

Diffused Heads: Diffusion Models Beat GANs on Talking-Face Generation

Michał Stypułkowski¹

michal.stypulkowski@cs.uni.wroc.pl

Konstantinos Vougioukas²

k.vougioukas@imperial.ac.uk

Sen He

senhe752@gmail.com

Maciej Zięba^{3,4}

maciej.zieba@pwr.edu.pl

Stavros Petridis²

sp104@imperial.ac.uk

Maja Pantic²

m.pantic@imperial.ac.uk

¹University of Wrocław

²Imperial College London

³Wrocław University of Science and Technology

⁴Tooploox

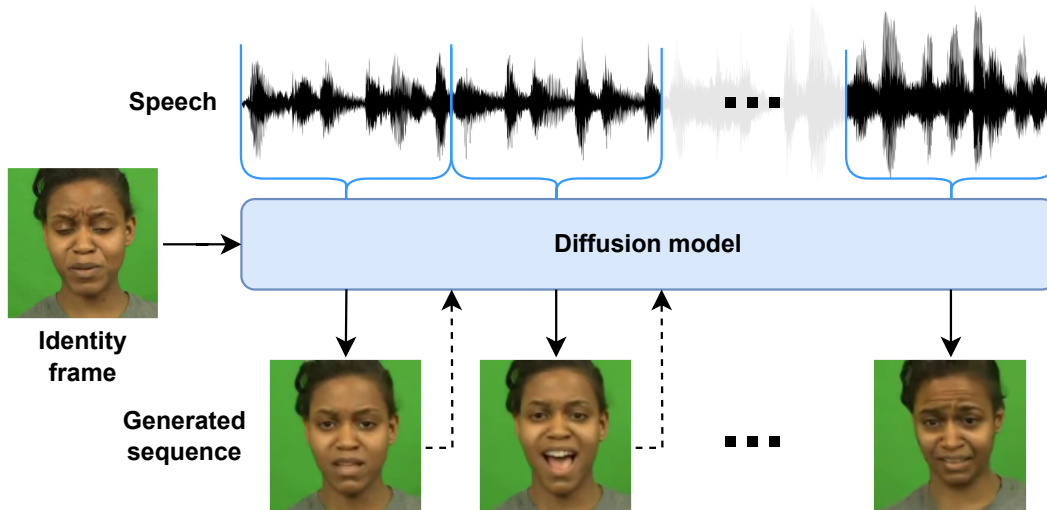


Figure 1. Overview of the proposed approach. Given a single identity frame and an audio clip containing speech, the model uses a diffusion model to sample consecutive frames in an autoregressive manner, preserving the identity, and modeling lip and head movement to match the audio input. Contrary to other methods, no additional guidance is required.

Abstract

Talking face generation has historically struggled to produce head movements and natural facial expressions without guidance from additional reference videos. Recent developments in diffusion-based generative models allow for more realistic and stable data synthesis and their performance on image and video generation has surpassed that of other generative models. In this work, we present an autoregressive diffusion model that requires only one identity image and audio sequence to generate a video of a realistic talking head. Our solution is capable of hallucinating head movements, facial expressions, such as blinks, and preserving a given background. We evaluate our model on two different datasets, achieving state-of-the-art results in

expressiveness and smoothness on both of them.¹

1. Introduction

Animation of faces from speech can have a broad scope of applications from an alternative to video compression during virtual calls with poor connectivity, to artistic animation for entertainment industry applications, e.g. movies, video games, and VR experience. Up to date, existing methods struggle to create naturally-looking faces that maintain genuine expressions and movements, while still requiring additional supervision during the generation process.

Deep generative models are constantly gaining popularity and achieving impressive results in image and video gen-

¹Project page: <https://mstypulkowski.github.io/diffusedheads/>.

eration tasks and have become the *de facto* standard for most facial animation systems. In particular speech-driven facial animation systems, which are a simple and effective way of producing character animations, have been revolutionized by the introduction of recent generative models such as Generative Adversarial Networks (GANs) [9]. GANs are known for being able to produce high-quality frames while simultaneously giving a large degree of control over the generation process [20, 24, 29].

Despite the powerful capabilities of GANs, their application to speech-driven video synthesis has several drawbacks. Firstly, GAN training is notoriously difficult, often requiring an extensive architectural search and parameter tuning to achieve convergence. The training stability of GAN-based facial animation methods can be improved through the use of additional guidance such as masks or driving frames to guide the generation process. However, this limits them to applications of facial reenactment and reduces their ability to produce original head motions and facial expressions. Furthermore, GAN training can often lead to mode collapse, *i.e.* a situation when the generator can't produce samples that cover the entire support of the data distribution and instead learns to generate only a few unique samples [1]. Finally, existing one-shot GAN-based solutions have problems with face distortion in the generated videos, especially when generating videos with large head motions. This is often solved by either switching to a few-shot approach (*i.e.* using several frames or a short clip) or relying on pre-trained face verification models that serve as oracles for maintaining identity consistency.

We address all of the above problems, proposing Diffused Heads - a frame-based diffusion model that produces realistic videos, requiring only one identity frame and a speech recording. Generated heads move and behave in a natural expressive way while still preserving the subject's identity and plausible lip sync. In contrast to most recent approaches [3, 18, 21, 27, 32, 39, 45, 49–51], we use Denoising Diffusion Probabilistic Models [15, 23] that utilize a variational approach instead of adversarial training and do not require stabilizing discriminators [27, 45, 50, 51]. To eliminate the problem of unnaturally-looking sequences, we introduce motion frames (see Section 4.2) that are guiding video creation. To maintain the consistency between the speech and the generated frames, we postulate to use audio embeddings extracted from a pre-trained temporal model injected into the model via our novel conditioning approach. Finally, instead of using a pre-trained oracle model, we introduce a simple modification of the loss function to preserve the consistency of the lip movement.

Contributions of our work are summarized as follows: 1) To the best of our knowledge, we present the first solution for talking-face generation based on diffusion models. 2) We enrich the diffusion model with motion frames

and audio embeddings in order to maintain the consistency of generated images. 3) Our approach is robust in terms of generalization, invariant on the source of identity frames and audio recordings.

2. Related work

The problem of speech-driven video synthesis was initially investigated in [48], where the authors discovered a strong correlation between acoustic and video features. Some of the earliest approaches utilized Hidden Markov Models (HMMs) to capture the dynamics of the video and speech sequences [36, 46, 47]. The authors of [36] used the compact feature representations of speech and video jointly as states of the fully-connected Markov model. In [47] the authors used HMMs to estimate the sequence of lip parameters. The authors of [46] proposed a coupled hidden Markov model (CHMM) approach to video-realistic speech animation, which realizes realistic facial animations driven by speaker-independent continuous speech.

Following the modern trends in machine learning, deep learning approaches gained the most promising results in the audio-based video synthesis domain. In [41] the authors propose to use a deep learning model that learns arbitrary nonlinear mappings from phoneme label input sequences to mouth movements in a way that accurately captures natural motion and visual coarticulation effects. In [19] a convolutional network is used to transform audio features to 3D meshes of a specific person. The authors of [35] presented unsupervised keypoint detection and warping for a motion transfer. Several approaches explore the variations of recurrent models [8, 25, 40, 51].

The most up-to-date approaches for speech-driven video synthesis are based on generative models. Variations of Generative Adversarial Networks (GANs) [9] were primarily applied to the problem of video generation [26, 29, 42]. The GAN-based approach for a speech-driven generation was introduced in [45]. The authors propose an end-to-end system that generates videos of a talking head, using only a still image of a person and an audio clip containing speech without relying on handcrafted intermediate features. They achieve it by utilizing temporal GAN that uses three discriminators focused on achieving detailed frames, audio-visual synchronization, and realistic expressions. In [27], the authors propose to incorporate an additional pre-trained Lip-Sync expert during the training to maintain the consistency of generated videos. The paper [3] introduces a 3D-aware generative network along with a hybrid embedding module that assures rhythmic head motion. The model presented in [50] modularizes audio-visual representations by devising an implicit low-dimension pose code to tackle the problem of rhythmic head motion. In StyleHEAT [49], the authors show how to utilize StyleGAN [20] model to create talking faces guided by speech embeddings but also con-

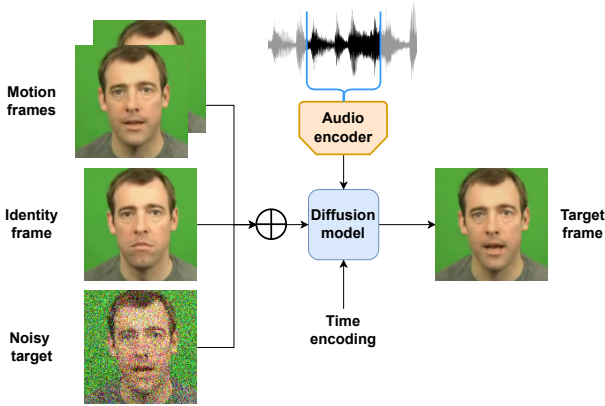


Figure 2. Training step of Diffused Heads. Our model learns to denoise one frame at a time, using identity and motion frames, and an audio embedding extracted from a pre-trained audio encoder. The identity frame informs the model what the face of interest is, and the motion frames are utilized to preserve the movement.

trolled by intuitive or attribute editing.

Some modern approaches utilize rendering networks to obtain more accurate face 3D representation. In [39] the authors introduce a novel video rendering network and a dynamic programming method to construct a temporally coherent and photo-realistic video. The authors of [10] propose to use an audio-conditioned implicit function to generate a dynamic neural radiance field, from which a high-fidelity talking-head video corresponding to the audio signal is synthesized using volume rendering. A Portrait Image Neural Renderer (PIRenderer) is introduced in [32] that controls the face motions with the parameters of a three-dimensional morphable face model. In [18], Implicit Emotion Displacement Learner, together with Dense Warping Field, are used to obtain high-quality images.

The Denoising Diffusion Probabilistic Models [15] are gaining popularity and often outperforming GANs in tasks like image synthesis [6], and other guided image generation tasks [22, 30, 31, 33]. Several attempts utilize that group of models in video generation [11, 14, 16, 17, 37].

To the best of our knowledge, there are no direct attempts to solve speech-driven video synthesis problems using diffusion models. Moreover, our method is the first one-shot approach that can hallucinate diverse head motions and does not require an actor to drive the movement via additional visual guidance input. The realism of the gestures is on par with or superior to that of facial reenactment methods.

3. Diffusion models

Let us assume we are given samples x_0 from a data distribution $x_0 \sim q(x_0)$. We can define a *forward process*

$q(x_{1:T}|x_0) := \prod_{t=1}^T q(x_t|x_{t-1})$ that gradually adds Gaussian noise to the data:

$$q(x_t|x_{t-1}) := \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t\mathbf{I}) \quad (1)$$

where T defines a number of diffusion steps, and $\{\beta_t\}_{t=1}^T$ is a noise schedule starting from low values that increase with t . Note that $\{\beta_t\}_{t=1}^T$ is known beforehand, so the entire forward process is fixed. For sufficiently large T and $\beta_t \in (0, 1)$, x_T becomes close to a sample drawn from isotropic Gaussian distribution, *i.e.* $x_T \sim \mathcal{N}(0, \mathbf{I})$.

An interesting property of the forward process is the fact that one can access its intermediate states in only one step, that is:

$$q(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)\mathbf{I}) \quad (2)$$

where $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$, and $\alpha_t = 1 - \beta_t$. It allows us to train the model more efficiently.

We are interested in learning how to denoise a sample drawn from Gaussian distribution back to the data. Note that $q(x_{t-1}|x_t)$ depends on the entire dataset and hence is intractable. Intuitively, learning how to denoise x_T into the underlying data point is only possible when the model is given explicit information about where the forward diffusion process started for it. Thus, we can additionally condition $q(x_{t-1}|x_t)$ on x_0 , making it tractable. Using Bayes' theorem we get:

$$q(x_{t-1}|x_t, x_0) = \mathcal{N}(x_{t-1}; \tilde{\mu}(x_t, x_0), \tilde{\beta}_t\mathbf{I}) \quad (3)$$

where:

$$\tilde{\beta}_t := \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t \quad (4)$$

$$\tilde{\mu}(x_t, x_0) := \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1 - \bar{\alpha}_t}x_0 + \frac{\sqrt{\bar{\alpha}_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t}x_t \quad (5)$$

Similarly to VAE framework, we define variational posteriors that approximate $q(x_{t-1}|x_t, x_0)$:

$$p_\theta(x_{t-1}|x_t) := \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)) \quad (6)$$

As in [23], $\mu_\theta(x_t, t)$ and $\Sigma_\theta(x_t, t)$ are further reparameterized into:

$$\mu_\theta(x_t, t) = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(x_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right) \quad (7)$$

$$\Sigma_\theta(x_t, t) = \exp(\nu \log \beta_t + (1 - \nu) \log \tilde{\beta}_t) \quad (8)$$

where $\epsilon_\theta(x_t, t)$ is the model's prediction of the Gaussian noise ϵ applied on x_0 in the process of getting x_t , and ν is an additional output of the model. Nichol & Dhariwal in [23] proposed to train μ_θ and Σ_θ separately, using L_{simple} and L_{vlb} respectively, where:

$$L_{simple} := \mathbb{E}_{t, x_0, \epsilon} [\|\epsilon - \epsilon_\theta(x_t, t)\|^2] \quad (9)$$

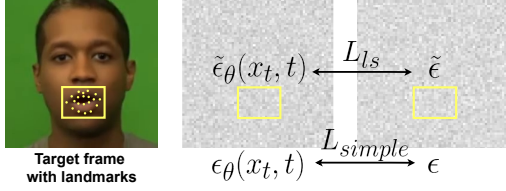


Figure 3. In addition to minimizing L2 distance between ground truth noise ϵ and predicted noise $\epsilon_\theta(x_t, t)$ in L_{simple} , we utilize the target frame’s landmarks to minimize lip sync loss L_{ls} between cropped ground truth noise $\tilde{\epsilon}$ and corresponding predicted area $\tilde{\epsilon}_\theta(x_t, t)$.

and L_{vlb} is the variational lower bound (VLB) defined as:

$$L_{vlb} := L_0 + L_1 + \dots + L_{T-1} + L_T \quad (10)$$

where:

$$L_0 := -\log p_\theta(x_0|x_1) \quad (11)$$

$$L_t := D_{KL}(q(x_t|x_{t+1}, x_0) \parallel p_\theta(x_t|x_{t+1})) \quad (12)$$

for $t \in \{1, \dots, T-1\}$

$$L_T := D_{KL}(q(x_T|x_0) \parallel p_\theta(x_T)) \quad (13)$$

For images, L_0 is a discretized Gaussian distribution as proposed in [15]. L_T is omitted because q has no trainable parameters and $p_\theta(x_T)$ is a Gaussian prior. All of the other terms are Kullback–Leibler divergences between two Gaussian distributions that can be written in a closed form.

In practice, a 2D UNet [34] with skip-connections, and attention layers is used as a backbone to predict both noise $\epsilon_\theta(x_t, t)$ and variance $\Sigma_\theta(x_t, t)$. Information about timestep t is injected using corresponding time embedding $\psi(t)$ and group normalization (GN):

$$h_{s+1} = t_s \text{GN}(h_s) + t_b \quad (14)$$

where h_s and h_{s+1} are consecutive hidden states of UNet, and $(t_s, t_b) = \text{MLP}(\psi(t))$, where MLP is a shallow neural network consisting of linear layers.

4. Method

Diffused Heads generates one frame at a time given an identity frame that stays fixed during the entire generation process, and a speech recording embedded using a pre-trained audio encoder. To achieve smoother and more expressive results, we inject additional information on past movement and future expressions by motion frames (Section 4.2) and audio embeddings (Section 4.3). Moreover, an additional lip sync loss (Section 4.4) is defined to force the model to pay more attention to the mouth region.

4.1. Training

We train a diffusion model to learn the distribution of frames extracted from videos. The training process is

shown in Figure 2. We randomly sample a video $X = \{x^{(1)}, \dots, x^{(n)}\}$ from the training set, and then a frame $x^{(k)}$ from X . n is the total number of frames. In addition to the standard diffusion model’s inputs, *i.e.* a time step t and the frame with added noise $x_t^{(k)}$ (following Equation (2)), to keep the actor’s identity, we concatenate $x_t^{(k)}$ with an *identity frame* x_{Id} channel-wise:

$$x_{In,t}^{(k)} := x_t^{(k)} \oplus_c x_{Id} \quad (15)$$

x_{Id} is randomly chosen from X . Selecting the identity frame randomly instead of $x^{(0)}$ during the training makes the model familiar to a larger variety of frames as input. In consequence, the generation’s robustness is improved.

To add temporal information, we split the corresponding audio sequence into chunks of equal length based on the number of frames in the video. Then, using an audio encoder from [45] pre-trained on the LRW [4] dataset, the audio chunks are encoded into audio embeddings $Y = \{y^{(1)}, \dots, y^{(n)}\}$. The details of our proposed audio conditioning method can be found in Section 4.3.

4.2. Motion frames

Even though temporal information is provided to the model by the audio encoder, it is not enough to generate smooth videos. To overcome this problem and preserve the motion, for the target frame $x^{(k)}$ we introduce *motion frames* $x_{Motion}^{(k)} = \bigoplus_c (\{x^{(k-m_x)}, \dots, x^{(k-1)}\})$, where m_x is the number of motion frames, and $\bigoplus_c(\cdot)$ is concatenation operation in channel dimension on all of its arguments. During our ablation study (Section 5.4), we found that the best value for m_x is 2.

If there are not enough frames preceding $x^{(k)}$, the most natural choice is to fill the remaining motion frames with duplicates of $x^{(0)}$. However, during sampling, we have no access to any ground truth frames, except the identity one. We also do not necessarily want the generated video to start with an exact facial expression as given in the identity frame, *e.g.* when the audio recording starts with silence and the person in the identity frame has their mouth open. Thus, to make the model robust on sample initialization, we utilize x_{Id} as a substitute for missing motion frames.

The motion frames are added to Equation (15) and get the final form of the direct input to the model:

$$x_{In,t}^{(k)} := x_t^{(k)} \oplus_c x_{Id} \oplus_c x_{Motion}^{(k)} \quad (16)$$

4.3. Speech conditioning

We propose to inject information from audio embedding $y^{(k)}$ by modifying Equation (14) into:

$$h_{s+1} = y_s^{(k)}(t_s \text{GN}(h_s) + t_b) + y_b^{(k)} \quad (17)$$




Figure 4. Comparison with other methods on LRW [4] (left) and CREMA [2] (right) datasets.

where $(y_s^{(k)}, y_b^{(k)}) = \text{MLP}(y^{(k)})$. In this setting, we shift and scale hidden states of the UNet with the information not only from the time encoding but the audio embedding as well. We found this approach works better in comparison to other conditioning methods, such as using just an additional scale on top of Equation (14) [28], and applying a multi-head attention mechanism with queries being a function of the audio embedding [33].

In contrast to motion frames which during sampling are only available for already processed frames, we have an access to the entire speech recording beforehand. To make use of it, we introduce *motion audio embeddings* that bring information from both past and future audio segments. We define them as a vector created by concatenating selected audio embeddings: $y_{Motion}^{(k)} = \bigoplus(\{y^{(k-m_y)}, \dots, y^{(k)}, \dots, y^{(k+m_y)}\})$, where m_y is the number of additional audio embeddings from one side. The details of our choice of m_y can be found in the ablation study in Section 5.4. Similarly to motion frames, if we run out of embeddings, we pad $y_{Motion}^{(k)}$ with either $y^{(0)}$ at the beginning or $y^{(n)}$ at the end. Finally, we use $y_{Motion}^{(k)}$ instead of $y^{(k)}$ in Equation (17).

4.4. Lip sync loss

Unlike other methods [3, 18, 27, 32, 39, 45, 50], we do not use any explicit loss function to promote better lip sync of generated samples. Solutions that rely on using dedicated perceptual losses based on pre-trained lipreading models have been effective in improving lip motion accu-

racy [27, 44]. However, Diffused Heads works on frames, not sequences so sequence-based losses can not be applied, and more importantly, during diffusion model training, the goal is to predict the noise that was used on the target frame. Getting back from predicted noise to initial x_0 , which is required to apply the perceptual loss, is not accurate enough in a single step, and computationally inefficient in more steps.

We introduce a simpler solution: an additional *lip sync loss* L_{ls} . During the training, we leverage facial landmarks to crop each frame around the mouth area, and minimize noise prediction in this region:

$$L_{ls} := \mathbb{E}_{t, x_0, \epsilon} [||\tilde{\epsilon} - \tilde{\epsilon}_\theta(x_t, t)||^2] \quad (18)$$

where $\tilde{\epsilon}$ and $\tilde{\epsilon}_\theta$ indicate cropped versions of ground truth and predicted noise, respectively. The process is visualized in Figure 3. With the lip loss, the model pays more attention to lip synchronization with audio embeddings, improving the overall perception of sampled videos. We weight L_{ls} with a constant λ_{ls} that leverages the model’s attention to details of a mouth region and the rest of the frame. We discuss the choice of λ_{ls} in Section 5.4.

The final optimization objective becomes:

$$L_{simple} + \lambda_{vfb} L_{vfb} + \lambda_{ls} L_{ls} \quad (19)$$

where L_{simple} and L_{vfb} are defined by Equations (9) and (10), respectively.

4.5. Sampling

For sampling, only an identity frame and audio embeddings extracted from a speech recording are required. We

		FVD ↓	FID ↓	Blinks/s	Blink dur.	OFM	F-MSE	AV off.	AV Conf. ↑	WER ↓
LRW	SDA	198.84	61.95	0.52	0.28	73.82	18.94	1	7.40	0.77
	MakeItTalk	269.29	7.57	0.09	0.28	57.21	3.44	-3	3.16	0.99
	Wav2Lip*	366.14	2.83	0.03	0.16	47.12	1.45	-2	6.58	0.51
	PC-AVS	153.12	11.96	0.20	0.16	69.59	17.13	-3	6.24	0.64
	EAMM	172.18	9.28	0.03	0.16	58.46	4.39	-3	3.83	0.95
	Ours	71.88	3.94	0.35	0.28	70.71	19.69	-2	4.61	0.77
	GT			0.53	0.28	72.02	27.34	-2	5.83	
CREMA	SDA	376.48	79.82	0.25	0.26	68.21	6.83	2	5.50	-
	MakeItTalk	256.88	17.26	0.02	0.80	62.36	2.07	-3	3.75	-
	Wav2Lip*	193.32	12.57	0	-	46.87	1.07	-2	6.68	-
	PC-AVS	333.94	22.53	0.02	0.20	70.36	6.93	-3	6.17	-
	EAMM	196.82	19.40	0	-	58.91	1.65	-2	4.26	-
	Ours	88.61	12.45	0.28	0.36	64.30	6.99	1	4.52	-
	GT			0.24	0.40	68.76	7.76	1	5.14	

Table 1. Comparison with other methods. The best scores are in dark green and bold, second bests are in light green. ↑ / ↓ indicate higher/lower is better, respectively. Lack of arrow indicates the closer to GT the better. *All of the other methods are one-shot. For fair comparison, Wav2Lip videos were generated using still images, *i.e.* only mouth regions change.

start video generation by initializing $x_{Motion}^{(0)}$ with copies of the identity frame. Each frame is sampled following the denoising process of diffusion models defined by variational posterior in Equation (6). After every step, we replace the latest motion frame with a synthesized one. $y_{Motion}^{(k)}$ follows the same procedure as during the training.

Generation of a single frame takes a significant amount of time since it requires the model to make a prediction for all of the diffusion time steps $\{1, \dots, T\}$. To speed up the process, methods like DDIM [38] or time step resampling can be used. In this work, we use the latter reducing sampling time by a factor of 5.

During experiments, we observed that our model sometimes failed when generating sudden head movements. It synthesizes sequences frame-by-frame, and any occurring errors accumulate in later steps. One of the associated problems is that during training all of the motion frames come from the dataset. Meanwhile, during generation, we use previously sampled frames that have some distortions. We hypothesize that with this setting, the motion frames and the identity frame are equally important in terms of extracting a person’s attributes.

To force the model to take more information on the person’s appearance from the identity frame, we convert each motion frame to a grayscale. The intuition behind this is that it should make it harder for the model to extract identity features (such as color) while pushing it to seek motion information instead. We found this solution to work well on more complex datasets with a big number of participants.

5. Experiments

We evaluate Diffused Heads on the most commonly used datasets for talking face generation: CREMA [2] and LRW [4]. We compare our method qualitatively and quantitatively to the current state-of-the-art in guided [18,27,50,51] and pose guidance-free [45] video synthesis. To experience the full quality of our results, readers are strongly encouraged to watch generated videos in the supplementary materials. We will release our code for public use.

5.1. Implementation Details

Our model is trained on 128x128 resolution videos. We use the same UNet [34] architecture as proposed in [6], with audio conditioning explained in Section 4.3. We use 256-512-768 channels for the input blocks with 2 ResNet [12] layers each. In the early stages of experiments, we found adding more attention layers worsened the quality of generated frames. Thus, we only use an attention layer with 4 heads and 64 head channels in the middle block.

5.2. Qualitative results

We present qualitative comparison with other methods on CREMA and LRW in Figure 4. Videos can be found in the supplementary materials. Diffused Heads generates videos that are hard to distinguish from real ones. The faces have natural expressions, eye blinks, and grimaces. The model is able to preserve smooth motion between frames and identity from a given input frame. There are hardly any artifacts, and difficult objects such as hair or glasses are generated accurately. Additionally, Diffused Head works well on challenging videos with people shown from a side view. The important thing to note is that our model does not suffer

Method	Score
PC-AVS [50]	34.95%
Diffused Heads (ours)	68.72%
Real videos	64.00%

Table 2. Turing test on LRW [4] dataset. 10 videos per method and 10 real ones (30 in total) were shown to 140 people. They were asked to vote on whether the samples were real or not. The scores indicate how authentic the videos seemed to the participants.

from mode collapse, and to prove our claim, we will share the entire generated test set on the project’s website.

5.3. Quantitative results

We compare Diffused Heads to other methods: SDA [45], Wav2Lip [27], MakeItTalk [51], PC-AVS [50], and EAMM [18]. Results can be found in Table 1. We use first frames and audio sequences from test splits of CREMA and LRW datasets to generate clips. For a fair comparison, the driving videos for PC-AVS and EAMM are chosen randomly from the test sets, and the audio recordings are the same for all of the models.

To measure quality of generated videos, we calculate Fréchet Inception Distance (FID) [13] and Fréchet Video Distance (FVD) [43]. FVD shares the same principles as FID but is extended to video domain, penalizing unnatural dynamics of consecutive frames.

To overcome one of the main challenges of video evaluation, we propose to use the average Optical Flow Magnitude (OFM) of consecutive frames, and Frame-wise Mean Square Error (F-MSE) to measure sequence smoothness (see Appendix for definitions). However, values close to zero are not desirable, since a single repeating image would get a perfect score. Hence, we treat them as population-reference metrics, *i.e.* we want them to be as close as possible to their corresponding ground truth values. Similarly, to evaluate expressiveness, we report population-reference metrics obtained using a blink detector: the average number of blinks per second and the median blink duration.

Lip (WER) and audio-visual (AV Offset, AV Confidence) syncs are measured using a pre-trained lipreader and Syncnet [5], respectively.

Unlike other approaches that mimic head movements from a driving video, our model hallucinates the head motion. As a result, metrics that measure deviation from a ground truth sequence (*e.g.* PSNR, SSIM) heavily penalize our method and are not suitable for evaluation.

Clips generated by our model achieve state-of-the-art scores in FVD and best or close-to-best scores in FID, blinks per second, blink duration, smoothness (OFM and F-MSE), and AV Offset. These metrics show our samples being the best looking and on par with head movements and facial expressions presented in the real videos while main-



Figure 5. Average magnitudes of optical flow and consecutive frames for 0 (top) and 2 (bottom) motion frames.

taining smooth motion and good AV sync. Worse performance in terms of WER is a consequence of the lack of an expert method supervising lip sync during training. For the same reason, lower values of AV Confidence are observed. However, they do not differ a lot from the ground truth and are still comparable to the best methods. It is also worth noting that SDA generates heavily cropped faces without any head movements. Additionally, high FID values of Wav2Lip are the effect of animating only mouth region. The remaining part is a fixed ground truth image which is rewarded by FID. This is reflected in worse FVD scores that capture unrealistic motion - in this case a lack of movement in the majority of a video. Diffused Heads produces much more realistic and vivid samples. Overall, our model achieves state-of-the-art on most of the metrics.

Moreover, Diffused Heads wins in the most important metric in visual data synthesis - human perception (see Table 2). To prove it, we conducted a Turing test on 140 participants. We picked 10 test videos from the LRW dataset generated by the current state-of-the-art method PC-AVS, 10 from our model, and 10 real ones. We sent them to both females and males from different backgrounds and countries. Each of them, after watching every one of the 30 videos, was asked to vote whether they were real or not. Diffused Heads performed much better than PC-AVS, and also achieved higher scores than real videos. LRW videos contain jitter due to the landmark detection and cropping procedure that was used by the authors of the original paper. This high-frequency noise is not captured by our model leading to smoother sequences (OFM and F-MSE metrics in Table 1) that may seem more natural.

5.4. Ablation study

We investigated the influence of a number of motion frames on video quality. We noticed that not using any led to almost random facial expressions. To maintain the motion, we experimented with up to 3 motion frames. In Figure 5, we show the comparison between generated videos without any blinks for 0 and 2 motion frames. The average magnitude of an optical flow of a video generated without any motion frames is much higher than the one for 2 motion frames. It is also uniformly high in all of the face regions. It indicates more random movements between consecutive

Motion audio embeddings	Motion grayscale	CPBD \uparrow	MSE \downarrow	SSIM \uparrow	PSNR \uparrow	LMD \downarrow	WER \downarrow
0	\times	0.0857	1194	0.5767	18.2912	3.0282	0.93
	\checkmark	0.0858	1148	0.5890	18.5003	2.8471	0.93
1	\times	0.0872	1228	0.5782	18.2309	2.8786	0.80
	\checkmark	0.0856	1131	0.5996	18.6589	2.6705	0.81
2	\times	0.0831	1275	0.5658	17.9912	3.0214	0.72
	\checkmark	0.0925	1025	0.6225	19.1072	2.5297	0.77
3	\times	0.0882	1266	0.5678	17.9945	2.9253	0.74
	\checkmark	0.0882	1184	0.5851	18.4350	2.7590	0.75

Table 3. Ablation study on LRW [4] dataset. \uparrow/\downarrow indicate higher/lower is better, respectively.



Figure 6. Results of generalization. The audio recordings were (from top): Korean female, German male, and English male. The first row was generated with a model trained on CREMA [2], and the last two with LRW [4] one. The first two rows used audio from AVSpeech [7], and for the last one, we used a custom image and recording.

frames. For 2 motion frames, we can spot the highest density around the mouth region, which is the desired behavior. We include the videos for 0, 1, and 2 motion frames in the supplementary materials where the difference is clearly visible. The model failed to work with 3 motion frames.

For the lip sync loss weight λ_{ls} , we observed values greater than 0.5 degraded the quality of results. The value of 0.2 gave very realistic results and the best WER score.

Finally, the number of motion audio embeddings and whether to use grayscale on motion frames turned out to be crucial. We present the numerical results of the ablation study on the LRW dataset in Table 3. Utilizing grayscale improves the quality of generated videos for every choice of the number of motion audio embeddings. For the latter, the best value to use was 2. We noticed that using grayscale on motion frames does not help in less diverse datasets, such as CREMA. It consists of videos of only 91 actors, making generalization to new faces much harder. For that reason, using RGB motion frames lets the model take more information on identity from both identity and motion frames.

5.5. Generalization

One of the main challenges in deep learning is the ability of models to generalize well to unseen data. We conducted experiments to show the robustness of Diffused Heads in this manner, and the results can be found in Figure 6. We

show that the model performs well when given part or even all of the input from a different source.

We investigated the behavior of our model with identity frames from CREMA and LRW, and audio recordings from AVSpeech [7]: a female speaking Korean, and a male German speaker. As the final test, we generated a video of a talking avatar, using an image synthesized by DALL-E 2 [31] and our own recorded speech. We include more examples in the supplementary video.

The results show Diffused Heads works well on data that comes from outside of a training distribution. The generated frames look pleasant, and the lip movement and facial expressions look natural. Surprisingly, our model was even able to successfully process the image of the avatar, even though it was given only human faces during the training.

5.6. Limitations

Despite Diffused Heads achieving state-of-the-art results, it still suffers from some limitations. The main challenge of our method is the length of generated videos. Since we do not provide any additional pose input or visual guidance for head movement and the model autoregressively generates frames, it fails to keep the initial quality for sequences longer than 8-9 seconds. Additionally, diffusion models suffer from long generation times in comparison to other generative models. For now, it is not possible to use

our approach in real-time applications, even though it is theoretically suitable for them. New metrics used in talking face generation task are also an open research problem.

6. Conclusions

In this work, we presented Diffused Heads: a frame-based method for talking face generation. To synthesize a video that is hard to distinguish by a human from a real one, it only needs one identity frame and an audio sequence containing speech. We evaluated our approach on 2 datasets with different levels of complexity, achieving state-of-the-art results on both of them. We supported this statement by conducting a Turing test on 140 participants showing our results to be indistinguishable from ground truth videos.

References

- [1] Martin Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks. *arXiv preprint arXiv:1701.04862*, 2017. [2](#)
- [2] Houwei Cao, David G Cooper, Michael K Keutmann, Ruben C Gur, Ani Nenkova, and Ragini Verma. Crema-d: Crowd-sourced emotional multimodal actors dataset. *IEEE transactions on affective computing*, 5(4):377–390, 2014. [5](#), [6](#), [8](#)
- [3] Lele Chen, Guofeng Cui, Celong Liu, Zhong Li, Ziyi Kou, Yi Xu, and Chenliang Xu. Talking-head generation with rhythmic head motion. In *European Conference on Computer Vision*, pages 35–51. Springer, 2020. [2](#), [5](#)
- [4] Joon Son Chung and Andrew Zisserman. Lip reading in the wild. In *A'sian conference on computer vision*, pages 87–103. Springer, 2016. [4](#), [5](#), [6](#), [7](#), [8](#)
- [5] J. S. Chung and A. Zisserman. Out of time: automated lip sync in the wild. In *Workshop on Multi-view Lip-reading, ACCV*, 2016. [7](#)
- [6] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021. [3](#), [6](#)
- [7] Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T Freeman, and Michael Rubinstein. Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation. *arXiv preprint arXiv:1804.03619*, 2018. [8](#)
- [8] Bo Fan, Lijuan Wang, Frank K Soong, and Lei Xie. Photoreal talking head with deep bidirectional lstm. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4884–4888. IEEE, 2015. [2](#)
- [9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. [2](#)
- [10] Yudong Guo, Keyu Chen, Sen Liang, Yong-Jin Liu, Hujun Bao, and Juyong Zhang. Ad-nerf: Audio driven neural radiance fields for talking head synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5784–5794, 2021. [3](#)
- [11] William Harvey, Saeid Naderiparizi, Vaden Masrani, Christian Weilbach, and Frank Wood. Flexible diffusion modeling of long videos. *arXiv preprint arXiv:2205.11495*, 2022. [3](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *Computer Science*, 2015. [6](#)
- [13] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017. [7](#)
- [14] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. [3](#)
- [15] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. [2](#), [3](#), [4](#)
- [16] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. *arXiv preprint arXiv:2204.03458*, 2022. [3](#)
- [17] Tobias Höppe, Arash Mehrjou, Stefan Bauer, Didrik Nielsen, and Andrea Dittadi. Diffusion models for video prediction and infilling. *arXiv preprint arXiv:2206.07696*, 2022. [3](#)
- [18] Xinya Ji, Hang Zhou, Kaisiyuan Wang, Qianyi Wu, Wayne Wu, Feng Xu, and Xun Cao. Eamm: One-shot emotional talking face via audio-based emotion-aware motion model. *arXiv preprint arXiv:2205.15278*, 2022. [2](#), [3](#), [5](#), [6](#), [7](#)
- [19] Tero Karras, Timo Aila, Samuli Laine, Antti Herva, and Jaakko Lehtinen. Audio-driven facial animation by joint end-to-end learning of pose and emotion. *ACM Transactions on Graphics (TOG)*, 36(4):1–12, 2017. [2](#)
- [20] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. [2](#)
- [21] Avisek Lahiri, Vivek Kwatra, Christian Frueh, John Lewis, and Chris Bregler. Lipsync3d: Data-efficient learning of personalized 3d talking faces from video using pose and lighting normalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2755–2764, 2021. [2](#)
- [22] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021. [3](#)
- [23] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. [2](#), [3](#)
- [24] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2085–2094, October 2021. [2](#)

- [25] Hai X Pham, Samuel Cheung, and Vladimir Pavlovic. Speech-driven 3d facial animation with implicit emotional awareness: a deep learning approach. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 80–88, 2017. 2
- [26] Hai X Pham, Yuting Wang, and Vladimir Pavlovic. Generative adversarial talking head: Bringing portraits to life with a weakly supervised neural network. *arXiv preprint arXiv:1803.07716*, 2018. 2
- [27] KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Namboodiri, and CV Jawahar. 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*, pages 484–492, 2020. 2, 5, 6, 7
- [28] Konpat Preechakul, Nattanat Chatthee, Suttisak Widadwongsa, and Supasorn Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10619–10629, 2022. 5
- [29] Albert Pumarola, Antonio Agudo, Aleix M Martinez, Alberto Sanfeliu, and Francesc Moreno-Noguer. Ganimation: Anatomically-aware facial animation from a single image. In *Proceedings of the European conference on computer vision (ECCV)*, pages 818–833, 2018. 2
- [30] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 3
- [31] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 3, 8
- [32] Yurui Ren, Ge Li, Yuanqi Chen, Thomas H Li, and Shan Liu. Pirenderer: Controllable portrait image generation via semantic neural rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13759–13768, 2021. 2, 3, 5
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 3, 5
- [34] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 4, 6
- [35] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. In *Conference on Neural Information Processing Systems (NeurIPS)*, December 2019. 2
- [36] AD Simons. Generation of mouthshape for a synthetic talking head. *Proc. of the Institute of Acoustics*, 1990. 2
- [37] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oran Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022. 3
- [38] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 6
- [39] Linsen Song, Wayne Wu, Chen Qian, Ran He, and Chen Change Loy. Everybody’s talkin’: Let me talk as you want. *IEEE Transactions on Information Forensics and Security*, 17:585–598, 2022. 2, 3, 5
- [40] Supasorn Suwajanakorn, Steven M Seitz, and Ira Kemelmacher-Shlizerman. Synthesizing obama: learning lip sync from audio. *ACM Transactions on Graphics (TOG)*, 36(4):1–13, 2017. 2
- [41] Sarah Taylor, Taehwan Kim, Yisong Yue, Moshe Mahler, James Krahe, Anastasio Garcia Rodriguez, Jessica Hodgins, and Iain Matthews. A deep learning approach for generalized speech animation. *ACM Transactions on Graphics (TOG)*, 36(4):1–11, 2017. 2
- [42] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018. 2
- [43] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphaël Marinier, Marcin Michalski, and Sylvain Gelly. Fvd: A new metric for video generation. 2019. 7
- [44] Konstantinos Vougioukas. *Generation of realistic human behaviour*. PhD thesis, Imperial College London, 2022. 5
- [45] Konstantinos Vougioukas, Stavros Petridis, and Maja Pantic. Realistic speech-driven facial animation with gans. *International Journal of Computer Vision*, 128(5):1398–1413, 2020. 2, 4, 5, 6, 7
- [46] Lei Xie and Zhi-Qiang Liu. A coupled hmm approach to video-realistic speech animation. *Pattern Recognition*, 40(8):2325–2340, 2007. 2
- [47] Eli Yamamoto, Satoshi Nakamura, and Kiyohiro Shikano. Lip movement synthesis from speech based on hidden markov models. *Speech Communication*, 26(1-2):105–115, 1998. 2
- [48] Hani Yehia, Philip Rubin, and Eric Vatikiotis-Bateson. Quantitative association of vocal-tract and facial behavior. *Speech Communication*, 26(1-2):23–43, 1998. 2
- [49] Fei Yin, Yong Zhang, Xiaodong Cun, Mingdeng Cao, Yanbo Fan, Xuan Wang, Qingyan Bai, Baoyuan Wu, Jue Wang, and Yujiu Yang. Styleheat: One-shot high-resolution editable talking face generation via pretrained stylegan. *arXiv preprint arXiv:2203.04036*, 2022. 2
- [50] Hang Zhou, Yasheng Sun, Wayne Wu, Chen Change Loy, Xiaogang Wang, and Ziwei Liu. Pose-controllable talking face generation by implicitly modularized audio-visual representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4176–4186, 2021. 2, 5, 6, 7
- [51] Yang Zhou, Xintong Han, Eli Shechtman, Jose Echevarria, Evangelos Kalogerakis, and Dingzeyu Li. Makeltalk: speaker-aware talking-head animation. *ACM Transactions on Graphics (TOG)*, 39(6):1–15, 2020. 2, 6, 7

A. Temporal metrics

We introduce two additional metrics to evaluate the smoothness of generated videos: frame-wise Optical Flow Magnitude (OFM) and Frame-wise Mean Square Error (F-MSE).

Let $X = \{x^{(1)}, \dots, x^{(n)}\}$ be a sequence of frames, and $x_{(i,j)}^{(k)}$ an individual pixel of the k -th frame of size $W \times H$. We can define OFM(X) and F-MSE(X) as:

$$\text{OFM}(X) = \frac{1}{WH(n-1)} \sum_{k=2}^n \sum_{i=1}^W \sum_{j=1}^H F(x_{(i,j)}^{(k-1)}, x_{(i,j)}^{(k)}) \quad (20)$$

$$\text{F-MSE}(X) = \frac{1}{WH(n-1)} \sum_{k=2}^n \sum_{i=1}^W \sum_{j=1}^H \|x_{(i,j)}^{(k-1)} - x_{(i,j)}^{(k)}\|_2^2 \quad (21)$$

where F is an optical flow magnitude between consecutive frames. We use OpenCV's `calcOpticalFlowFarneback` with parameters: `pyr_scale=0.5`, `levels=7`, `winsize=5`, `iterations=15`, `poly_n=5`, and `poly_sigma=1.2`.

To simplify the notation, we assumed a single channel per pixel. For more channels, the metrics are defined analogically, with additional averaging over channel dimension. Finally, the reported values are averages over all of the videos in the test set.

The proposed metrics are population-based. We want their values calculated on generated videos to be as close as possible to the ones on ground truth sequences. It is worth noting that a sequence containing just a single repeating frame would get a value of 0 in both of the metrics making them impractical.