Yishun Dou, Chao Duan, and Qingshan Deng. 2021. *Imitating Arbitrary Talking Style for Realistic Audio-Driven Talking Face Synthesis*. Association for Computing Machinery, New York, NY, USA, 1478–1486. <https://doi.org/10.1145/3474085.3475280>
  - [52] Li Xu, Jimmy SJ Ren, Ce Liu, and Jiaya Jia. 2014. Deep Convolutional Neural Network for Image Deconvolution. In *Advances in Neural Information Processing Systems*, Vol. 27. Curran Associates, Inc.
  - [53] Bin-Cheng Yang. 2019. Super Resolution Using Dual Path Connections (*MM '19*). Association for Computing Machinery, New York, NY, USA, 1552–1560. <https://doi.org/10.1145/3343031.3350878>
  - [54] Xin-Wei Yao, Ohad Fried, K. Fatahalian, and Maneesh Agrawala. 2020. Iterative Text-based Editing of Talking-heads Using Neural Retargeting. *ArXiv* abs/2011.10688 (2020).
  - [55] Chenxu Zhang, Yifan Zhao, Yifei Huang, Ming Zeng, Saifeng Ni, Madhukar Budagavi, and Xiaohu Guo. 2021. FACIAL: Synthesizing Dynamic Talking Face with Implicit Attribute Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. 3867–3876.
  - [56] Shifeng Zhang, Xiangyu Zhu, Zhen Lei, Hailin Shi, Xiaobo Wang, and S. Li. 2017. S<sup>3</sup>FD: Single Shot Scale-Invariant Face Detector. *2017 IEEE International Conference on Computer Vision (ICCV)* (2017), 192–201.
  - [57] Hang Zhou, Yasheng Sun, Wayne Wu, Chen Change Loy, Xiaogang Wang, and Ziwei Liu. 2021. Pose-Controllable Talking Face Generation by Implicitly Modularized Audio-Visual Representation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
  - [58] Yang Zhou, Xintong Han, Eli Shechtman, Jose Echevarria, Evangelos Kalogerakis, and Dingzeyu Li. 2020. MakeItTalk: Speaker-Aware Talking-Head Animation. *ACM Transactions on Graphics* 39, 6 (2020).