

LiP-Flow: Learning Inference-time Priors for Codec Avatars via Normalizing Flows in Latent Space

Emre Aksan^{§1}, Shugao Ma³, Akin Caliskan^{§2}, Stanislav Pidhorskyi³,
Alexander Richard³, Shih-En Wei³, Jason Saragih³, and Otmar Hilliges¹

¹ ETH Zürich, Department of Computer Science

² CVSSP, University of Surrey

³ Meta Reality Labs Research

Abstract. Neural face avatars that are trained from multi-view data captured in camera domes can produce photo-realistic 3D reconstructions. However, at inference time, they must be driven by limited inputs such as partial views recorded by headset-mounted cameras or a front-facing camera, and sparse facial landmarks. To mitigate this asymmetry, we introduce a prior model that is conditioned on the runtime inputs and tie this prior space to the 3D face model via a normalizing flow in the latent space. Our proposed model, LiP-Flow, consists of two encoders that learn representations from the rich training-time and impoverished inference-time observations. A normalizing flow bridges the two representation spaces and transforms latent samples from one domain to another, allowing us to define a latent likelihood objective. We trained our model end-to-end to maximize the similarity of both representation spaces and the reconstruction quality, making the 3D face model aware of the limited driving signals. We conduct extensive evaluations where the latent codes are optimized to reconstruct 3D avatars from partial or sparse observations. We show that our approach leads to an expressive and effective prior, capturing facial dynamics and subtle expressions better. Check out our [video](#) for an overview.

1 Introduction

VR telepresence promises immersive social interactions. To experience the presence of others as genuine, such a system must provide photo-realistic 3D renderings, capture fine-grained details such as pores and hair, and produce subtle facial dynamics. Codec Avatars [32,47] have been shown to be a promising direction towards this goal. Such systems implement 3D avatars via the decoder network of a variational auto-encoder (VAE) [26]. Per subject, high-quality 3D face avatars are trained with multi-view imagery from a 40+ camera dome [32].

In prior work [8,16,32,33,40], the network learns to decode latent codes into shape and view-conditioned appearance. Importantly, each latent code is a projection of the geometry and unwrapped texture which are only available during *training* (Fig. 1-a). At *inference-time*, the driving signal is impoverished – consisting of a set of partial face images captured from cameras mounted on a VR

[§] This work was performed during an internship at Meta Reality Labs Research.

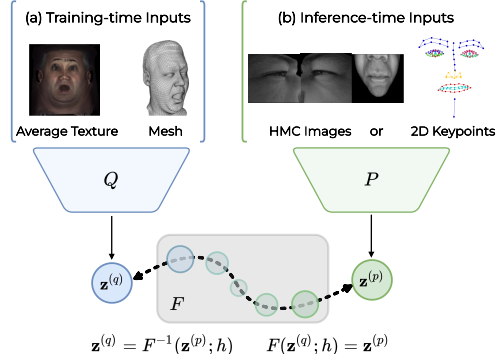


Fig. 1. Information asymmetry between the (a) training- and (b) inference-time observations. We learn a normalizing flow model (F) across latent spaces, minimizing the discrepancy between the base network’s Deep Appearance Model (DAM) encoder [32], DAM-encoder (Q), and an inference-time encoder (P), namely the HMC-encoder or KPT-encoder conditioned on partial head-mounted camera (HMC) images or sparse 2D keypoints (KPT), respectively.

headset [47], facial landmarks [13] (Fig. 1-b) or a single-view from a frontal camera [12]. To maximize the reconstruction quality, the avatar model is often trained in isolation, and typically a separate run-time encoder for the driving signal is used. We show that this separate training scheme limits performance since the constrained run-time conditions are not considered during decoder training.

In this work, we introduce LiP-Flow, a learned inference-time prior, based on normalizing flows, to bridge this information gap for 3D face models. Specifically, we incorporate a prior model into the Deep Appearance Model (DAM) pipeline [32] with minimal modifications such that the decoder’s high reconstruction quality is preserved. We make use of this prior model in *inference-time* optimization tasks. This allows us to optimize latent codes to capture facial dynamics and subtle expressions even when only impoverished inputs are available (Fig. 2). It is effective even on less challenging neutral expressions, while it is more pronounced on peak expressions. We simulate the run-time conditions via sparse 2D keypoints (KPT) and synthetically generated partial head-mounted camera (HMC) images, allowing to define (virtual) correspondences between training and inference data and to run tightly controlled evaluations in all settings.

Data-driven priors have many uses in under-constrained problems. Given partial or imperfect observations, the task is the reconstruction of a clean signal or inference of the latent factors leading to it. This task has been studied in the context of related areas, amongst others 3D lifting [10,36], 3D pose and shape estimation from images [14,37,44,52], and refining [31], image editing [6,35], or prediction of the next step from the past in sequence models [2,3,16,29,37]. Our focus lies particularly on learning latent priors conditioned on limited observations [2,17,37] and using prior models for inference in a more compact representation space [14,28,36,37]. In our work, we learn a prior that is conditioned on the HMC views or facial landmarks. We aim to regularize the main avatar model by incorporating the run-time observations into the training pipeline and increase reconstruction quality in inference-time optimization.

Our work introduces a novel approach for learning of conditional priors. To minimize the discrepancy between the training- and inference-time representations, we learn a mapping across the two latent spaces via a conditional normalizing flow (Fig. 1) [49,7]. It is noteworthy to mention that both the training- and




Fig. 2. Driving a 3D avatar with impoverished inputs. For the DAM (c) and our proposed model LiP-Flow (e), we first optimize latent codes to reconstruct the inference-time observations (b), and then reconstruct the full face. In (d), we train a regressor to directly predict latent codes from the inference-time inputs (b) by using the pretrained DAM-encoder (Q) in (c).

inference-time encoders, our latent flow network as well as the decoder are jointly trained. Our training objective consists of reconstruction terms for appearance and geometry, and a latent likelihood term, replacing the KL-divergence (KL-D) term in the VAE framework. Specifically, we optimize the likelihood of the latent samples from the DAM-encoder (Q) under the learned inference-time prior (LiP) distribution estimated from the HMC views or 2D keypoints. This is enabled by the bijective nature of the underlying flow, allowing us to map samples from one space to another. We note that the information asymmetry is pronounced in our task both in terms of modality and amount of carried information (*i.e.*, geometry and texture vs. HMC images or keypoints). The KL-D objective enforces distribution similarity between the latent-spaces which we show experimentally may be a too strong of an assumption, given the differences in input distributions. We relax this assumption by leveraging a normalizing flow to bridge between the latent spaces without enforcing distributional similarity. We show that our approach constitutes a powerful means to deal with such large discrepancies.

We first demonstrate how our latent space formulation improves the reconstruction performance of the base DAM network. Next, we compare different approaches for learning a prior model conditioned on the HMC images or 2D keypoints. We evaluate the quality of the prior models via an analysis-by-synthesis technique [14, 34]. More specifically, we aim to recover the underlying latent code from HMC images or facial landmarks. We treat the pre-trained decoder (*i.e.*, the 3D face model) as a black box renderer to produce an image for a given latent code, view direction, and camera pose. In this task, in addition to a reconstruction objective, we also leverage the learned prior via a likelihood objective on the latent codes. We show that our flow-based prior model yields latent codes that reconstruct the 3D face faithfully (Fig. 2). Our contributions can be summarized as follows: (1) a novel approach for learning conditional priors that can be integrated seamlessly into a base architecture for the task of 3D face avatar generation and animation, (2) a more expressive and flexible representation space compared to KL-D driven latent spaces, (3) better reconstruction quality than the baseline approaches in tightly controlled extensive evaluations.

2 Related work

3D Face Appearance Models. Modeling of 3D human faces for animation has long been an active area of research. Early work models 3D human faces with linear combinations of blendshape vectors representing meshes and texture (*e.g.* [9]). We refer to Lewis *et al.* [30] for a review. Such models are limited in their expressiveness: a large number of blendshapes are needed to create high-fidelity animations. Recent works [11,20] leverages 3D parametric face models to predict texture and shape for novel identities. Many deep appearance face models utilize deep generative neural networks such as VAEs [5,13,32,33,45] or GANs [1,15,41,42] with facial expressions represented as compact *latent codes*.

In particular *Codec Avatars* [32,33] trained with high resolution multi-view images can photo-realistically render human faces. Face animations can also be generated when conditioned on HMC images [16,47], frontal-view images [12] or even audio [39]. However, the significant discrepancy between training- and inference-time inputs is a key challenge for achieving high fidelity facial animation for telepresence applications. To alleviate this problem, we explore novel approaches augmenting the Codec Avatar pipeline via conditional priors..

Latent Code Optimization. Generative models like GANs [22] and VAEs [26] can represent the underlying data manifold in a latent space. They have been used as data-driven priors in solving inverse problems, via optimization in the more abstract latent space. The standard Gaussian latent prior of a VAE [36] and a hierarchical VAE [31] replace a GMM prior in the output space [10] for modeling and reconstructing 3D human pose. Chen *et al.* [14] also use a VAE decoder and regularize the latent codes explicitly by evaluating the latent likelihood under the VAE prior. In [6,35], GANs and in [23], denoising autoencoders are treated as a block-box sample generator for image in-painting and denoising tasks. Similarly, GANFit [21] models facial texture via a GAN, and jointly optimizes the corresponding latent code and parameters of a 3D morphable model. Our work differs from those in that we learn a conditional prior model allowing us to calculate the likelihood of the latent code during optimization. HuMoR [37] also provides a learned prior that affords likelihood evaluation. It is conditioned on the 3D pose in the previous frame and regularized via KL-D. Instead, we propose a flow-based approach for learning conditional priors in settings where the inference-time data differs significantly from the training data.

Normalizing flow priors have been proposed for inverse problems [4,48,51]. In [4], an invertible Glow model [25] replaces GANs for image inpainting. Similarly in [28], a normalizing flow model is used for 2D-3D lifting. Zanfir *et al.* [51] introduce a 3D human pose prior by learning a mapping between a Gaussian latent space and 3D human poses via normalizing flows. These works make use of the generative nature of unconditional flows and replace GANs or VAEs 1-to-1. The LSGM [46] learns an invertible mapping between a standard Normal prior and an encoder space via a diffusion process. In our work, we introduce a different setting where we jointly learn separate representation spaces with a large discrepancy in the inputs and the mapping between them.

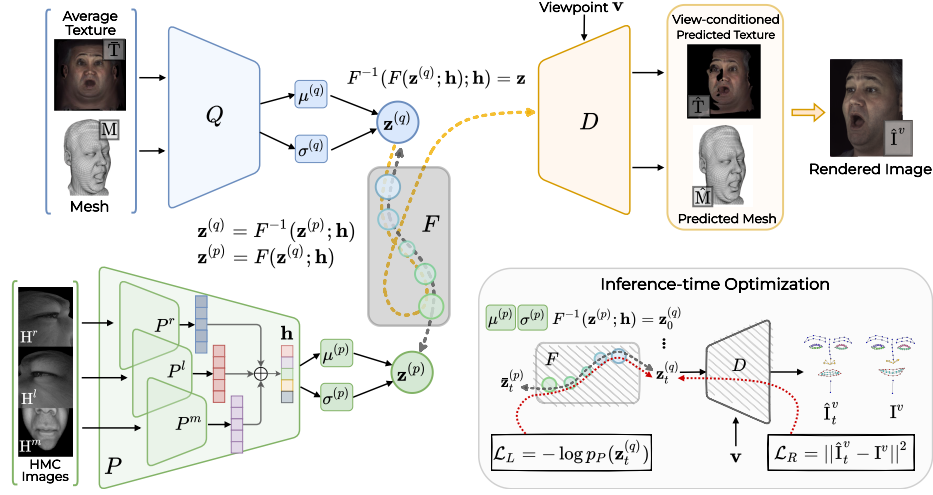


Fig. 3. LiP-Flow with HMC inputs. Without losing generality, we introduce an **HMC-encoder** (P), treated as a prior model conditioned on the HMC images, which can be replaced with a **KPT-encoder** (P) conditioned on 2D keypoints. The prior model (P) is trained end-to-end with the base **DAM-encoder** (Q) and the **decoder** (D). The bijective function (F) learns to map samples from one space to another. We train the decoder (D) with latent samples $\mathbf{z}^{(q)}$ only. To get the reconstruction training signal in the flow (F), the input code follows the flow cycle $\mathbf{z} = F^{-1}(F(\mathbf{z}^{(q)}; \mathbf{h}); \mathbf{h})$. **Inference-time Optimization.** To faithfully reconstruct the 3D face only from runtime inputs, we optimize the latent codes to fit the observations such as HMC images or 2D keypoints. We first estimate a prior distribution (P) with the inference-time encoders **HMC-encoder** or **KPT-encoder**. We then use this prior to initialize the latent code \mathbf{z}_0 and evaluate the likelihood of the latent code \mathbf{z}_t where t corresponds to the optimization step. The red dashed lines denote the gradients of the respective objective terms, namely the latent likelihood \mathcal{L}_L (Eq. 9) and the reconstruction loss \mathcal{L}_R .

3 Method

Our work extends the Deep Appearance Model (DAM) [32] by introducing a runtime encoder (LiP-encoder) that is conditioned on the impoverished inference-time inputs. We aim to learn models that are capable of driving avatars via partial HMC images or sparse facial landmarks. Our paper showcases two modalities for different yet related tasks, including but not limited to HMC- and KPT-encoder taking HMC views (see Fig. 3) and 2D keypoints, respectively. In the remainder of the paper, the LiP-encoder term refers to either of the HMC- or KPT-encoder. We treat a LiP-encoder as a conditional prior and train it end-to-end together with the DAM network via our flow-based latent space formulation.

In the following sections, we first explain the problem setup and provide an overview of the base 3D face model, DAM [32], in Sec. 3.1. We then present our approach, LiP-Flow, in Sec. 3.3 along with the HMC- and KPT-encoder networks in Sec. 3.2. We also describe alternative techniques to drive the avatar with inference-time inputs in Sec. 4. Architecture details are in the supp. mat.

3.1 Background: Deep Appearance Model

We follow the data preprocessing steps described in [32] and use the same data, and training routines for all the models. Specifically, frames are unwrapped into a texture, using the tracked geometry. The DAM-encoder (Q) (Fig. 4-a) takes the geometry $\mathbf{M} \in \mathbb{R}^{7306 \times 3}$ and the average texture across all the views $\bar{\mathbf{T}} \in \mathbb{R}^{3 \times 1024 \times 1024}$, and parameterizes a Normal distribution of latent codes $\mathbf{z} \in \mathbb{R}^{256}$:

$$\boldsymbol{\mu}^{(q)}, \boldsymbol{\sigma}^{(q)} = Q(\bar{\mathbf{T}}, \mathbf{M}), \quad \mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}^{(q)}, \boldsymbol{\sigma}^{(q)}). \quad (1)$$

The decoder D decodes a latent sample \mathbf{z} for a given camera view-vector \mathbf{v} into geometry $\hat{\mathbf{M}}$ and view-specific texture $\bar{\mathbf{T}}^v$ accounting for view-dependent effects by $\hat{\mathbf{M}}, \bar{\mathbf{T}}^v = D(\mathbf{z}, \mathbf{v})$. The reconstructed image is rendered with the decoded texture, mesh, and camera pose information. The training objective consists of reconstruction and latent space regularization terms with separate weights:

$$\mathcal{L} = \sum_v \lambda_I \mathcal{L}_I + \lambda_M \mathcal{L}_M + \lambda_L \mathcal{L}_L, \quad (2)$$

$$\mathcal{L}_I = \left\| \left(\mathbf{I}^v - \hat{\mathbf{I}}^v \right) \odot \mathbf{W}^v \right\|^2, \quad \mathcal{L}_M = \left\| \mathbf{M} - \hat{\mathbf{M}} \right\|^2, \quad (3)$$

$$\mathcal{L}_L = D_{KL} \left(\mathcal{N}(\boldsymbol{\mu}^{(q)}, \boldsymbol{\sigma}^{(q)}) \parallel \mathcal{N}(\mathbf{0}, \mathbf{I}) \right), \quad (4)$$

where \mathcal{L}_I is the image reconstruction loss between the rendered $\hat{\mathbf{I}}^v$ and the ground-truth image $\mathbf{I}^v \in \mathbb{R}^{3 \times 1334 \times 2048}$, given view \mathbf{v} . The view-dependent image masks \mathbf{W}^v removes the background. We apply additional supervision for the predicted mesh via a geometry reconstruction loss \mathcal{L}_M . DAM uses a standard Gaussian prior (Fig. 4-a) where its latent space is regularized via KL-D (\mathcal{L}_L). This formulation already achieves high-quality reconstructions during training. However, we show that the prior $\mathcal{N}(\mathbf{0}, \mathbf{I})$ is not effective enough when recovering underlying latent codes only from the impoverished observations (see Fig. 2-c). for the latent code fitting task with only partial observations (see Fig. 2-c).

3.2 Inference-time Encoders (LiP-encoder)

We define the HMC- and the KPT-encoders conditioned on the HMC images and the 2D keypoints, respectively. Note that these networks are also used in the baselines illustrated in Fig. 4-b and Fig. 4-c.

HMC-encoder. The set of HMC views at every frame consists of partial and non-overlapping images of mouth \mathbf{H}^m , left eye \mathbf{H}^l and right eye \mathbf{H}^r regions (Fig. 5) where $\mathbf{H}^* \in \mathbb{R}^{480 \times 640}$. The HMC-encoder consists of small networks for each view which are then fused via a 3-channel attention operation:

$$\mathbf{h}^* = P^*(\mathbf{H}^*), \quad \mathbf{h} = \text{Attention}(\mathbf{h}^m, \mathbf{h}^l, \mathbf{h}^r; \mathbf{W}^A), \quad (5)$$

where \mathbf{h}^* denotes the representations for each HMC view. The hidden representation and the attention weights are of shape $\mathbf{h} \in \mathbb{R}^{1024}$ and $\mathbf{W}^A \in \mathbb{R}^{1024 \times 3}$, respectively. The attention weights are data-agnostic trainable parameters. The attention operation learns a weighted mixture of HMC views.




Fig. 4. An overview of latent space concepts. Note that the models are depicted with HMC inputs. The same concepts also apply to the KPT-encoder with keypoint inputs. Details are skipped for brevity. The decoder D decodes a view-specific texture and mesh for a given view vector \mathbf{v} and a latent code \mathbf{z} . (a) The Deep Appearance Model (DAM) following the conditional VAE framework with $\mathcal{N}(0, I)$ prior. (b) Replacing the DAM-encoder Q with the HMC-encoder P . The face avatar model (i.e., the decoder) is trained with the HMC images directly. (c) Learning to regress the latent codes for HMC views by using a pretrained DAM-encoder Q . (d) Training the HMC-encoder as a conditional prior model by minimizing the KL-divergence objective. (e) Our LiP-Flow introduces a normalizing flow bridging the latent space of Q and P .

KPT-encoder. 2D keypoint samples, \mathbf{K} , are in pixel coordinates normalized between 0 and 1, consisting of facial landmarks except the iris (Fig. 5) where $\mathbf{K} \in \mathbb{R}^{136 \times 2}$. We pass the keypoint samples to a ResNet [24] with fully connected layers. Our KPT-encoder maps the keypoint inputs to an over-parameterized representation $\mathbf{h} \in \mathbb{R}^{1024}$.

Latent Space. We model a probabilistic latent space and parameterize a Gaussian distribution as a function of the run-time input context \mathbf{h} :

$$\boldsymbol{\mu}^{(p)}, \boldsymbol{\sigma}^{(p)} = P(\mathbf{h}), \quad \mathbf{z}^{(p)} \sim \mathcal{N}(\boldsymbol{\mu}^{(p)}, \boldsymbol{\sigma}^{(p)}). \quad (6)$$

Without loss of generality, P encapsulates the HMC- and KPT-encoders and the layers that compute $\boldsymbol{\mu}^{(p)}$ and $\boldsymbol{\sigma}^{(p)}$.

3.3 Learning Conditional Priors

Instead of a standard Gaussian prior, we propose a learned prior, which has been shown to be more expressive [17, 43]. Here, the prior model is conditioned on the HMC images or 2D keypoints. We also study several architectures for this task, illustrated in Fig. 4 and explained in Sec. 4.

LiP-KL (Fig. 4-d). In previous studies [17, 43, 2, 37], conditional priors replace the VAE’s isotropic Gaussian prior with a learned prior. Similarly in our work, the latent objective \mathcal{L}_L in Eq. 4 takes the form of

$$\mathcal{L}_L = D_{KL} \left(\mathcal{N}(\boldsymbol{\mu}^{(q)}, \boldsymbol{\sigma}^{(q)}) \parallel \mathcal{N}(\boldsymbol{\mu}^{(p)}, \boldsymbol{\sigma}^{(p)}) \right), \quad (7)$$

where the KL-D objective D_{KL} enforces similarity between Q and P . We build a variant of our model with the KL objective (see Fig. 4-d). Ground-truth correspondences between the inputs of the LiP-encoders and the DAM-encoder allow us to learn a mapping between the latent spaces. Note that the distributional similarity is a plausible assumption in the prior works where the inputs are often

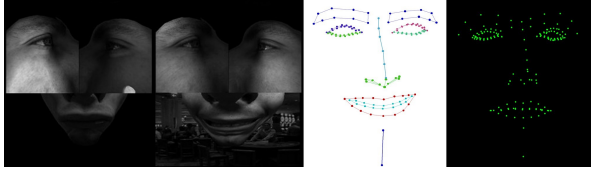


Fig. 5. Driving signals. Synthetic HMC images with lighting and background augmentations (left) and 2D key-points (right).

consecutive timesteps. However, in our case, the discrepancy between the Q and the P inputs is considerably larger than in prior work.

LiP-Flow (Fig. 4-e). To further mitigate the asymmetry of the Q and P spaces, we propose to bridge the two spaces and remove the assumption that they must be the same (*i.e.*, KL-D). To achieve this, we introduce a normalizing flow network F to transform latent samples from one space to another (Fig. 4-e). Normalizing flows are typically used to represent complex and intractable data distributions with a simple parametric distribution such as the standard Gaussian [18,38,19,25]. In our work, we use a conditional normalizing flow [18,49,7] enabling the transformation:

$$\bar{\mathbf{z}}^{(p)} = F(\mathbf{z}^{(q)}; \mathbf{h}), \quad F = f_K \circ \dots \circ f_2 \circ f_1, \quad (8)$$

where the bijection $F : Q \rightarrow P$ is a composition of K transformations f_k . For a given pair of inference- and training-time inputs, we first draw a $\mathbf{z}^{(q)}$ sample (Eq. 1) and estimate a prior distribution by using the LiP-encoder, parameterizing it as a Gaussian (Eq. 6). The $\mathbf{z}^{(q)}$ sample is then transformed to the P space corresponding to $\bar{\mathbf{z}}^{(p)}$ (Eq. 8). Finally, we calculate the log-likelihood of the latent sample $\mathbf{z}^{(q)}$ under the prior distribution P such that:

$$\log p_P(\mathbf{z}^{(q)}) = \log p_P(\bar{\mathbf{z}}^{(p)}) + \log \left(\det \left| \frac{\partial F(\mathbf{z}^{(q)}; \mathbf{h})}{\partial \mathbf{z}^{(q)}} \right| \right), \quad (9)$$

where $\log p_P(\bar{\mathbf{z}}^{(p)}) = \log \mathcal{N}(\bar{\mathbf{z}}^{(p)}; \boldsymbol{\mu}^{(p)}, \boldsymbol{\sigma}^{(p)})$ (see Eq. 6).

In LiP-Flow, the latent training objective \mathcal{L}_L then becomes the negative log-likelihood $-\log p_P(\mathbf{z}^{(q)})$ in Eq. 9, replacing the D_{KL} term. We use the flow F to project the sample $\mathbf{z}^{(q)}$ onto the P space. Hence, this objective enforces $\mathbf{z}^{(p)}$ and $\bar{\mathbf{z}}^{(p)}$ to be similar *after* the transformation of $\mathbf{z}^{(q)}$ via F (Eq. 8). However, the formulation does not restrict Q or P to any particular structure.

We would like to note that this likelihood can increase arbitrarily by contracting the P distribution, resulting in a trivial solution. This behavior is prevented by the determinant of the Jacobian term in Eq. 9, penalizing the contraction and encouraging expansion [18].

To keep the modifications to the base model at a minimum, the decoder D is trained with samples from the DAM-encoder Q only, as in the original pipeline. We do not use any samples from the LiP-encoders P . In other words, the LiP-encoders learn to complete the partial observations via the latent objective (Eq. 9). Finally, to attain the reconstruction signal in our LiP-Flow’s latent flow network F , the decoder’s latent code \mathbf{z} follows the flow cycle: $\mathbf{z} = F^{-1}(F(\mathbf{z}^{(q)}; \mathbf{h}); \mathbf{h})$. We provide ablations in Sec. 4.4.

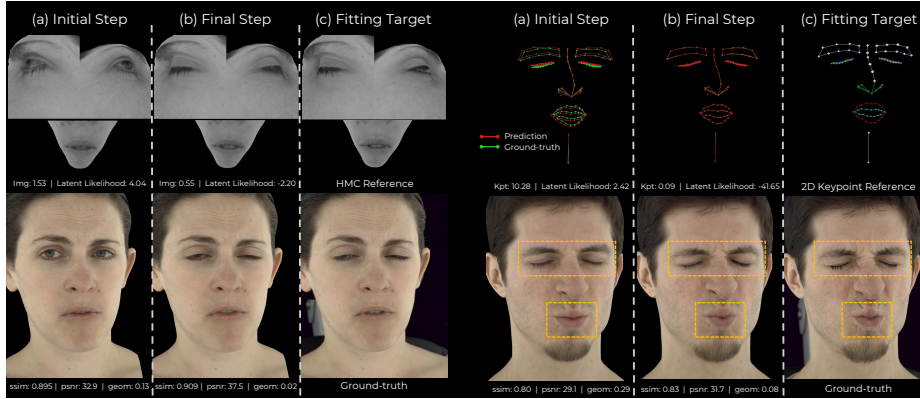


Fig. 6. Reconstruction from inference-time observations with LiP-Flow. We optimize a latent code to reconstruct the 3D avatar (bottom) from imperfect fitting targets (c), such as HMC images (left) or 2D keypoints (right). The latent code is initialized with a sample from the HMC-encoder’s or the KPT-encoder’s prior. We provide the renderings from the initial latent code in columns (a) both in the target’s representation and in the frontal view. The optimization considers the fitting loss between the targets and the HMC (*i.e.* “Img” loss) or 2D keypoint (*i.e.* “Kpt” loss) projections (Eq. 11), and the latent likelihood under the corresponding prior model (Eq. 12). Renderings after the optimization are in the “Final Step” columns (b).

4 Results

We experimentally answer the following questions: (1) How does end-to-end training of the conditional priors affect the base model’s performance? (2) What is the most effective approach of leveraging limited driving signals? (3) Which of the priors is the most useful in inference-time optimization tasks?

The base DAM network consists of a view-conditioned decoder and an encoder expecting rich training data. To evaluate the DAM in run-time conditions, we follow analysis-by-synthesis and optimize latent codes to reconstruct impoverished inference-time observations. For a fair comparison, we also introduce the HMC- or the KPT-encoders into the DAM, leveraging the run-time inputs in different ways and allowing us to make latent code predictions directly.

DAM with an inference-time encoder (Fig. 4-b). The naive way to attain a face avatar from the inference-time inputs is to use the ground-truth correspondences between the multi-view images and the synthetic HMC data or 2D keypoints. To do so, we replace the DAM-encoder with an HMC- or a KPT-encoder and train a conditional VAE as before. We use this model to verify the hypothesis that the limited run-time observations do not carry enough information to faithfully build a high-quality avatar model.

Inference-time encoder as regressor (Fig. 4-c). To analyse the effectiveness of treating the inference-time encoders as a prior and inference-time optimization, we train the KPT- and HMC-encoders to regress the latent codes of the pre-trained DAM-encoder. We follow the probabilistic approach in Eq. 6 and use the negative log-likelihood training objective.

Models	Decoding DAM-encoder (Q)		
	PSNR \uparrow	SSIM \uparrow	Geom \downarrow
(a) DAM	36.05	0.893	0.015
(b) DAM-HMC Enc.	33.93	0.872	0.186
(c) DAM-HMC Reg.	n/a	n/a	n/a
(d) LiP-KL-HMC	35.79	0.890	0.018
(e) LiP-Flow-HMC	36.21	0.895	0.014
(b) DAM-KPT Enc.	33.50	0.866	0.142
(c) DAM-KPT Reg.	n/a	n/a	n/a
(d) LiP-KL-KPT	35.71	0.888	0.021
(e) LiP-Flow-KPT	36.22	0.895	0.014

Table 1. Reconstruction. We evaluate the DAM-encoder (Q) and the decoder (D) by reconstructing the latent codes estimated from the training-time inputs to see how the inference-time encoder affects the base model, DAM. We report the performance with both the HMC-encoder and the KPT-encoder (Fig. 4). The DAM-HMC Enc. and the DAM-KPT Enc. base-lines (b) take HMC images and 2D keypoints as inputs, respectively.

In our work, the latent space formulation is the only difference across models (Fig. 4). To tightly control the settings and simulate the application conditions, we introduce facial landmark and synthetic HMC datasets (Fig. 5). The synthetic HMC images are generated by re-projecting the multi-view camera images into virtual head-mounted camera views with various lighting, background and head-set slop augmentations. For the 2D keypoints, we project a predetermined set of mesh vertices onto the image plane. Both the 2D keypoints and synthetic HMC images provide correspondences to ground truth 3D full face observations, enabling careful study of the methods. We report average performance over 4 subjects in the PSNR and SSIM image metrics, and per vertex geometry loss.

In Sec. 4.1, we present reconstruction performance of the models in the training setup. We then compare our model LiP-Flow against the baselines via inference-time optimization with limited observations in Sec. 4.2. We report the performance of the models with both the HMC-encoder and KPT-encoder as the prior model. Finally, we provide insights on dynamics of our model in Sec. 4.3.

In our supp. mat., we provide experiment details, additional evaluations (Sec. B), qualitative results (Sec. E) and an ablation (Sec. A) where we use our flow-based formulation with a $\mathcal{N}(\mathbf{0}, \mathbf{I})$ instead of a conditional prior.

4.1 Reconstruction Quality

We evaluate the reconstruction performance of the 3D face model in the training setup. More specifically, the decoder (D) reconstructs the latent codes from the DAM-Encoder (Q) where we pass the average texture and the 3D geometry (Eq. 1). We aim to find the effect of end-to-end training of the conditional priors on the base model. Tab. 1 provides the reconstruction results when the HMC- and KPT-encoders are treated as priors and trained along with the DAM.

In both setups, our LiP-Flow outperforms the baselines. The LiP-KL reduces the reconstruction performance of the base model DAM, suggesting that the combination of conditional prior (P) and the KL objective are detrimental. In contrast, our flow-based prior improves the performance of the DAM, implying that LiP-Flow learns a more expressive latent space than the LiP-KL. Also note the consistent performance of our flow-based formulation in the HMC- and KPT-encoder settings. Our learned conditional priors tackle the difficulty of modeling the limited inference-time data, and the base model’s DAM-encoder (Q) and the decoder (D) enjoy a more expressive latent space. Since the HMC- and

Models	HMC Observations		
	PSNR \uparrow	SSIM \uparrow	Geom \downarrow
(a) DAM	31.65	0.858	0.676
(b) DAM-HMC Enc.	33.93	0.872	0.186
(c) DAM-HMC Reg.	34.16	0.872	0.248
(d) LiP-KL-HMC	34.36	0.881	0.090
(e) LiP-Flow-HMC	34.98	0.885	0.087

Table 2. HMC inputs. We use the HMC-encoder (P) to estimate a prior distribution from the HMC images. For the baselines (b) and (c), a latent sample is decoded directly, whereas for (a), (d) and (e), the latent code is first optimized to reconstruct the given HMC images (Fig. 6).

Models	Keypoint (KPT) Observations		
	PSNR \uparrow	SSIM \uparrow	Geom \downarrow
(a) DAM	31.62	0.865	0.431
(b) DAM-KPT Enc.	33.50	0.866	0.142
(c) DAM-KPT Reg.	33.67	0.864	0.151
(d) LiP-KL-KPT	35.08	0.886	0.089
(e) LiP-Flow-KPT	35.55	0.891	0.053

Table 3. 2D keypoint inputs. We use the KPT-encoder (P) to estimate a prior distribution from the keypoints. For the baselines (b) and (c), a latent sample is decoded directly, whereas for (a), (d) and (e), the latent code is first optimized to reconstruct the given keypoints (Fig. 6).

KPT-encoders are required to extract the relevant information only from the limited inputs, the “DAM-HMC Enc.” and “DAM-KPT Enc.” baselines suffer from underfitting, degrading the performance of the DAM.

4.2 Inference-time Optimization

Here, we analyze the models in the inference-time conditions. We treat the pre-trained decoder as a black box renderer for a given latent code, view direction, and camera pose and fit the latent code to reconstruct the given target:

$$\mathbf{z} = \arg \min_{\mathbf{z}} \mathcal{L}_R + \lambda_L \mathcal{L}_L, \quad (10)$$

$$\mathcal{L}_R = \left\| \left(\mathbf{I}^v - \hat{\mathbf{I}}^v \right) \odot \mathbf{W}^v \right\|^2, \quad (11)$$

$$\mathcal{L}_L = -\log p_P(\mathbf{z}) = -\log \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}^{(p)}, \boldsymbol{\sigma}^{(p)}), \quad (12)$$

where $\boldsymbol{\mu}^{(p)}$ and $\boldsymbol{\sigma}^{(p)}$ are predicted by the HMC- or KPT-encoders. The learned prior is leveraged via the latent likelihood term \mathcal{L}_L (Eq. 12). Similarly, the latent codes are initialized with the predicted mean $\boldsymbol{\mu}^{(p)}$. Note that for the base model DAM (Fig. 4-a), the prior is a standard Gaussian. In this task, we use only the HMC- or KPT-encoders and the decoder (D) and report the performance on the fully rendered images including the parts that are not visible in the fitting targets (see Fig. 6). Evaluations without \mathcal{L}_L is provided in Sec. B in supp. mat.

For all models, we assume that the decoder’s input \mathbf{z} stems from Q . Hence, in LiP-Flow, the initial latent sample is mapped to the Q space via the inverse of Eq. 8. The likelihood is computed via Eq. 8 and Eq. 9 (cf. Fig. 3).

HMC Targets. We use the partial HMC views as the fitting target and use the HMC-conditioned priors (see Fig. 6-left). We decode the latent code \mathbf{z} for 3 view vectors and render with the corresponding HMC parameters. For each rendered HMC view, the image fitting objective \mathcal{L}_R (Eq. 11) is calculated on the visible HMC regions only. This task evaluates the models in terms of multi-view consistency under partial views and the reconstruction quality from incomplete

Models	Frontal-view Fitting		
	PSNR \uparrow	SSIM \uparrow	Geom \downarrow
(a) DAM	35.16	0.888	0.212
(b) DAM-HMC Enc.	35.36	0.883	0.047
(c) DAM-HMC Reg.	36.28	0.895	0.057
(d) LiP-KL-HMC	36.08	0.892	0.039
(e) LiP-Flow-HMC	36.50	0.898	0.022
(b) DAM-KPT Enc.	34.75	0.874	0.059
(c) DAM-KPT Reg.	36.31	0.895	0.030
(d) LiP-KL-KPT	35.94	0.890	0.041
(e) LiP-Flow-KPT	36.53	0.898	0.025

Table 4. Fitting to frontal-view images. We evaluate all the models (a-e) in the inference-time optimization task with frontal-view targets. First, the latent codes are fitted to the frontal-view images and then evaluated by decoding in left, right, and frontal views. Note that the prior input is different from the fitting targets. It is conditioned on the HMC images or 2D keypoints.

observations. The prior term \mathcal{L}_L (Eq. 12) is of higher importance in this task as the HMC views are unseen and very different from the decoder’s training views.

In Tab. 2, our flow-based (e) LiP-Flow outperforms the baselines in both the image and geometry metrics. We observe a significant gap between the base model DAM and the models with a learned prior, showing that the standard Gaussian prior is not effective for optimization when only partial information is available. This is partly mitigated by the variants (b) DAM-HMC Enc. and (c) DAM-HMC Reg. which estimate latent codes directly from the observations.

2D Keypoint Targets. We evaluate the models by fitting the latent codes to sparse 2D keypoint observations where the priors and the inference-time encoders are also conditioned on the keypoints (see Fig. 6-right). After decoding a latent code, a predetermined set of vertices on the predicted mesh is projected onto the image plane. The reconstruction term L_R in Eq. 11 is the ℓ^2 norm between the given and the projected keypoints. This means that the latent code is optimized solely based on the geometry and the improved image quality can be attributed to the learned correlations between the texture and the geometry.

The evaluations on keypoints in Tab. 3 are inline with the HMC setup (Tab. 2) where our (e) LiP-Flow achieves the best performance again, demonstrating the generalization ability of our flow-based latent formulation. Moreover, the geometry error is improved for all the models as the underlying prior model and the fitting objective ($\mathcal{L}_R + \mathcal{L}_L$) explicitly use geometric cues.

Frontal-view Targets. In this task, we use frontal-view images as the fitting targets for all the models. The models have access to the frontal views targets, and depending on the configuration, the HMC images or the 2D keypoints as the driving signal. This setup allows us to evaluate the models with more informative observations and assess the learned priors when the fitting target is different from the prior inputs. Tab. 4 summarizes the results for both the HMC- and KPT-encoders as well as the base DAM. Our LiP-Flow achieves the best performance in both image and geometry metrics. The DAM’s performance significantly improves when the fitting target is more informative frontal views, implying that the DAM is not able to handle impoverished observations. We also evaluate the DAM baselines (b and c) via latent code fitting. Similar to the DAM, (b) the DAM-HMC Enc. and the DAM-KPT Enc. use a standard




Fig. 7. Latent space visualization in PCA (Left) Samples in the P space. (Right) Same samples in the Q space after applying our latent transformation via flow F . For a given HMC sample (color-coded), the HMC-encoder predicts a prior distribution. For each of the 15 HMC inputs, we visualize 100 latent codes (see Sec. C in supp. mat.).

Gaussian prior. Our proposed LiP-Flow outperforms the DAM and its variants, showing the advantage of our flow-based conditional prior over the $\mathcal{N}(\mathbf{0}, \mathbf{I})$ prior.

4.3 Latent Space

In Fig. 7, we visualize the P and Q latent spaces in 2D PCA space. Removing the similarity constraint imposed by the KL-divergence loss results in a highly distorted representation space for the prior model (P). It also allows the prior model to assign different amounts of variance to the samples (*i.e.*, the green sample has the largest variance). Considering the performance of our model in the fitting tasks, this unstructured nature of the prior space seems to be effective. In the Q space, however, the same latent samples form well-separated clusters for different inputs, indicating that LiP-Flow captures semantics in Q .

4.4 Ablations

In Tab. 5, we ablate our learned prior in the HMC fitting task. It significantly improves the performance when the latent codes are initialized from the prior and the latent likelihood is considered. Random initializations from the $\mathcal{N}(\mathbf{0}, \mathbf{I})$ causes higher geometry error while the latent likelihood \mathcal{L}_L is highly important for image quality. We also present two variants of our model. We first train “LiP-Flow + Decoding P ” by decoding samples both from the P and Q spaces, causing a detrimental effect. This setup encourages the LiP-encoders (P) to predict the incomplete information in the output space rather than the more compact latent space, limiting the prior’s capacity. We then train “LiP-Flow - Flow Cycle” by ignoring the flow cycle (cf. Fig. 3) and decoding Q samples directly. The reconstruction signal helps the latent flow network to learn transformations that are more accurately reflected in the output space.

4.5 Discussion

Our evaluations provide evidence for the effectiveness of our LiP-Flow. When evaluated with the limited inference-time observations, our proposed conditional prior via our flow-based formulation is superior compared to the standard Normal prior (see Tables 2 and 3, and Fig. 2). This is achieved without sacrificing the base DAM’s performance in training-time setup (Tab. 1). We show that the DAM’s performance can be improved by leveraging the inference-time encoders. However, learning an inference-time regressor (DAM-HMC Reg. or DAM-KPT

Models	HMC Fitting		
	PSNR \uparrow	SSIM \uparrow	Geom \downarrow
DAM	31.09	0.850	0.522
LiP-Flow-Flow Cycle	33.07	0.856	0.108
LiP-Flow+Decoding P	33.03	0.862	0.125
LiP-Flow-prior init- \mathcal{L}_L	28.51	0.821	0.841
LiP-Flow-prior init	32.29	0.862	0.405
LiP-Flow- \mathcal{L}_L	30.34	0.845	0.212
LiP-Flow	33.61	0.874	0.060

Table 5. Ablation. (Top) Training our model by ignoring the flow cycle $\mathbf{z} = F^{-1}(F(\mathbf{z}^{(q)}; \mathbf{h}); \mathbf{h})$, and decoding latent samples from the P space as well. (Bottom) Evaluating our LiP-Flow by not using initializations from the learned prior and the latent likelihood objective \mathcal{L}_L (Eq. 12). Ablation is performed on one subject.

Reg.) separately is not as effective as end-to-end training of the inference-time encoders by treating them as priors. Our LiP-Flow yields consistently the best results via inference-time optimization, indicating that our flow-based formulation improves both the prior’s and the decoder’s performance.

5 Limitations and Future Work

Learning conditional priors relies on the correspondences between the training and inference data, which is not always feasible. For example, in VR telepresence with head-worn displays, we expect a domain gap when attempting to directly replace synthetic HMC images with real ones. Considering the extensive use of synthetic data in other domains [50,27], we think that introducing domain adaptation techniques [47,32] to our pipeline could be a promising direction to make use of unpaired HMC data. Another potential direction is a temporal extension where our flow network considers latent samples from the past frames.

6 Conclusion

Our work introduces a novel representation learning approach to bridge the information asymmetry between the training and inference domains for 3D face avatar models. We propose end-to-end training of the inference-time models as a prior together with the main avatar model to make it aware of the impoverished driving signals. Our prior model is tied to the main model via a normalizing flow which learns to map samples from the prior’s representation space to the main model’s, allowing us to define a latent likelihood objective. We present two related tasks where we augment a 3D face avatar model via a learned prior conditioned on either partial HMC images or sparse 2D keypoints. We experimentally show that our formulation yields an expressive and flexible latent space. In tightly controlled evaluations, our model LiP-Flow outperforms the base model as well as a set of carefully designed baselines in reconstruction and various inference-time optimization tasks. Importantly, our approach does not require modifications to the base model or additional training objectives.

Acknowledgments

This project is partially funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme grant agreement No 717054.



References

1. Abrevaya, V.F., Boukhayma, A., Wuhler, S., Boyer, E.: A decoupled 3d facial shape model by adversarial training. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (October 2019) [4](#)
2. Aksan, E., Hilliges, O.: Stcn: Stochastic temporal convolutional networks. arXiv preprint arXiv:1902.06568 (2019) [2](#), [7](#)
3. Aksan, E., Pece, F., Hilliges, O.: Deepwriting: Making digital ink editable via deep generative modeling. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. pp. 1–14 (2018) [2](#)
4. Asim, M., Daniels, M., Leong, O., Ahmed, A., Hand, P.: Invertible generative models for inverse problems: mitigating representation error and dataset bias. In: International Conference on Machine Learning. pp. 399–409. PMLR (2020) [4](#)
5. Bagautdinov, T., Wu, C., Saragih, J., Fua, P., Sheikh, Y.: Modeling facial geometry using compositional vaes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018) [4](#)
6. Bau, D., Strobel, H., Peebles, W., Wulff, J., Zhou, B., Zhu, J.Y., Torralba, A.: Semantic photo manipulation with a generative image prior. arXiv preprint arXiv:2005.07727 (2020) [2](#), [4](#)
7. Bhattacharyya, A., Hanselmann, M., Fritz, M., Schiele, B., Straehle, C.N.: Conditional flow variational autoencoders for structured sequence prediction. arXiv preprint arXiv:1908.09008 (2019) [2](#), [8](#)
8. Bi, S., Lombardi, S., Saito, S., Simon, T., Wei, S.E., Mcphail, K., Ramamoorthi, R., Sheikh, Y., Saragih, J.: Deep relightable appearance models for animatable faces. ACM Transactions on Graphics (TOG) **40**(4), 1–15 (2021) [1](#)
9. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques. p. 187–194. SIGGRAPH '99 (1999) [4](#)
10. Bogu, F., Kanazawa, A., Lassner, C., Gehler, P., Romero, J., Black, M.J.: Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In: European conference on computer vision. pp. 561–578. Springer (2016) [2](#), [4](#)
11. Bühler, M.C., Meka, A., Li, G., Beeler, T., Hilliges, O.: Varitex: Variational neural face textures. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13890–13899 (2021) [4](#)
12. Cao, C., Agrawal, V., De La Torre, F., Chen, L., Saragih, J., Simon, T., Sheikh, Y.: Real-time 3d neural facial animation from binocular video. ACM Transactions on Graphics (TOG) **40**(4), 1–17 (2021) [2](#), [4](#)
13. Chandran, P., Bradley, D., Gross, M., Beeler, T.: Semantic deep face models. In: 2020 International Conference on 3D Vision (3DV). pp. 345–354. IEEE (2020) [2](#), [4](#)
14. Chen, X., Dong, Z., Song, J., Geiger, A., Hilliges, O.: Category level object pose estimation via neural analysis-by-synthesis. In: European Conference on Computer Vision. pp. 139–156. Springer (2020) [2](#), [3](#), [4](#)
15. Cheng, S., Bronstein, M., Zhou, Y., Kotsia, I., Pantic, M., Zafeiriou, S.: Meshgan: Non-linear 3d morphable models of faces (2019) [4](#)
16. Chu, H., Ma, S., De la Torre, F., Fidler, S., Sheikh, Y.: Expressive telepresence via modular codec avatars. In: European Conference on Computer Vision. pp. 330–345. Springer (2020) [1](#), [2](#), [4](#)
17. Chung, J., Kastner, K., Dinh, L., Goel, K., Courville, A.C., Bengio, Y.: A recurrent latent variable model for sequential data. Advances in neural information processing systems **28**, 2980–2988 (2015) [2](#), [7](#)

18. Dinh, L., Krueger, D., Bengio, Y.: Nice: Non-linear independent components estimation. arXiv preprint arXiv:1410.8516 (2014) [8](#)
19. Dinh, L., Sohl-Dickstein, J., Bengio, S.: Density estimation using real nvp. arXiv preprint arXiv:1605.08803 (2016) [8](#)
20. Feng, Y., Feng, H., Black, M.J., Bolkart, T.: Learning an animatable detailed 3d face model from in-the-wild images. *ACM Transactions on Graphics (TOG)* **40**(4), 1–13 (2021) [4](#)
21. Gecer, B., Ploumpis, S., Kotsia, I., Zafeiriou, S.: Ganfit: Generative adversarial network fitting for high fidelity 3d face reconstruction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1155–1164 (2019) [4](#)
22. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. *Advances in neural information processing systems* **27** (2014) [4](#)
23. Guo, B., Han, Y., Wen, J.: Agem: Solving linear inverse problems via deep priors and sampling. *Advances in Neural Information Processing Systems* **32**, 547–558 (2019) [4](#)
24. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016) [7](#)
25. Kingma, D.P., Dhariwal, P.: Glow: Generative flow with invertible 1x1 convolutions. arXiv preprint arXiv:1807.03039 (2018) [4](#), [8](#)
26. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013) [1](#), [4](#)
27. Kocabas, M., Huang, C.H.P., Tesch, J., Muller, L., Hilliges, O., Black, M.J.: Spec: Seeing people in the wild with an estimated camera. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 11035–11045 (2021) [14](#)
28. Kolotouros, N., Pavlakos, G., Jayaraman, D., Daniilidis, K.: Probabilistic modeling for human mesh recovery. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 11605–11614 (2021) [2](#), [4](#)
29. Lai, G., Li, B., Zheng, G., Yang, Y.: Stochastic wavenet: A generative latent variable model for sequential data. arXiv preprint arXiv:1806.06116 (2018) [2](#)
30. Lewis, J.P., Anjyo, K., Rhee, T., Zhang, M., Pighin, F., Deng, Z.: Practice and theory of blendshape facial models. In: *Eurographics* (2014) [4](#)
31. Li, J., Villegas, R., Ceylan, D., Yang, J., Kuang, Z., Li, H., Zhao, Y.: Task-generic hierarchical human motion prior using vaes. In: *2021 International Conference on 3D Vision (3DV)*. pp. 771–781. IEEE (2021) [2](#), [4](#)
32. Lombardi, S., Saragih, J., Simon, T., Sheikh, Y.: Deep appearance models for face rendering. *ACM Transactions on Graphics (TOG)* **37**(4), 1–13 (2018) [1](#), [2](#), [4](#), [5](#), [6](#), [14](#)
33. Ma, S., Simon, T., Saragih, J., Wang, D., Li, Y., De La Torre, F., Sheikh, Y.: Pixel codec avatars. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 64–73 (2021) [1](#), [4](#)
34. Nair, V., Susskind, J., Hinton, G.E.: Analysis-by-synthesis by learning to invert generative black boxes. In: *International Conference on Artificial Neural Networks*. pp. 971–981. Springer (2008) [3](#)
35. Pan, X., Zhan, X., Dai, B., Lin, D., Loy, C.C., Luo, P.: Exploiting deep generative prior for versatile image restoration and manipulation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2021) [2](#), [4](#)

36. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10975–10985 (2019) 2, 4
37. Rempe, D., Birdal, T., Hertzmann, A., Yang, J., Sridhar, S., Guibas, L.J.: Humor: 3d human motion model for robust pose estimation. *arXiv preprint arXiv:2105.04668* (2021) 2, 4, 7
38. Rezende, D., Mohamed, S.: Variational inference with normalizing flows. In: *International conference on machine learning*. pp. 1530–1538. PMLR (2015) 8
39. Richard, A., Lea, C., Ma, S., Gall, J., De la Torre, F., Sheikh, Y.: Audio-and gaze-driven facial animation of codec avatars. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. pp. 41–50 (2021) 4
40. Schwartz, G., Wei, S.E., Wang, T.L., Lombardi, S., Simon, T., Saragih, J., Sheikh, Y.: The eyes have it: An integrated eye and face model for photorealistic facial animation. *ACM Transactions on Graphics (TOG)* **39**(4), 91–1 (2020) 1
41. Shamaï, G., Slossberg, R., Kimmel, R.: Synthesizing facial photometries and corresponding geometries using generative adversarial networks. *ACM Trans. Multimedia Comput. Commun. Appl.* **15**(3s) (2019) 4
42. Slossberg, R., Shamaï, G., Kimmel, R.: High quality facial surface and texture synthesis via generative adversarial networks. In: *Proceedings of the European Conference on Computer Vision (ECCV) Workshops (September 2018)* 4
43. Sohn, K., Lee, H., Yan, X.: Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems* **28**, 3483–3491 (2015) 7
44. Spurr, A., Song, J., Park, S., Hilliges, O.: Cross-modal deep variational hand pose estimation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 89–98 (2018) 2
45. Tewari, A., Zollöfer, M., Kim, H., Garrido, P., Bernard, F., Perez, P., Christian, T.: MoFA: Model-based Deep Convolutional Face Autoencoder for Unsupervised Monocular Reconstruction. In: *The IEEE International Conference on Computer Vision (ICCV)* (2017) 4
46. Vahdat, A., Kreis, K., Kautz, J.: Score-based generative modeling in latent space. *Advances in Neural Information Processing Systems* **34** (2021) 4
47. Wei, S.E., Saragih, J., Simon, T., Harley, A.W., Lombardi, S., Perdoch, M., Hypes, A., Wang, D., Badino, H., Sheikh, Y.: Vr facial animation via multiview image translation. *ACM Transactions on Graphics (TOG)* **38**(4), 1–16 (2019) 1, 2, 4, 14
48. Whang, J., Lindgren, E., Dimakis, A.: Composing normalizing flows for inverse problems. In: *International Conference on Machine Learning*. pp. 11158–11169. PMLR (2021) 4
49. Winkler, C., Worrall, D., Hoogetboom, E., Welling, M.: Learning likelihoods with conditional normalizing flows. *arXiv preprint arXiv:1912.00042* (2019) 2, 8
50. Wood, E., Baltrusaitis, T., Hewitt, C., Dziadzio, S., Cashman, T.J., Shotton, J.: Fake it till you make it: Face analysis in the wild using synthetic data alone. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3681–3691 (2021) 14
51. Zanfir, A., Bazavan, E.G., Xu, H., Freeman, W.T., Sukthankar, R., Sminchisescu, C.: Weakly supervised 3d human pose and shape reconstruction with normalizing flows. In: *European Conference on Computer Vision*. pp. 465–481. Springer (2020) 4

52. Zhou, X., Leonardos, S., Hu, X., Daniilidis, K.: 3d shape estimation from 2d landmarks: A convex relaxation approach. In: proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4447–4455 (2015) [2](#)