



HAL
open science

4D Facial Expression Diffusion Model

Kaifeng Zou, Sylvain Faisan, Boyang Yu, Sebastien Valette, Hyewon Seo

► **To cite this version:**

Kaifeng Zou, Sylvain Faisan, Boyang Yu, Sebastien Valette, Hyewon Seo. 4D Facial Expression Diffusion Model. ACM Transactions on Multimedia Computing, Communications and Applications, 2023, 10.1145/3653455 . hal-04301288

HAL Id: hal-04301288

<https://hal.science/hal-04301288>

Submitted on 17 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

4D Facial Expression Diffusion Model

KAIFENG ZOU, ICube Laboratory, University of Strasbourg, France

SYLVAIN FAISAN, ICube Laboratory, University of Strasbourg, France

BOYANG YU, ICube Laboratory, University of Strasbourg, France

SÉBASTIEN VALETTE, CREATIS, CNRS, INSA-Lyon, Lyon, France, France

HYEWON SEO, ICube Laboratory, University of Strasbourg, France

Facial expression generation is one of the most challenging and long-sought aspects of character animation, with many interesting applications. The challenging task, traditionally having relied heavily on digital craftspersons, remains yet to be explored. In this paper, we introduce a generative framework for generating 3D facial expression sequences (i.e. 4D faces) that can be conditioned on different inputs to animate an arbitrary 3D face mesh. It is composed of two tasks: (1) Learning the generative model that is trained over a set of 3D landmark sequences, and (2) Generating 3D mesh sequences of an input facial mesh driven by the generated landmark sequences. The generative model is based on a Denoising Diffusion Probabilistic Model (DDPM), which has achieved remarkable success in generative tasks of other domains. While it can be trained unconditionally, its reverse process can still be conditioned by various condition signals. This allows us to efficiently develop several downstream tasks involving various conditional generation, by using expression labels, text, partial sequences, or simply a facial geometry. To obtain the full mesh deformation, we then develop a landmark-guided encoder-decoder to apply the geometrical deformation embedded in landmarks on a given facial mesh. Experiments show that our model has learned to generate realistic, quality expressions solely from the dataset of relatively small size, improving over the state-of-the-art methods. Videos and qualitative comparisons with other methods can be found at <https://github.com/ZOUKaifeng/4DFM>. Code and models will be made available upon acceptance.

CCS Concepts: • **Computing methodologies** → Machine learning; Computer graphics.

Additional Key Words and Phrases: diffusion model, neural networks, facial expression, generative model

ACM Reference Format:

Kaifeng Zou, Sylvain Faisan, Boyang Yu, Sébastien Valette, and Hyewon Seo. 2023. 4D Facial Expression Diffusion Model. *J. ACM* 37, 4, Article 111 (August 2023), 24 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

3D facial expression synthesis is a fundamental, long-sought problem in face animation and recognition. Due to the inherent subtlety and sophistication of facial expressions, as well as our sensitivity to them, the task is extremely challenging yet highly beneficial to various multimedia applications, such as creating virtual humans in games and films, developing virtual avatars in immersive virtual and augmented reality experiences, implementing chatbots in digital marketing,

Authors' addresses: Kaifeng Zou, ICube Laboratory, University of Strasbourg, 300 bd Sébastien Brant, Illkirch, France, 67402, kaifeng.zou@unistra.fr; Sylvain Faisan, faisan@unistra.fr, ICube Laboratory, University of Strasbourg, 300 bd Sébastien Brant, Illkirch, France, 67402; Boyang Yu, ICube Laboratory, University of Strasbourg, 1 Pl. de l'Hôpital, Strasbourg, France, b.yu@unistra.fr; Sébastien Valette, CREATIS, CNRS, INSA-Lyon, Lyon, France, 21 Av. Jean Capelle O, Villeurbanne, France, sebastien.valette@creatis.insa-lyon.fr; Hyewon Seo, ICube Laboratory, University of Strasbourg, 1 Pl. de l'Hôpital, Strasbourg, France, seo@unistra.fr.

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 ACM 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.

© 2023 Association for Computing Machinery.

0004-5411/2023/8-ART111 \$15.00

<https://doi.org/XXXXXXX.XXXXXXX>

and recognizing, managing, or tracking of emotional states for education and training purposes. It has traditionally relied on time- and skill-intensive design work by trained CG (Computer Graphics) artists, who utilized dedicated software equipped with geometric deformation and shape interpolation. The prevailing shape and motion capture technology has shifted this paradigm, allowing the algorithmic reconstruction of 3D face shape and motion of real people. Multi-view acquisition systems or 4D (3D+time) laser scans have been employed to capture the dynamic face geometry of the performer, even in real-time scenarios [5, 15, 17, 65]. Such motion capture techniques have been supported by deformation transfer or animation retargeting [6, 60, 91, 99], enabling the reuse of captured animations on new individuals or imaginary faces without having to capture shapes specific to the new target face. This line of research has facilitated the construction of 4D face datasets and the derivation of priors from the data. Consequently, this advance has greatly amplified data-driven approaches to face modeling, paving the way for recent deep learning-based methods. Reconstructive methods train the model to regress the 3D facial geometry from a 2D video, often in a frame-by-frame manner [18, 25, 31, 88] without modeling the temporal aspect of shape evolution. Generative models like Generative Adversarial Nets (GANs) [28] and Variational Autoencoders (VAEs) [45] have been adapted to learn and sample the distribution of facial shapes and expression poses. However, with a few exceptions [80, 94], most existing works focus on the generation of 2D facial poses [69, 97], or expression videos [8, 61, 89, 94, 95], leaving the challenging task of 4D geometry modeling largely unexplored.

In this paper, we address the challenging problem of 3D dynamic facial expression generation, one that has not yet received a lot of attention. Most available 3D facial expression datasets [12, 16, 24, 73, 101] come in the form of dense triangular meshes containing thousands of vertices. It is computationally expensive to train a generative model directly using all the vertices. Therefore, similarly to most successful models for 3D facial animation generation, we use a set of predefined 3D face landmarks to represent a face. Typically, landmarks are located on facial features that are highly mobile during animation, such as the face outline, eyes, nose, and mouth. The specific aim of the 3D facial animation generation is to learn a model that can generate facial expressions that are realistic, appearance-preserving, rich in diversity, with various ways to condition it such as categorical expression labels. Prior works that have attempted to model the temporal dimension of the face animation [61, 62, 80, 93] mostly leverage auto-regressive approaches, such as Long short-term memory (LSTM) [39] and Gated recurrent units (GRUs) [13]. Here, we propose to use a Denoising Diffusion Probabilistic Model (DDPM) [36, 82, 85], a generative approach that has achieved remarkable success in several domains. A DDPM has the nice property of being trainable unconditionally whereas the reverse process can still be conditioned using, a classifier-guidance [20], for instance. This allows us to define the following paradigm: a DDPM is learned unconditionally and several downstream tasks associated with several conditional generations are developed from the same learned model, such as expression control (with label or text), expression filling (with partial sequence(s)), or geometry-adaptive generation (with facial geometry). This makes the proposed approach highly flexible and efficient, benefiting from the generative power of diffusion models while circumventing their limitations of being resource-hungry and difficult to control.

We note that, concurrent to this work, several works have also adopted diffusion models for human motion generation [43, 87, 100]. However, to the best of our knowledge, we are the first to adapt diffusion models to 3D face expression generation. More importantly, although approaches developed in [43, 87, 100] enable different forms of conditioning, they require the diffusion model to be retrained for each way of conditioning.

While the task of 3D facial animation generation has been reduced to the estimation of a temporal sequence of 3D face landmark sets, it is then necessary, in a second task, to compute a sequence of

animated meshes. We use an encoder-decoder model similar to [62], which retargets the expression of a 3D face landmark set to the neutral 3D face mesh by computing its per-vertex displacement, in a frame-by-frame manner. Unlike [62], however, we take into account the different morphological shapes of the neutral mesh to adapt the estimation of per-vertex displacements. Results thus obtained validate the effectiveness of the proposed approach.

In summary, our key contributions are as follows: (1) We successfully use a DDPM to propose an original solution to the conditional generation of 3D facial animation. To the best of our knowledge, it is the first to adopt a diffusion-based generative framework in 4D face modeling. (2) We train a DDPM unconditionally and develop several downstream tasks by conditioning the reverse process. In addition to improving the efficiency of training, this paradigm makes the approach highly versatile and easily applicable to other downstream tasks. (3) In various evaluations, the landmark sequence generation and landmark-guided mesh deformation outperform SOTA methods.

2 RELATED WORK

2.1 Denoising Diffusion Probabilistic Models

3D facial expression generation is achieved here by adapting a DDPM. The surge in popularity of diffusion models is primarily due to this model [36], which has made significant advancements in high-resolution image generation.

A Denoising Diffusion Probabilistic Model (DDPM) is a Markov chain aimed at synthesizing data (that resembles the training data) from noise within a finite number of sampling steps. The model's parameters undergo training via variational inference so as to reverse a diffusion process – a Markov chain that gradually adds noise to the initial data. Such diffusion models have a wide range of applications across various fields. In computer vision, diffusion models have demonstrated impressive capabilities for image generation [72, 75, 78], super-resolution [37, 48, 79], image inpainting [53], image translation [68, 77], and semantic segmentation [4, 9, 29]. They also demonstrated capability in handling 3D data to generate and complete point clouds [54, 55, 102], as well as handling time series data for human motion generation [87, 100], time series forecasting and imputation [1], and video generation [32, 35, 38, 98]. Additionally, diffusion models have demonstrated significant potential in the realm of natural language processing with numerous applications and use cases [2, 11, 21, 27, 51]. They find application in audio synthesis as well [46].

However, diffusion models involve a complex generation process and suffer from slow sampling and have no encoding functionality. Several approaches have been proposed to address the limitations of DDPM. One such approach is Denoising Diffusion Implicit Model (DDIM) [84], which accelerates the sampling process by reducing the number of required sampling steps. DDIM also allows for deterministic reverse processing, enabling various image editing possibilities [33]. However, DDIM inversion can result in instability and distorted reconstructions. To address this issue, [92] proposes a novel approach inspired by the coupling layers in normalizing flow models [22] providing mathematically exact inversion. In addition, several strategies exist to improve the performance of DDPMs, such as learning the variances of the reverse diffusion process, using a cosine noise schedule, and adding extra loss terms to optimize the variational lower-bound [59].

Note that each generative model has its own advantages and disadvantages. For example, VAE [45] has nice encoding capabilities; however, it tends to lose high-frequency information of images. On the other hand, GANs [28] have the ability to produce high-quality images, but they are challenging to train and often prone to mode collapse. Therefore, each method has its own suitable application scenarios. For instance, diffusion models are preferred when prioritizing image quality over generation time. At the same time, efforts are underway to integrate different models aimed

at overcoming their respective limitations, as seen in initiatives like DiffuseVAE [63] and Diffusion-GAN [96].

2.2 Conditional generation

The conditional generation has been investigated since its inception, and its importance continues to increase. The generation can be conditioned by various modalities, ranging from a simple label to an image, or a sentence, which guides or controls the generation process. For example, image translation involves creating images conditioned on other images. Text-to-image generation uses text to control the characteristics of the generated images. In image captioning, conversely, the image conditions the text generation.

The first approach to achieve conditional generation involves the use of conditional generative models, which are trained by incorporating specific conditions. For instance, in the context of diffusion models, a commonly used approach is to train a label embedding and combine it with the time step embedding. The resulting combined information is then fed into the noise approximator, which generates the conditional noise samples [51]. Several models have been proposed to accommodate various forms of label conditions for GANs or VAEs, such as CVAE [83], CGAN [57], PPGN [58], and CVAE-GAN[3]. The neural network must include a mechanism to effectively incorporate the conditional information. One straightforward approach is to concatenate the label with an intermediate output of the model. In the context of conditional human motion generation [64], learnable tokens are used. These tokens, of the same size as the latent representation, are added as offsets to the latent space. We employ a comparable mechanism to integrate conditional information in our landmark-to-mesh deformation task (Sec. 3.3), focusing on generating a sequence of animated meshes from a sequence of landmarks, conditioned on the morphological shape of a neutral mesh.

Another line of approach related to ours is conditioning models that were originally learned unconditionally. This enables us to perform conditional generation using high-performance pretrained generative models. Note that the training data used for learning the “independent conditioning task” is not required to be the same as the data used for training the generative model. Moreover, this paradigm enhances training efficiency, especially when there are multiple ways to apply conditioning. In this work, we unconditionally train a DDPM, after which we delve into several downstream tasks by conditioning the reverse process in different ways. One commonly employed approach in conditioning generation within diffusion models involves classifier-guided sampling [20, 85]. To condition a pre-trained diffusion model, it leverages the gradients of an independent classifier. This entails training a classifier on noisy data and subsequently using the gradients to guide the diffusion sampling process towards a specific class label.

Note that conditioning an unconditional model can also be achieved with other generative models than DDPM, especially through latent code manipulation [14, 81]. These methods begin by computing the latent representation of an image, and then adjusting this representation in a way that carries semantic meaning.

2.3 Facial expression generation

Deep generative models [28, 44, 45, 74, 82] have proven effective at high-quality image synthesis, such as content-preserving image rendering with different styles, and the generation of images depicting learned objects. For 2D images, these models have also shown to be beneficial to facial expression transfer and expression editing tasks. However, the majority of existing solutions address the problem of *static* expression synthesis. Here we review some recent advances achieved in *dynamic* facial expression generation, i.e. modeling and predicting the temporal evolution of poses

elicited by facial expressions.

2D facial expression video generation. Building on the achievements in deep image generation, recent works have tackled the problem of generating 2D facial expression videos [8, 61, 89, 94, 95]. Methods like [42] focus on facial expression transfer, where the dynamics of the facial expression is transferred from a source to a target face. In other works, the dynamics of facial expression is encoded using recurrent generators. For example, Wang et al [93] proposed to generate sequence of 2D landmarks by a conditional LSTM, which is trained over a manifold of landmark images found by a VAE. Addressing the one-to-many generation problem, they generate multiple sequences given a neutral face image, which are translated to face videos. Recent methods seek to learn the spatiotemporal embedding and generate all frames simultaneously, rather than generating frames sequentially. G³AN [94] learns disentangled representations of the facial video, appearance and motion by a three-stream generator, each taking into account appearance features, motion features, and the fusion of the two toward video generation. Other works have investigated modelling and learning the appearance and the motion in a disentangled manner. MoCoGAN [89] decomposes the video into content and motion: An image-based generator generates the content, an image representing the identity of the person, while a GRU-based motion GAN generates the motion-dependent appearance variations in videos through sequential sampling from the motion space. ImaGINator [95] presents an image to video generation consisting of an image encoder and a video decoder. The appearance extracted from the encoder is fused with motion feature in multi-level feature space via a spatio-temporal fusion in the video decoder to simultaneously generate appearance and motion, i.e. video. Two stream discriminator assesses the authenticity of the generated images and videos, respectively. Such an encoder-decoder based approach has been also adopted in FEV-GAN[8], where an encoder is trained to extract identity- and spatial- features from the neutral face image, and a video decoder utilizes them along with a label to generate realistic video.

Generative models for body animations. Several works have shown the sequence generation on the body motion [30, 47, 49, 64, 86]. This is mainly due to the compact, readily available skeleton-based representation of the body [52], the relatively large set of action vocabulary, and the availability of rich 3D body motion datasets [40, 47, 56, 66, 70]. Unfortunately it is not yet the case with the 3D facial expression.

Dynamic 3D facial expression synthesis. To our knowledge, dynamic 3D facial expression synthesis has not been fully explored. [67] synthesizes realistic high resolution facial expressions by using a deep mesh encoder-decoder like architecture to estimate the displacements which are then added to a neutral face frame. [80] deploys LSTMs to estimate the facial landmark changes, which are then used to guide the deformation of a neutral mesh via a Radial Basis Function network. However, both works focus on the displacement estimation for a given expression and do not consider conditional generations. The closest work to ours is Motion3DGAN [62] which extends the aforementioned MotionGAN [61] to model the dynamics of 3D landmarks. The learned distribution of 3D expression dynamics by a WGAN over the hypersphere space is sampled with a condition to generate landmark sequences, which are then fed into a mesh decoder to deform a neutral 3D face mesh frame-by-frame. Our work has several advantages over their work. Benefiting from the power of diffusion models, we model the input distribution without requiring any extra preprocessing, and can learn from sequences of different lengths. Moreover, our framework offers a highly versatile and efficient alternative, as we train a DDPM unconditionally and different conditional generations can be performed solely during the reverse process in a plug-and-play manner.

Given the scarcity of existing work on 3D facial animation generation, we compare our work with some generator models originally dedicated to human motion synthesis, including Action2motion [30] and ACTOR [64].

3 METHOD

At the core of our approach is a DDPM-based model to generate a 3D landmark sequence $x = \{L_1, \dots, L_F\}$ where a frame $L_f \in \mathbb{R}^{N \times 3}$ (for $f = 1$ to F) represents the 3D coordinates of N landmarks. Note that the 3D arrangement of a landmark set L_f implicitly encodes the geometric information specific to the facial anatomy of an individual, and can be viewed as a mixture of the facial identity shape at a neutral pose L and the pose-induced shape change, i.e. $L_f = \Delta L_f + L$. The method is composed of two tasks: First, a DDPM is trained unconditionally (Sec.3.1), whereas conditional generations are obtained by conditioning the reverse process. Different forms of conditioning can be performed, leading to several downstream tasks (Sec.3.2). Afterwards, our landmark-guided encoder-decoder (Sec.3.3) estimates ΔM_f at each frame (for $f = 1$ to F), using a target neutral face mesh M and ΔL_f as input. The desired animation mesh sequence $\{M_1, \dots, M_F\}$ is obtained by adding the estimated displacement ΔM_f to M at each corresponding frame, i.e. $M_f = M + \Delta M_f$. The overview of the proposed method is illustrated in Fig.1.

Note that directly training from and generating full meshes may be beneficial but raises technical issues since the model becomes computationally and memory intensive. An alternative is to utilize diffusion models directly in the latent space of autoencoders [75], or a pre-constructed parameter space of 3D face. Our work can be viewed as akin to the latter approach, except that we use a heuristically defined feature space, i.e., the landmark space, instead of a learned latent space. This choice has been validated by the quality of the reconstruction obtained by the landmark-guided encoder-decoder (Tab. 5).

3.1 Denoising Diffusion Probabilistic Models

DDPMs are latent variable models where the latent variables x_t (for $t = 1$ to T) have the same dimension as the original data $x_0 \sim q(x_0)$. In our work, x_0 is a landmark-based facial animation data: $x_0 = \{L_1, \dots, L_F\}$. Note that it is contrary to most prior works which generate only the displacements ΔL_f [61, 62, 80]. Training our model to generate L_f directly allows it to learn to produce quality expressions that are consistent with the inherent facial morphology.

DDPM consists of two distinct processes. The unparameterized diffusion process is a Markov chain, that gradually adds noise to the initial data. Conversely, the parameterized reverse process involves a Markov chain specifically designed to generate, within a limited number of sampling iterations, data resembling the training dataset and originating from noise. The parameters of the reverse process are trained so as to reverse the diffusion process.

More formally, the joint distribution $p_\theta(x_{0:T})$ from which we derive the likelihood $p_\theta(x_0) = \int p_\theta(x_{0:T}) dx_{1:T}$ is the reverse process whereas the approximate posterior $q(x_{1:T}|x_0)$ is the forward process or diffusion process. The diffusion process produces gradually noisier samples (x_1, x_2, \dots, x_T) by adding Gaussian noise to the initial data x_0 according to a variance schedule β_1, \dots, β_T [36]:

$$q(x_{1:T}|x_0) = \prod_{t=1}^T q(x_t|x_{t-1}) \quad (1)$$

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t I). \quad (2)$$

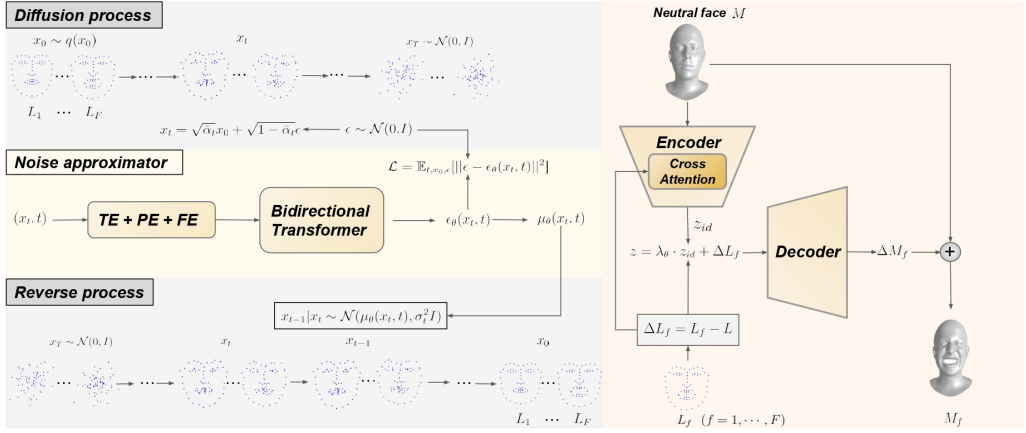


Fig. 1. Overview of the proposed approach. Generally, the diffusion process is used to train the noise approximator while the reverse process is used to sample x_0 from the distribution q . But some tasks developed in Sec. 3.2 require both processes for sampling. The bidirectional transformer takes as input the sum of the outputs of three embedding layers: the temporal embedding layer (TE) that takes as input t , the positional encoding layer (PE) that takes as input an integer sequence from 1 to F , and the feature embedding layer (FE) that takes x_t . The landmark-guided encoder-decoder retargets the expression of L_f onto the input mesh M to estimate M_f at each frame.

We can derive from Eq. 2 the following property [36] which allows us to train the diffusion model efficiently at an arbitrary time step t :

$$q(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)I), \quad (3)$$

where $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ and $\alpha_t = 1 - \beta_t$.

x_T follows a near-isotropic Gaussian distribution provided that a well-behaved schedule is defined and that T is sufficiently large. DDPM [36] uses this property to sample the target distribution q ($x_0 \sim q(x_0)$). This is achieved by reversing the diffusion process: It begins by sampling x_T from $\mathcal{N}(0, I)$. Next, the reverse process generates progressively less-noisy samples $x_{T-1}, x_{T-2}, \dots, x_1$ until $x_0 \sim q(x_0)$ is obtained, by repeatedly sampling x_{t-1} from $p_\theta(x_{t-1}|x_t)$ by using Eq. 5. This reverse process is formally defined as a Markov chain with learned Gaussian transitions whose mean and variance are estimated by a neural network of parameter θ :

$$p_\theta(x_{0:T}) = p(x_T) \prod_{t=1}^T p_\theta(x_{t-1}|x_t) \quad (4)$$

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)), \quad (5)$$

where $p(x_T) = \mathcal{N}(x_T; 0, I)$. As in [36], we set $\Sigma_\theta(x_t, t)$ to $\sigma_t^2 I$. This is a reasonable choice for generating quality samples, provided that T is chosen to be sufficiently large [59]. Note that estimating $\Sigma_\theta(x_t, t)$ allows sampling with many fewer steps [59].

Several possibilities can be considered to parameterize $\mu_\theta(x_t, t)$ in Eq. 5. [36] shows that approximating the noise ϵ that appears in the following equation:

$$x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, \quad (6)$$

is a suitable choice, especially when combined with a simple loss function (See Eq. 8). Note that Eq. 6 is a different way of writing Eq. 3 ($\epsilon \sim \mathcal{N}(0, I)$). Finally, the term $\mu_\theta(x_t, t)$ can be computed from the approximation of ϵ , denoted as $\epsilon_\theta(x_t, t)$:

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

Diffusion models can be trained by optimizing the usual variational bound on negative log-likelihood, but we adopt here the simplified objective function proposed in [36]:

$$\mathbb{E}_{t, x_0, \epsilon} [\|\epsilon - \epsilon_\theta(x_t, t)\|^2], \quad (8)$$

where the term x_t is computed from Eq. 6.

Many previous works [20, 36, 59, 75], especially those for modeling 2D images, have utilized a UNet-like structure[76] to model the mean $\mu_\theta(x, t)$ or the noise $\epsilon_\theta(x, t)$. Here we employ a bidirectional transformer (BiT) [19] to efficiently capture the temporal characteristics of x_t .

3.2 Downstream tasks

The DDPM is trained in an unconditional manner, and from this single trained model, various downstream tasks can be derived. Here, we demonstrate expression control tasks through labels or texts, and expression filling using partial sequence, offering tools for creators to explore new expression sequences. The last downstream task, geometry-adaptive generation, enables us to generate a landmark sequence conditioned to a facial anatomy. The pseudo code for each task can be found in Appendix C.

Conditioning on expression label (label control). The task is to perform a conditional generation according to the expression label y . Conditioning the reverse process of an unconditional DDPM is achieved by using the classifier-guidance [20, 51, 85]. First, we train a classifier that predicts the label y given a latent variable x_t (and t). Here the classification is conducted with a BiT by adopting the usual approach of adding an extra learnable classification token [19]. Note that the BiT discussed here should be distinguished from the other BiT in the diffusion model and is used to condition its reverse process. It is achieved by sampling x_t according to the distribution:

$$p_{\theta, \phi}(x_t | x_{t+1}, y) \propto p_\theta(x_t | x_{t+1}) p_\phi(y | x_t), \quad (9)$$

where ϕ represents the parameters of the classifier. Sampling of Eq. 9 can be achieved approximately [82] by sampling from a Gaussian distribution similar to the unconditional transition operator $p_\theta(x_t | x_{t+1})$, but with its mean shifted by a quantity proportional to $\Sigma_\theta(x_t, t) \nabla_{x_t} p_\phi(y | x_t)$.

Instead of sampling Eq. 9, we used an alternative way, as proposed in [51]: x_t is computed so as to maximize the \log of Eq. 9. A hyperparameter λ is used to adjust the trade-off between fluency ($p_\theta(x_t | x_{t+1})$) and control ($p_\phi(y | x_t)$), leading to a stochastic decoding method that balances maximizing and sampling $p_{\theta, \phi}(x_t | x_{t+1}, y)$. As in [51], optimization is achieved by running 3 steps of the Adagrad [23] update for each diffusion step (Alg. 1 of App. C).

Conditioning on text (text control). We also use in this task a BiT guidance, but instead of estimating a label from x_t and t , the BiT outputs a vector of dimension 512 (the softmax layer is removed). As in [86], the BiT is trained so as to increase the cosine similarity between its output and the textual features extracted with CLIP [71] from the text associated with x_0 . Conditioning the reverse process according to the text c is then achieved (Alg. 2 of App. C) by adapting the procedure presented for the label control: x_t is computed so that it maximizes:

$$\lambda \cdot \log(p_\theta(x_t | x_{t+1})) + \cos(\text{BiT}(x_t, t), \text{CLIP}(c)). \quad (10)$$

Conditioning on partial sequence (expression filling). Similarly to inpainting whose purpose is to predict missing pixels of an image using a mask region as a condition, this task aims to predict missing frames of a temporal sequence by leveraging known frames as a condition. The sequence x_0 is composed of F frames, some of which are unknown. Let \mathcal{S}_K and \mathcal{S}_U denote respectively the set of indices associated with known and unknown frames, and let $x|_{\mathcal{S}}$ denote the subsequence containing only the frames of x whose indices belong to \mathcal{S} .

Since $x_0|_{\mathcal{S}_K}$ is known, note that $x_t|_{\mathcal{S}_K}$ can be drawn according to Eq. 3. Indeed, each component of x_t can be drawn independently since $q(x_t|x_0)$ is an isotropic normal distribution. Sampling from the reverse process conditioned on a partial sequence can also be achieved as follows: X_T is first determined: $x_T|_{\mathcal{S}_U}$ is drawn from $\mathcal{N}(0, I)$ and $x_T|_{\mathcal{S}_K}$ according to Eq. 3. Then, computing x_t from x_{t+1} is achieved in two steps: First, a temporal sequence \hat{x}_t is simply drawn from $p_\theta(\cdot|x_{t+1})$ (it is the way to compute x_t in the usual case). $x_t|_{\mathcal{S}_U}$ is set to $\hat{x}_t|_{\mathcal{S}_U}$, while for known frames, $x_t|_{\mathcal{S}_K}$ is directly drawn according to Eq. 3 (Alg. 3 in App. C). Despite its simplicity, this strategy gives satisfactory results as we will demonstrate through qualitative validation in later sections of this paper, provided that the partial sequence is of sufficient length.

Geometry-adaptive generation. Given the facial geometry of a specific subject, a generation can be performed as a special case of expression filling: \mathcal{S}_K is set to $\{1\}$ or to $\{F\}$ (F is the sequence length) and the unique known frame associated with $x_0|_{\mathcal{S}_K}$ is set to the neutral face L of the subject. The remaining sequence is considered as unknown, for which the model performs an expression filling.

However, we observed that the generated frames may not always smoothly connect to the given frame, a problem that did not arise when the partial sequence remained long enough. In the context of image inpainting, [53] also shows that the simple sampling strategy used for the expression filling task may introduce disharmony. A more sophisticated approach has been proposed so as to harmonize the conditional data $x_t|_{\mathcal{S}_K}$ with the generated one $x_t|_{\mathcal{S}_U}$ [53]. In order to achieve better convergence properties of the algorithm while maintaining its simplicity, we derive the sequence with five iterations, each with a slight modification: For the first iteration, $x_T|_{\mathcal{S}_U}$ is drawn, as previously, from $\mathcal{N}(0, I)$. For the following iterations, $x_T|_{\mathcal{S}_U}$, as $x_T|_{\mathcal{S}_K}$, is drawn according to Eq. 3 where x_0 is the result obtained from the previous iteration. By doing so, we expect $x_T|_{\mathcal{S}_U}$ and $x_T|_{\mathcal{S}_K}$ to be harmonized progressively, thus leading to the improved harmonization of $x_t|_{\mathcal{S}_U}$ and $x_t|_{\mathcal{S}_K}$ along the iterations.

Note that this process can also be easily guided by a classifier (as in the label control) so as to generate a desired facial expression starting from a given facial anatomy (See Alg. 4 in App. C).

3.3 Landmark-guided mesh deformation

To obtain the full mesh sequence $\{M_1, \dots, M_F\}$ from $\{L_1, \dots, L_F\}$, one could use existing fitting methods such as FLAME [50] or DL-3DMM [26] so as to preserve both the facial anatomy and the expression encoded in the landmark frames. However, the meshes generated through the linear blending models tend to lack intricate details of facial geometry, resulting in dull, lifeless shapes. Thus, in our work, we retarget the expression encoded in L_f to the facial geometry given as a (realistic) input mesh M , as in [62]. The mesh M is assumed to be at its neutral pose with a predefined topology [50]. Each mesh frame M_f should retain the facial identity shape M , combined with the expression-driven shape change encoded in $\Delta L_f = L_f - L$ (ΔL_f represents landmark displacements at f -th frame). This is achieved by our encoder-decoder network that takes both M and ΔL_f as input and predicts ΔM_f at each frame, which is respectively added to M to obtain the final mesh sequence: $M_f = M + \Delta M_f$. This is similar to the Sparse2Dense mesh decoder proposed in [62], except that only ΔL_f (and not M) is used to predict ΔM_f in their work. In our approach, on

the other hand, we take into account the different morphological shapes of the neutral mesh M to adapt the estimation of per-vertex displacements ΔM_f .

In order to benefit from the consistent and quality expressions adapted to the facial morphology by the DDPM, one can extract a landmark set L_M from a mesh M , perform the geometry-adaptive task on it to generate a sequence involving L_M , and target it to M by the landmark-guided mesh deformation.

Encoder & decoder. Inspired by the Sparse2Dense mesh decoder of [62], we develop an encoder-decoder architecture based on spiral operation layers. The encoder contains a backbone consisting of five spiral operation layers [7] that extracts the features of M . In addition, we propose to incorporate a cross-attention mechanism [90] to account for the possible influence of the characteristics of M on the impact of ΔL_f on each vertex of M : It enables us to find the relevant features of the mesh M that can help predict a latent representation (of ΔM_f) according to ΔL_f . More specifically, the *query* is derived from a linear embedding of ΔL_f (computed by a fully-connected layer FC) and the *key, value* pairs from the output of the backbone (i.e. features of M) denoted as F_e . The output of the attention layer writes:

$$\text{softmax} \left(\frac{FC(\Delta L_f) \cdot F_e^T}{\sqrt{d}} \right) F_e, \quad (11)$$

where d is the dimension of F_e . Then a linear layer maps the vector of Eq. 11 to the *identity-aware* representation z_{id} , which is further shifted by the landmark displacement ΔL_f to obtain the final latent representation: $z = \lambda_\theta \cdot z_{id} + \Delta L_f$, where the weight parameter λ_θ is a learnable parameter.

We use the same decoder as [62]. It consists of a linear layer and five spiral operation layers. It takes the latent representation z as input and outputs the per-vertex displacement ΔM_f . M_f is then set to $M + \Delta M_f$. The model is learned using the loss function proposed in [62].

4 EXPERIMENTAL SETTING

As proposed in [36], we set a linear noise schedule starting from $\beta_1 = 1e - 4$ to $\beta_T = 0.02$, and σ_t^2 is set to β_t . T is set to 2000. We train the model on 200K iterations with a learning rate of $1e - 4$ and a batch size of 256. The hyperparameter λ that is used to guide the sampling of the reverse process is set to 0.01 as in [51].

CoMA dataset [73] is a commonly used 4D facial expression dataset in face modeling tasks [7, 41], consisting of over a hundred 3D facial animation sequences captured from 12 subjects, each performing 12 facial actions (“high smile”, “mouth up”, etc.). Each data is composed of a triangular mesh of 5023 vertices undergoing some deformation elicited by an expression.

BU-4DFE dataset [101] contains a total of 606 sequences of 83 landmarks extracted from a sequence of 3D facial scans. Six basic emotional expressions (“anger”, “disgust”, “fear”, “happy”, “sad”, and “surprise”) of 101 subjects have been recorded.

To efficiently capture the essence of expressions while maintaining a small data size, we use 68 landmarks defined by FLAME[50] for both databases. The different sequences have been manually divided into sequences of around 40 frames. Some of them start from the neutral pose and end with a maximal expression intensity (we call them sequences of type N2E), while others evolve from the maximal expression intensity to the neutral face (they are called of type E2N). Since a majority of the methods utilized for comparison necessitate the sequences to share a same length, linear interpolation has been carried out so as to obtain sequences of length 40. Note, however, that our method can deal with sequences of various lengths (See App. A). As a result, our dataset comprises 689 sequences from the CoMA dataset and 1,212 sequences from the BU-4DFE dataset, totaling 1,901 sequences with 18 facial actions. Different sequences have been used depending on

the specific task at hand. Unless otherwise specified, solely the sequences from the CoMA dataset have been utilized.

5 RESULTS INVOLVING VARIOUS CONDITIONAL GENERATIONS

Here we describe the results we obtained on the various conditional generations. Throughout this section, a classifier that predicts the expression from a sequence independently of its type (see Section 4) is called a classifier of type I (order-Insensitive), whereas a classifier of type S (order-Sensitive) predicts both the expression class and the expression type (either N2E or E2N).

To assess the fidelity of the generated expressions in relation to the control, an independent classifier (which we denote as IC) is trained to predict the label from a sequence x_0 . We use one LSTM layer followed by a linear layer, as in [62]. The model’s ability to generate a desired expression is assessed by the classification accuracy of the IC tested on the generated expressions. Additionally, the quality and the diversity of the generated sequences are assessed by using the Frechet Inception Distance (FID) score [34], a widely used metric for comparing the distribution of fake data with that of real data. It is computed from the output of the linear layer of the IC.

5.1 Label control

The proposed approach is compared with several SOTA methods that perform conditional sequence generation: Action2Motion [30], Motion3DGAN [62] and ACTOR [64]. The BiT-based classifier is used to guide the reverse process, as well as the IC are of type I. Quantitative evaluation results as measured by the classification accuracy and the FID score are summarized in Table 1, which confirms that the proposed approach outperforms all SOTA methods. Fig. 2 shows some illustrative results: Our model generates various realistic and quality expressions adapted to different facial geometries. Videos presented in the project website (<https://github.com/ZOUKaifeng/4DFM>) demonstrate the generated expressions and offer qualitative comparisons among these methods. Visual observation of these videos confirms that sequences generated by our approach are more expressive. The diversity of the generated sequences in terms of both expression and facial anatomy is illustrated in Appendix B.

Table 1. Performance of different methods for generating desired expressions has been evaluated by measuring the classification accuracy and the FID score. We report as ground truth the FID and the accuracy computed on the test dataset, assuming that an ideal method could have generated it.

Model	CoMA		BU-4DFE	
	Acc	FID	Acc	FID
Ground truth	83.78%	2.77	99.51%	6.02
A2M [30]	52.36%	29.44	80.83%	19.64
MoGAN [62]	80.76%	7.72	99.26%	13.29
ACTOR [64]	81.40%	7.11	99.13%	14.56
Ours	84.97%	6.79	99.89%	12.37

5.2 Text control

To demonstrate this task, we have increased the vocabulary of our dataset by merging CoMA and BU-4DFE. In the first experiment, the raw text label is used to condition the animation (we call it *raw text* task) and the IC used for the evaluation is of type I. In the second experiment, the



Fig. 2. Animated mesh sequences guided by the label “mouth side” (top), “mouth extreme” (middle), and “cheeks in” (bottom). The meshes are obtained by retargeting the expression of the generated x_0 on different neutral faces.

description of a sequence is enriched to be a short sentence such as “from the neutral face to the raw text label”, or “from the raw text label to the neutral face” (we call it *enriched text* task) and the IC used for the evaluation is of type S.

We compare our results with those of MotionClip [86]. Quantitative results are shown in Table 2. Classification accuracies obtained with the proposed method are slightly higher than those of MotionClip, with FID scores significantly lower. Sequences created by MotionClip are actually realistic but the FID scores are high, due to the lack of diversity in the generated sequences.

Fig. 3 shows illustrative examples obtained with the proposed approach. Note that our model is able to create animated meshes that combine different types of expressions by compositing a text combining different types of expressions. For the complete sequences as well as the qualitative comparisons, readers may refer to the project website.

Table 2. Quantitative evaluation of the text control task. Classification accuracy and FID are computed for the raw text task (rtt, left) and for the enriched text task (ent, right).

	Acc (rtt)	FID	Acc (ent)	FID
Ground truth	86.02%	3.67	74.40%	4.56
MotionClip	80.67%	42.19	58.33%	38.83
Ours	82.01%	9.46	64.38%	11.34

5.3 Expression filling

Given a partial sequence of an expression, the model can fill up the missing frames. Three experiments have been conducted: In the filling from the beginning (FFB) or the filling from the end (FFE) cases, the length l of the partial sequence is drawn uniformly in [10, 30]. In the filling from



Fig. 3. Text-driven generation results obtained by the *enriched text* task (“from neutral face to bareteeth” (top)), and by the *raw text* task (“angry mouth down” (middle), “disgust high smile” (bottom)). The input texts used for the *raw text* task are the combinations of two terms used for training. For instance, “disgust high smile” is a new description that hasn’t been seen before, which combines “disgust” and “high smile”.

the middle (FFM) case, l frames have been given at the beginning and at the end of the sequence, respectively. l is uniformly sampled in [5, 15].

The proposed approach for expression sequence filling is compared with a mean imputation strategy. To evaluate the result, an IC (of type I) is trained, so as to check if the filled data has the same expression class as the original one. Results are shown in Table 3. The expression label of the partial sequence is well-captured and reflected in the filled part, leading to an improved classification accuracy especially for the FFM case, where the classification accuracy is comparable to that obtained for the ground truth (Table 1). Classification accuracies obtained in the FFE and FFB cases are lower due to the content of the sequences. As an example, when the partial sequence is associated with the beginning of a sequence of type N2E, it may be composed, at worst, of neutral faces only, or at best of less expressive faces. This is worsened by the fact that sometimes certain expressions appear only at the end of the sequences. This is contrary to the FFM case, where the partial sequence contains both the neutral and the most expressive poses.

Finally, there is a significant improvement of FID score after filling with the proposed approach. Furthermore, our videos presented on the project website illustrate that the generated sequences are smoothly connected to the given partial sequence.

5.4 Geometry-adaptive generation

We have conducted the geometry-adaptive generation task by using classifier guidance so as to generate a desired facial expression from a given facial anatomy (Alg. 4 of App. C). The BiT used for guidance and the IC used for evaluation are both of type S. \mathcal{S}_K is set to $\{1\}$ if the chosen label is associated with N2E sequences, and to $\{F\}$ otherwise.

Quantitative results are shown in Table 4. The classification accuracy is close to the ground truth, and the visual inspection of the video sequences on the project website shows no gap between the generated frames and the enforced one.

Table 3. Quantitative evaluation of the expression filling task for three different locations of the missing part. Accuracy and FID are computed on the sequences obtained by the mean imputation strategy, and by our diffusion model. Note that accuracy is 83.78% and FID is 2.77 for the ground truth in all cases (FFE, FFM, FFB).

	Mean Imputation		Ours	
	Acc	FID	Acc	FID
FFE	60.15%	25.67	67.18%	5.51
FFM	56.25%	17.68	85.93%	5.06
FFB	53.90%	27.32	70.31%	5.22

Table 4. Quantitative evaluation of the geometry-adaptive generation task.

	Acc	FID
Ground truth	71.01%	5.57
Geometry-adaptive	70.43%	9.26

While App. B illustrates the diversity of generated expressions when the model is conditioned on the expression label, we study here the same type of diversity but when the facial geometry of a specific subject is enforced in the conditioning process. To this end, a landmark set L_M has been extracted from a given mesh M . The geometry-adaptive generation task is performed so as to generate a sequence containing L_M , and exhibiting an expression corresponding to a given label y . Then, the generated sequence is targeted to M with the landmark-guided mesh deformation.

Fig. 4 illustrates the variety of expressions we thus obtained by using a same facial anatomy L_M and a same label y (either “eyebrow” or “high smile”), which confirms that the proposed approach is able to generate expression sequences with sufficient level of diversity, even if a same facial anatomy is used for conditioning.




Fig. 4. Diversity of expressions generated with the label “eyebrow” (left), and “high smile” (right) in the geometry-adaptive generation task. All illustrated sequences are of type N2E. Note that eyebrows can be either lowered (the second and third rows) or raised (the first row). Although the poses of maximal expression intensity look all similar in the three sequences of “high smile”, their temporal properties are significantly different.

6 RESULTS RELATED TO LANDMARK-GUIDED MESH DEFORMATION

To the best of our knowledge, only [62] and our work estimate M_f from M and ΔL_f . Note that both approaches use spiral convolution. For the comparative experiments, we also adapt two autoencoders: CoMA [73], which uses Chebyshev convolution and a mesh pooling, and the autoencoder proposed in [10] (the encoder and decoder consisting of three layers of linear, nonlinear, and linear activation units, respectively). Both decoders, which originally take the latent representation of the input mesh as input, have been modified so as to consume the concatenation of the latent representation with ΔL_f .

We conducted two series of experiments: Either 3 expressions (expression split) or 3 subjects (identity split) have been excluded from the training set, and the performance of the model is evaluated on the excluded data. The mean per-vertex Euclidean error between the generated meshes and their ground truth has been measured to assess the performance.

Quantitative results are shown in Table 5. While the three methods based on spiral convolution generally yield effective results, our approach outperforms the others, thus confirming the advantage of the cross-attention layer, in particular.

Table 5. Per-vertex reconstruction error (mm).

Method	Expression split	Identity split
Linear [10]	0.67 ± 0.76	0.73 ± 0.77
CoMA[73]	0.58 ± 0.63	0.63 ± 0.67
S2D [62]	0.52 ± 0.59	0.55 ± 0.62
Ours(w/o attention)	0.54 ± 0.59	0.57 ± 0.64
Ours	0.45 ± 0.51	0.50 ± 0.58

We propose to complement our quantitative analysis by a qualitative comparison of the different methods. As the "Expression split" and "Identity split" experiments yield very similar results, we focus solely on the "Identity split" experiment in the following.

Fig. 5 depicts the ground truth mesh (a) as well as the meshes generated with several approaches (b-e). Each vertex of a generated mesh is assigned a color representing the Euclidean distance to its counterpart on the ground truth mesh. As expected, the errors appear mainly on the regions that have been deformed to attain the expression. In Fig. 5, targeting an expression close to the neutral pose (first row) leads to tiny errors, whereas targeting an expression "mouth extreme" leads to errors that are mostly located near the mouth. Our approach achieves the best performance in this qualitative error measure, confirming the quantitative results described above.

Here we were able to deploy the reconstruction error as evaluation metric, since ΔL_f and M pertain to the same individual. Besides, our encoder-decoder is also capable of retargeting landmark sequences to different facial meshes, a task previously shown by S2D [62]. As demonstrated in the videos on the project website, it produces results that are visually pleasing and qualitatively comparable to S2D. However, defining an evaluation metric for this task remains a challenging future endeavor.

7 CONCLUSION

We have presented a generator model to synthesize 3D dynamic facial expressions. The dynamics of facial expressions is first learned unconditionally, from which a series of downstream tasks are developed to synthesize an expression sequence conditioned on various condition signals. Also proposed is a robust face deformation scheme guided by the landmark set, which contributes

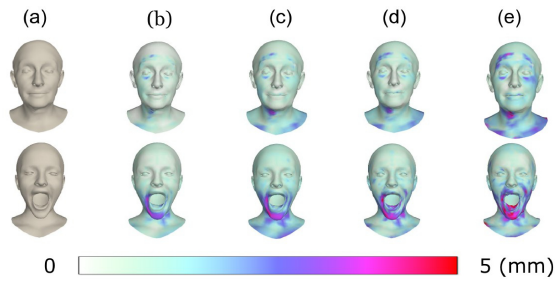


Fig. 5. Qualitative comparison of our method (b) with S2D (c), CoMA (d), and Linear (e) in the landmark-guided deformation of a given mesh. The ground truth meshes are given in the first column (a). The expression of the first row is close to the neutral face and that of the second row is taken from a sequence labeled as “mouth extreme”.

to a higher reconstruction validity. Experimental results show that the proposed method can produce plausible face meshes of diverse types of expressions on different subjects. In addition, it outperforms SOTA models both qualitatively and quantitatively. As has been demonstrated, our expression generation framework is versatile and can be used in many application scenarios including, but not limited to, label-guided generation, text-driven generation, geometry-adaptive generation, or expression filling.

ACKNOWLEDGMENTS

This work was funded by the TOPACS ANR-19-CE45-0015 project, the binational project “Synthetic Data Generation and Sim-to-Real Adaptive Learning for Real-World Human Daily Activity Recognition of Human-Care Robots (21YS2900)” granted by ETRI, South Korea, and Human4D ANR-19-CE23-0020 project funded by the French National Research Agency (ANR).

REFERENCES

- [1] Juan Miguel Lopez Alcaraz and Nils Strodthoff. 2022. Diffusion-based time series imputation and forecasting with structured state space models. *arXiv preprint arXiv:2208.09399* (2022).
- [2] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne van den Berg. 2021. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems* 34 (2021), 17981–17993.
- [3] Jianmin Bao, Dong Chen, Fang Wen, Houqiang Li, and Gang Hua. 2017. CVAE-GAN: fine-grained image generation through asymmetric training. In *Proceedings of the IEEE international conference on computer vision*. 2745–2754.
- [4] Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khrulkov, and Artem Babenko. 2021. Label-efficient semantic segmentation with diffusion models. *arXiv preprint arXiv:2112.03126* (2021).
- [5] Thabo Beeler, Fabian Hahn, Derek Bradley, Bernd Bickel, Paul Beardsley, Craig Gotsman, Robert Sumner, and Markus Gross. 2011. High-Quality Passive Facial Performance Capture using Anchor Frames. *ACM Trans. Graph.* 30 (07 2011), 75. <https://doi.org/10.1145/2010324.1964970>
- [6] V. Blanz, C. Basso, T. Poggio, and T. Vetter. 2003. Reanimating Faces in Images and Video. *Computer Graphics Forum* 22, 3 (2003), 641–650. <https://doi.org/10.1111/1467-8659.t01-1-00712>
- [7] Giorgos Bouritsas, Sergiy Bokhnyak, Stylianos Ploumpis, Michael Bronstein, and Stefanos Zafeiriou. 2019. Neural 3d morphable models: Spiral convolutional networks for 3d shape representation learning and generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7213–7222.
- [8] Hamza Bouzid and Lahoucine Ballihi. 2022. Facial Expression Video Generation Based-On Spatio-temporal Convolutional GAN: FEV-GAN. *Intelligent Systems with Applications* (2022), 200139.
- [9] Emmanuel Asiedu Brempong, Simon Kornblith, Ting Chen, Niki Parmar, Matthias Minderer, and Mohammad Norouzi. 2022. Denoising Pretraining for Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4175–4186.

- [10] Dan Casas and Miguel A Otaduy. 2018. Learning nonlinear soft-tissue dynamics for interactive avatars. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 1, 1 (2018), 1–15.
- [11] Ting Chen, Ruixiang Zhang, and Geoffrey Hinton. 2022. Analog bits: Generating discrete data using diffusion models with self-conditioning. *arXiv preprint arXiv:2208.04202* (2022).
- [12] Shiyang Cheng, Irene Kotsia, Maja Pantic, and Stefanos Zafeiriou. 2018. 4dfab: A large scale 4d database for facial expression analysis and biometric applications. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5117–5126.
- [13] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the Properties of Neural Machine Translation: Encoder–Decoder Approaches. In *SSST@EMNLP*.
- [14] Edo Collins, Raja Bala, Bob Price, and Sabine Süsstrunk. 2020. Editing in Style: Uncovering the Local Semantics of GANs. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [15] D. Cosker, E. Krumhuber, and A. Hilton. 2011. A FACS valid 3D dynamic action unit database with applications to 3D dynamic morphable facial modeling. In *2011 International Conference on Computer Vision*. 2296–2303. <https://doi.org/10.1109/ICCV.2011.6126510>
- [16] Darren Cosker, Eva Krumhuber, and Adrian Hilton. 2011. A FACS valid 3D dynamic action unit database with applications to 3D dynamic morphable facial modeling. In *2011 international conference on computer vision*. IEEE, 2296–2303.
- [17] D. DeCarlo and D. Metaxas. 1996. The Integration of Optical Flow and Deformable Models with Applications to Human Face Shape and Motion Estimation. In *2013 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE Computer Society, Los Alamitos, CA, USA, 231. <https://doi.org/10.1109/CVPR.1996.517079>
- [18] Yu Deng, Jiaolong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. 2019. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*.
- [19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [20] Prafulla Dhariwal and Alexander Nichol. 2021. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems* 34 (2021), 8780–8794.
- [21] Sander Dieleman, Laurent Sartran, Arman Roshannai, Nikolay Savinov, Yaroslav Ganin, Pierre H Richemond, Arnaud Doucet, Robin Strudel, Chris Dyer, Conor Durkan, et al. 2022. Continuous diffusion for categorical data. *arXiv preprint arXiv:2211.15089* (2022).
- [22] Laurent Dinh, David Krueger, and Yoshua Bengio. 2014. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516* (2014).
- [23] John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research* 12, 7 (2011).
- [24] Gabriele Fanelli, Juergen Gall, Harald Romsdorfer, Thibaut Weise, and Luc Van Gool. 2010. A 3-d audio-visual corpus of affective communication. *IEEE Transactions on Multimedia* 12, 6 (2010), 591–598.
- [25] Yao Feng, Fan Wu, Xiaohu Shao, Yanfeng Wang, and Xi Zhou. 2018. Joint 3d face reconstruction and dense alignment with position map regression network. In *Proceedings of the European conference on computer vision (ECCV)*. 534–551.
- [26] Claudio Ferrari, Giuseppe Lisanti, Stefano Berretti, and Alberto Del Bimbo. 2017. A dictionary learning-based 3D morphable shape model. *IEEE Transactions on Multimedia* 19, 12 (2017), 2666–2679.
- [27] Shansan Gong, Mukai Li, Jiantao Feng, Zhiyong Wu, and LingPeng Kong. 2022. Diffuseq: Sequence to sequence text generation with diffusion models. *arXiv preprint arXiv:2210.08933* (2022).
- [28] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.
- [29] Alexandros Graikos, Nikolay Malkin, Nebojsa Jojic, and Dimitris Samaras. 2022. Diffusion models as plug-and-play priors. *arXiv preprint arXiv:2206.09012* (2022).
- [30] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. 2020. Action2motion: Conditioned generation of 3d human motions. In *Proc. ACM Multimedia*. 2021–2029.
- [31] Jianzhu Guo, Xiangyu Zhu, Yang Yang, Fan Yang, Zhen Lei, and Stan Z Li. 2020. Towards fast, accurate and stable 3d dense face alignment. In *European Conference on Computer Vision*. Springer, 152–168.
- [32] William Harvey, Saeid Naderiparizi, Vaden Masrani, Christian Weilbach, and Frank Wood. 2022. Flexible diffusion modeling of long videos. *arXiv preprint arXiv:2205.11495* (2022).
- [33] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. 2022. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626* (2022).
- [34] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* 30 (2017).

- [35] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. 2022. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303* (2022).
- [36] Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* 33 (2020), 6840–6851.
- [37] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. 2022. Cascaded Diffusion Models for High Fidelity Image Generation. *J. Mach. Learn. Res.* 23, 47 (2022), 1–33.
- [38] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. 2022. Video diffusion models. *arXiv preprint arXiv:2204.03458* (2022).
- [39] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation* 9, 8 (11 1997), 1735–1780. <https://doi.org/10.1162/neco.1997.9.8.1735> arXiv:https://direct.mit.edu/neco/article-pdf/9/8/1735/813796/neco.1997.9.8.1735.pdf
- [40] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. 2013. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence* 36, 7 (2013), 1325–1339.
- [41] Zi-Hang Jiang, Qianyi Wu, Keyu Chen, and Juyong Zhang. 2019. Disentangled representation learning for 3d face shape. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 11957–11966.
- [42] Hyeonwoo Kim, Pablo Garrido, Ayush Tewari, Weipeng Xu, Justus Thies, Matthias Nießner, Patrick Pérez, Christian Richardt, Michael Zollöfer, and Christian Theobalt. 2018. Deep Video Portraits. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 163.
- [43] Jihoon Kim, Jiseob Kim, and Sungjoon Choi. 2023. FLAME: free-form language-based motion synthesis & editing. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence (AAAI'23/IAAI'23/EAAI'23)*. AAAI Press, Article 927, 9 pages. <https://doi.org/10.1609/aaai.v37i7.25996>
- [44] Diederik P Kingma, Shakir Mohamed, Danilo Jimenez Rezende, and Max Welling. 2014. Semi-supervised learning with deep generative models. In *Advances in neural information processing systems*. 3581–3589.
- [45] Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*. <http://arxiv.org/abs/1312.6114>
- [46] Zhifeng Kong, Wei Ping, Jiaji Huang, Kexin Zhao, and Bryan Catanzaro. 2020. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761* (2020).
- [47] Buyu Li, Yongchi Zhao, Shi Zhelun, and Lu Sheng. 2022. Danceformer: Music conditioned 3d dance generation with parametric motion transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 1272–1279.
- [48] Haoying Li, Yifan Yang, Meng Chang, Shiqi Chen, Huajun Feng, Zhihai Xu, Qi Li, and Yueting Chen. 2022. Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing* 479 (2022), 47–59.
- [49] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. 2021. Ai choreographer: Music conditioned 3d dance generation with aist++. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 13401–13412.
- [50] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. 2017. Learning a model of facial shape and expression from 4D scans. *ACM Trans. Graph.* 36, 6 (2017), 194–1.
- [51] Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori B Hashimoto. 2022. Diffusion-LM Improves Controllable Text Generation. *arXiv preprint arXiv:2205.14217* (2022).
- [52] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. 2015. SMPL: A skinned multi-person linear model. *ACM transactions on graphics (TOG)* 34, 6 (2015), 1–16.
- [53] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. 2022. RePaint: Inpainting using Denoising Diffusion Probabilistic Models. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 11451–11461.
- [54] Shitong Luo and Wei Hu. 2021. Diffusion probabilistic models for 3d point cloud generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2837–2845.
- [55] Zhaoyang Lyu, Zhifeng Kong, Xudong Xu, Liang Pan, and Dahua Lin. 2021. A conditional point diffusion-refinement paradigm for 3d point cloud completion. *arXiv preprint arXiv:2112.03530* (2021).
- [56] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. 2019. AMASS: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF international conference on computer vision*. 5442–5451.
- [57] Mehdi Mirza and Simon Osindero. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014).
- [58] Anh Nguyen, Jeff Clune, Yoshua Bengio, Alexey Dosovitskiy, and Jason Yosinski. 2017. Plug & play generative networks: Conditional iterative generation of images in latent space. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4467–4477.

- [59] Alexander Quinn Nichol and Prafulla Dhariwal. 2021. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*. PMLR, 8162–8171.
- [60] Jun-Yong Noh and Douglas Fidele. 2000. Animated Deformations with Radial Basis Functions. In *In ACM Virtual Reality and Software Technology (VRST)*. 166–174.
- [61] Naima Otberdout, Mohammed Daoudi, Anis Kacem, Lahoucine Ballihi, and Stefano Berretti. 2020. Dynamic facial expression generation on hilbert hypersphere with conditional wasserstein generative adversarial nets. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2020).
- [62] Naima Otberdout, Claudio Ferrari, Mohamed Daoudi, Stefano Berretti, and Alberto Del Bimbo. 2022. Sparse to Dense Dynamic 3D Facial Expression Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 20385–20394.
- [63] Kushagra Pandey, Avideep Mukherjee, Piyush Rai, and Abhishek Kumar. 2022. Diffusevae: Efficient, controllable and high-fidelity generation from low-dimensional latents. *arXiv preprint arXiv:2201.00308* (2022).
- [64] Mathis Petrovich, Michael J Black, and Gül Varol. 2021. Action-conditioned 3d human motion synthesis with transformer vae. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 10985–10995.
- [65] F. Pighin, R. Szeliski, and D. H. Salesin. 1999. Resynthesizing facial animation through 3D model-based tracking. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, Vol. 1. 143–150 vol.1. <https://doi.org/10.1109/ICCV.1999.791210>
- [66] Matthias Plappert, Christian Mandery, and Tamim Asfour. 2016. The KIT motion-language dataset. *Big data* 4, 4 (2016), 236–252.
- [67] Rolandos Alexandros Potamias, Jiali Zheng, Stylianos Ploumpis, Giorgos Bouritsas, Evangelos Ververas, and Stefanos Zafeiriou. 2020. Learning to Generate Customized Dynamic 3D Facial Expressions. In *Computer Vision – ECCV 2020: 16th European Conference* (Glasgow, United Kingdom). Springer-Verlag, 278–294. https://doi.org/10.1007/978-3-030-58526-6_17
- [68] Konpat Preechakul, Nattanat Chatthee, Suttisak Wizadwongsa, and Supasorn Suwajanakorn. 2022. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10619–10629.
- [69] Albert Pumarola, Antonio Agudo, Aleix M Martinez, Alberto Sanfeliu, and Francesc Moreno-Noguer. 2018. Ganimation: Anatomically-aware facial animation from a single image. In *Proc. European conference on computer vision*. 818–833.
- [70] Abhinanda R Punnakal, Arjun Chandrasekaran, Nikos Athanasiou, Alejandra Quiros-Ramirez, and Michael J Black. 2021. BABEL: Bodies, action and behavior with english labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 722–731.
- [71] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*. PMLR, 8748–8763.
- [72] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125* (2022).
- [73] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J Black. 2018. Generating 3D faces using convolutional mesh autoencoders. In *Proceedings of the European conference on computer vision (ECCV)*. 704–720.
- [74] Danilo Rezende and Shakir Mohamed. 2015. Variational inference with normalizing flows. In *International conference on machine learning*. PMLR, 1530–1538.
- [75] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10684–10695.
- [76] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*. Springer, 234–241.
- [77] Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. 2022. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH 2022 Conference Proceedings*. 1–10.
- [78] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. 2022. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. *arXiv preprint arXiv:2205.11487* (2022).
- [79] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. 2022. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022).
- [80] Hyewon Seo and Guoliang Luo. 2021. Generating 3D Facial Expressions with Recurrent Neural Networks. In *Intelligent Scene Modeling and Human-Computer Interaction*. Springer International Publishing, 181–196. https://doi.org/10.1007/978-3-030-71002-6_11

- [81] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. 2020. Interpreting the Latent Space of GANs for Semantic Face Editing. In *CVPR*.
- [82] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*. PMLR, 2256–2265.
- [83] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. 2015. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems* 28 (2015), 3483–3491.
- [84] Jiaming Song, Chenlin Meng, and Stefano Ermon. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020).
- [85] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456* (2020).
- [86] Guy Tevet, Brian Gordon, Amir Hertz, Amit H Bermano, and Daniel Cohen-Or. 2022. MotionCLIP: Exposing Human Motion Generation to CLIP Space. *arXiv preprint arXiv:2203.08063* (2022).
- [87] Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. 2023. Human Motion Diffusion Model. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=SJ1kSyO2jwu>
- [88] Xiaoguang Tu, Jian Zhao, Mei Xie, Zihang Jiang, Akshaya Balamurugan, Yao Luo, Yang Zhao, Lingxiao He, Zheng Ma, and Jiashi Feng. 2020. 3D face reconstruction from a single image assisted by 2D face images in the wild. *IEEE Transactions on Multimedia* 23 (2020), 1160–1172.
- [89] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. 2018. MoCoGAN: Decomposing Motion and Content for Video Generation. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1526–1535. <https://doi.org/10.1109/CVPR.2018.00165>
- [90] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*. 5998–6008.
- [91] Daniel Vlasic, Matthew Brand, Hanspeter Pfister, and Jovan Popovic. 2006. Face Transfer with Multilinear Models. *ACM Transactions on Graphics* 24 (07 2006). <https://doi.org/10.1145/1185657.1185864>
- [92] Bram Wallace, Akash Gokul, and Nikhil Naik. 2022. EDICT: Exact Diffusion Inversion via Coupled Transformations. *arXiv preprint arXiv:2211.12446* (2022).
- [93] Wei Wang, Xavier Alameda-Pineda, Dan Xu, Pascal Fua, Elisa Ricci, and Nicu Sebe. 2018. Every smile is unique: Landmark-guided diverse smile generation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*. 7083–7092.
- [94] Yaohui Wang, Piotr Bilinski, Francois Bremond, and Antitza Dantcheva. 2020. G3AN: Disentangling Appearance and Motion for Video Generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [95] Yaohui Wang, Piotr Bilinski, Francois Bremond, and Antitza Dantcheva. 2020. Imaginator: Conditional spatio-temporal gan for video generation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 1160–1169.
- [96] Zhendong Wang, Huangjie Zheng, Pengcheng He, Weizhu Chen, and Mingyuan Zhou. 2022. Diffusion-gan: Training gans with diffusion. *arXiv preprint arXiv:2206.02262* (2022).
- [97] Y. Yan, Ying Huang, Si Chen, Chunhua Shen, and Hanzi Wang. 2020. Joint Deep Learning of Facial Expression Synthesis and Recognition. *IEEE Transactions on Multimedia* 22 (2020), 2792–2807. <https://api.semanticscholar.org/CorpusID:211044056>
- [98] Ruihan Yang, Prakhar Srivastava, and Stephan Mandt. 2022. Diffusion probabilistic modeling for video generation. *arXiv preprint arXiv:2203.09481* (2022).
- [99] Juyong Zhang, Keyu Chen, and Jianmin Zheng. 2020. Facial Expression Retargeting From Human to Avatar Made Easy. *IEEE Transactions on Visualization and Computer Graphics* 28 (2020), 1274–1287. <https://api.semanticscholar.org/CorpusID:221103938>
- [100] M. Zhang, Z. Cai, L. Pan, F. Hong, X. Guo, L. Yang, and Z. Liu. 5555. MotionDiffuse: Text-Driven Human Motion Generation With Diffusion Model. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 01 (jan 5555), 1–15. <https://doi.org/10.1109/TPAMI.2024.3355414>
- [101] Xing Zhang, Lijun Yin, Jeffrey F Cohn, Shaun Canavan, Michael Reale, Andy Horowitz, and Peng Liu. 2013. A high-resolution spontaneous 3d dynamic facial expression database. In *IEEE workshops on automatic face and gesture recognition*. 1–6.
- [102] Linqi Zhou, Yilun Du, and Jiajun Wu. 2021. 3d shape generation and completion through point-voxel diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 5826–5835.

A TRAINING WITH AND GENERATION OF SEQUENCES OF ARBITRARY LENGTH

Since our noise approximator is a bidirectional transformer, it can take sequences of arbitrary length as input —it can be trained using sequences of any length, and we can sample from the resulting model to obtain sequences of desired lengths (The length of x_0 will be that of x_T). In the same way, as a bidirectional transformer is used also to guide the reverse process, it can guide the reverse process with any length for x_T . Consequently, tasks related to label control, text control, and geometry-adaptive generation can generate sequences of any desired length. Furthermore, the sequences that have to be filled with the expression filling task can be of any length. For the sake of simplicity, we describe here only the label control task. The noise approximator and the classifier used for the guidance are either trained using sequences of a fixed length ($F = 40$) or variable lengths (F is uniformly distributed in the interval $[35, 45]$).

The performance of both models is evaluated when outputting sequences of length in $[35, 45]$. The performance is evaluated as in Sec. 5.1, except that the independent classifier is trained with sequences of variable length (F is uniformly distributed in $[35, 45]$). Results are shown in Fig.6.

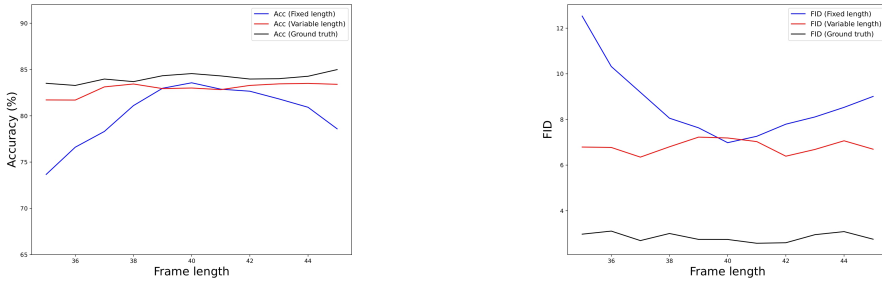


Fig. 6. Quantitative evaluation of the label control task for models trained with sequences of a fixed length ($F = 40$) or variable lengths. Performance is evaluated on generated sequences of different lengths using, as in Sec. 5.1, the classification accuracy (left) and the FID score (right).

When generating sequences of different lengths is required, training with variable lengths helps the model to perform better. Moreover, the results obtained with the model trained with sequences of variable length are satisfactory: the achieved accuracy is similar to that of the ground truth. Moreover, the FID obtained for a length frame of 40 is similar to that calculated with the model dedicated to output sequences of length 40.

B DIVERSITY OF THE GENERATED SEQUENCES WHEN CONDITIONING ON EXPRESSION LABEL

We study in this section the diversity of the generated sequences both in terms of facial anatomy (L) and in terms of expression (ΔL_f) in the label control task. As a reminder, the 3D arrangement of a landmark frame L_f can be regarded as the combination of the facial anatomy (at a neutral pose L) and the expression-driven shape change applied to it, i.e. $L_f = \Delta L_f + L$.

Since the proposed landmark-guided mesh deformation retargets the expression $\Delta L_f = L_f - L$ onto a new face anatomy given as a mesh M , it is used hereafter to illustrate the diversity of the generated expressions but it is not adapted to analyze the facial anatomy of the generated L . To

show the diversity of facial anatomy generated by our model, we use the FLAME model[50] to compute the facial mesh from the landmark set of neutral pose¹.

Fig. 7 presents three illustrative neutral faces L that we generated by conditioning the reverse process on the same expression label “mouth open”. Both landmark set and the FLAME-fitted mesh are shown, for each face. (The neutral face L associated with a generated sequence x_0 is set to either L_1 or L_F , depending on the sequence type.) Additionally, the diversity in the generated expression is illustrated in Fig. 8. The apparent distinction among these results demonstrate that the proposed approach is able to generate sequences of rich diversity, both in terms of facial anatomy and expression (This is due to the input noise x_T that is sampled from $\mathcal{N}(0, I)$).

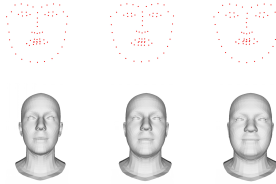


Fig. 7. Diversity of facial anatomy in the generated expressions. We use FLAME model to compute facial meshes from the landmark sets, for the visualization purpose.




Fig. 8. Diversity of expressions generated with the label “mouth side” (left), and “mouth open” (right) in the label control task. Note that generated sequences can be either of type E2N or N2E.

¹We can note that the meshes generated from FLAME lack certain details of the facial geometry, resulting in dull, lifeless shapes. Furthermore, FLAME takes about 470s to fit one sequence, while the proposed landmark-guided mesh deformation needs only about 1.30s.

C PSEUDO CODE FOR EACH DOWNSTREAM TASK

Algorithm 1 Label control

Input: Label y .
Output: Sequence x_0 (corresponding to label y).

- 1: $x_T \sim N(0, I)$
- 2: **for** $t = T, \dots, 1$ **do**
- 3: ▷ Estimation of $p_\theta(\cdot | x_t)$
- 4: Compute $\epsilon_\theta(x_t, t)$
- 5: Compute $\mu_\theta(x_t, t): \mu_\theta(x_t, t) = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{1-\alpha_t}} \epsilon_\theta(x_t, t) \right)$
- 6: ▷ Sampling from $p_\theta(\cdot | x_t)$
- 7: $z \sim N(0, I)$ if $t > 1, 0$ otherwise
- 8: Set \hat{x}_{t-1} to $\mu_\theta(x_t, t) + \sigma_t z$
- 9: ▷ Optimization: optimization procedure is initialized with \hat{x}_{t-1}
- 10: $x_{t-1} = \underset{x}{\operatorname{argmax}} \left[\lambda \log(p_\theta(x | x_t)) + \log(p_\phi(y | x, t-1)) \right]$

return x_0

Algorithm 2 Text control

Input: Text c .
Output: Sequence x_0 (corresponding to text c).

- 1: $x_T \sim N(0, I)$
- 2: **for** $t = T, \dots, 1$ **do**
- 3: ▷ Estimation of $p_\theta(\cdot | x_t)$
- 4: Compute $\epsilon_\theta(x_t, t)$
- 5: Compute $\mu_\theta(x_t, t): \mu_\theta(x_t, t) = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{1-\alpha_t}} \epsilon_\theta(x_t, t) \right)$
- 6: ▷ Sampling from $p_\theta(\cdot | x_t)$
- 7: $z \sim N(0, I)$ if $t > 1, 0$ otherwise
- 8: Set \hat{x}_{t-1} to $\mu_\theta(x_t, t) + \sigma_t z$
- 9: ▷ Optimization: optimization procedure is initialized with \hat{x}_{t-1}
- 10: $x_{t-1} = \underset{x}{\operatorname{argmax}} \left[\lambda \log(p_\theta(x | x_t)) + \cos(\operatorname{BiT}(x, t-1), \operatorname{CLIP}(c)) \right]$

return x_0

Algorithm 3 Sequence filling

Input: Partial sequence $x_0|_{S_K}$
Output: Completed sequence x_0

- 1: $x_T|_{S_U} \sim N(0, I)$
- 2: $x_T|_{S_K} = \sqrt{\alpha_T} x_0|_{S_K} + \sqrt{1-\alpha_T} \epsilon, \epsilon \sim N(0, I)$
- 3: **for** $t = T, \dots, 1$ **do**
- 4: ▷ Estimation of $p_\theta(\cdot | x_t)$
- 5: Compute $\epsilon_\theta(x_t, t)$
- 6: Compute $\mu_\theta(x_t, t): \mu_\theta(x_t, t) = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{\beta_t}{\sqrt{1-\alpha_t}} \epsilon_\theta(x_t, t) \right)$
- 7: ▷ Sampling from $p_\theta(\cdot | x_t)$
- 8: $z \sim N(0, I)$ if $t > 1, 0$ otherwise
- 9: Set \hat{x}_{t-1} to $\mu_\theta(x_t, t) + \sigma_t z$
- 10: $x_{t-1}|_{S_U} = \hat{x}_{t-1}|_{S_U}$
- 11: **if** $t > 1$ **then** ▷ if $t = 1, x_0|_{S_K}$ is already properly set.
- 12: $x_{t-1}|_{S_K} = \sqrt{\alpha_{t-1}} x_0|_{S_K} + \sqrt{1-\alpha_{t-1}} \epsilon, \epsilon \sim N(0, I)$

return x_0

Algorithm 4 Geometry-adaptive generation with label control

Input: Label y and partial sequence $x_0|_{S_K}$. S_K is either $\{1\}$ or $\{F\}$ and the unique frame associated with $x_0|_{S_K}$ is a neutral one.

Output: Completed sequence x_0 (corresponding to label y)

```

1: for  $i = 1$  to  $5$  do
2:   if  $i == 1$  then
3:      $x_T|_{S_U} \sim N(0, I)$ 
4:      $x_T|_{S_K} = \sqrt{\bar{\alpha}_T}x_0|_{S_K} + \sqrt{1 - \bar{\alpha}_T}\epsilon, \epsilon \sim N(0, I)$ 
5:   else
6:      $x_T = \sqrt{\bar{\alpha}_T}x_0 + \sqrt{1 - \bar{\alpha}_T}\epsilon, \epsilon \sim N(0, I)$ 
7:   for  $t = T, \dots, 1$  do
8:     Compute  $\epsilon_\theta(x_t, t)$  ▷ Estimation of  $p_\theta(\cdot|x_t)$ 
9:     Compute  $\mu_\theta(x_t, t): \mu_\theta(x_t, t) = \frac{1}{\sqrt{\bar{\alpha}_t}} \left( x_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right)$ 
10:     $z \sim N(0, I)$  if  $t > 1, 0$  otherwise ▷ Sampling from  $p_\theta(\cdot|x_t)$ 
11:    Set  $\hat{x}_{t-1}$  to  $\mu_\theta(x_t, t) + \sigma_t z$ 
12:     $\hat{x}_{t-1} = \operatorname{argmax}_x \left[ \lambda \log(p_\theta(x|x_t)) + \log(p_\phi(y|x, t-1)) \right]$  ▷ Optimization: optimization procedure is initialized with  $\hat{x}_{t-1}$ 
13:     $x_{t-1}|_{S_U} = \hat{x}_{t-1}|_{S_U}$ 
14:    if  $t > 1$  then ▷ if  $t = 1, x_0|_{S_K}$  is already properly set.
15:       $x_{t-1}|_{S_K} = \sqrt{\bar{\alpha}_{t-1}}x_0|_{S_K} + \sqrt{1 - \bar{\alpha}_{t-1}}\epsilon, \epsilon \sim N(0, I)$ 
16: return  $x_0$ 

```
