

Pix2NPHM: Learning to Regress NPHM Reconstructions From a Single Image

Simon Giebenhain¹ Tobias Kirschstein¹ Liam Schoneveld²
Davide Davoli^{3*} Zhe Chen² Matthias Nießner¹

¹Technical University of Munich

²Woven by Toyota

³Toyota Motor Europe



Figure 1. **Pix2NPHM** is a feed-forward network that predicts NPHM [12] latent codes from a single image. The latent codes can be further optimized at test-time to obtain more detailed 3D reconstructions. Here, we show mesh overlays showcasing well-aligned fittings of diverse head shapes and expressions under strong lighting conditions and occlusions. Website: <https://simongiebenhain.github.io/Pix2NPHM/>

Abstract

Neural Parametric Head Models (NPHMs) are a recent advancement over mesh-based 3d morphable models (3DMMs) to facilitate high-fidelity geometric detail. However, fitting NPHMs to visual inputs is notoriously challenging due to the expressive nature of their underlying latent space. To this end, we propose Pix2NPHM, a vision transformer (ViT) network that directly regresses NPHM parameters, given a single image as input. Compared to existing approaches, the neural parametric space allows our method to reconstruct more recognizable facial geometry and accurate facial expressions. For broad generalization, we exploit domain-specific ViTs as backbones, which are pretrained on geometric prediction tasks. We train Pix2NPHM on a mixture of 3D data, including a total of over 100K NPHM registrations that enable direct supervision in SDF space, and large-scale 2D video datasets, for which normal estimates serve as pseudo ground truth geometry. Pix2NPHM not only allows for 3D reconstructions at interactive frame rates, it is also possible to improve geometric fidelity by a subsequent inference-time optimization against estimated surface normals and canonical point maps. As a result, we achieve unprecedented face reconstruction quality that can run at scale on in-the-wild data.

1. Introduction

Reconstructing faces in 3D, tracking facial movements, and ultimately extracting animation signals for virtual avatars are fundamental problems in many domains such as the computer games and movie industry, telecommunication, and AR/VR. Arguably the most relevant sub-task is 3D face reconstruction from a single image due to the vast availability of image collections as well as straight-forward extensions to sequential tracking.

In order to solve the underconstrained reconstruction problem, 3d morphable models (3DMMs) [2] have evolved as industry and research standard due to their concise low-dimensional parametric representation, which lead to a plethora of algorithms build on top of 3DMMs. With the advancement of deep-learning methods, photometric tracking [45] approaches, have been augmented with additional priors, such as facial landmark detection, or direct 3DMM parameter regression from RGB signal [8, 10, 39, 56, 63]. Recently, additional priors, such as dense landmarks [41, 52] and surface normals [15] have further improved reconstructions. Due to such methods that enable fitting in even the most challenging scenarios, 3DMMs have become an essential component of photo-realistic avatars [14, 33], generalized avatars [4, 23, 24], and even controllable generative

* Providing contracted services for Toyota

diffusion models for faces [22, 32, 42, 43, 60].

While 3DMMs have achieved great success in these domains, we argue that their concise parametric representation comes at the cost of geometric expressiveness – i.e., modern 3DMMs, such as FLAME [25], are unable to model high-fidelity geometric detail. Therefore, a more recent line of work has developed neural parametric head models (NPHMs) [12, 13, 54, 57] for increased representational capacity, as shown in Fig. 2. This increased model capacity, however, makes image-based reconstruction challenging due its expressive parameter space. MonoNPHM [13] has attempted to reconstruct NPHM parameters from a single image. However, their purely photometric fitting approach remained slow and brittle in real-world applications. To this end, we propose a robust and high-fidelity fitting framework, yielding a first-class tool for face reconstruction and tracking based on NPHM [12, 13].

Our approach addresses the two main challenges of neural parametric model fitting: underconstrained optimization and reconstruction speed. This is achieved by tailoring a transformer-based feed-forward predictor for NPHM parameters from a single image. As a highly data-driven approach, large-scale high-quality training data is essential. To this end, we curated a large collection of publicly available 3D face datasets and fitted MonoNPHM against it, resulting in a total of 102K registrations, which will be shared with the research community. Despite these efforts, we find that training on large-scale 2D video datasets using a self-supervised geometric loss based on estimated surface normals [15] further improves generalization. Furthermore, we observe strong generalization improvements by replacing a generic DINOv2 [29] backbone with a ViT [9] pre-trained on per-pixel geometric prediction tasks, such as surface normal or canonical point map regression. Our feed-forward estimator produces state-of-the-art (SotA) results, and renders Pix2NPHM a first-class choice for 3D face reconstruction, due to its fidelity, ease of use, and robustness to diverse input scenarios, as showcased in Fig. 1. Moreover, we show that our results can be refined by a few optimization steps at inference time against estimated surface normals to increase fidelity, as shown in Fig. 2. Together, these insights improve 3D reconstruction on the NeRSemble single-image face reconstruction (SVFR) [15] benchmark by 21%, and neutral reconstruction improves by 6% compared to the best public method on the NoW [39] benchmark. To summarize, our main contributions are as follows:

- Pix2NPHM is the first feed-forward regressor for NPHM parameters, which, for the first time, enables robust, accurate and fast NPHM reconstructions from a single image.
- We curate a large mixture of high-quality 3D datasets, with over 100K NPHM registrations.
- For training on 2D data we formulate a novel self-supervised loss using estimated surface normals.




Figure 2. **Motivation:** Single-image 3DMM regressors are limited by their underlying 3DMM. More detailed reconstructions can be obtained by replacing FLAME [25] with NPHM [12], and running inference-time optimization further increase fidelity (see right).

- Replacing DINOv2 with geometrically pre-trained, face-specific ViTs improves generalization.

2. Related Work

Optimization-Based Methods Early mesh-based 3D Morphable Models (3DMMs) [2, 19, 25] represent facial shape and appearance via linear PCA spaces. Fitting these models to images or videos is typically approached through iterative optimization [2, 16, 33, 45, 63], usually by minimizing photometric error between renderings and the input. Because photometric losses are under-constrained and sensitive to illumination and noise, many works incorporate geometric priors such as sparse landmarks or semantic segmentations [16, 33, 45, 63] for stability. More recent methods leverage dense geometric cues to further improve robustness and accuracy, such as dense landmarks [52], UV flow [41], or a combination of UV maps and surface normals [15].

Regression-Based Methods Regression-based approaches train deep networks to directly predict 3DMM parameters, avoiding iterative optimization and enabling real-time inference. Fully supervised approaches [17, 35, 36, 46, 47, 63] rely on registered 3D scans or synthetic data, but are limited by the scarcity of high-quality annotations. Self-supervised methods learn from in-the-wild images by combining differentiable rendering with photometric constraints [8, 10, 11, 34, 39, 40, 44, 56]. EMOCA [8] extends this by incorporating emotion-aware supervision for more expressive results, and SPECTRE [11] integrated lipreading-based cues to better capture mouth movements and speech-related expressions. Token-

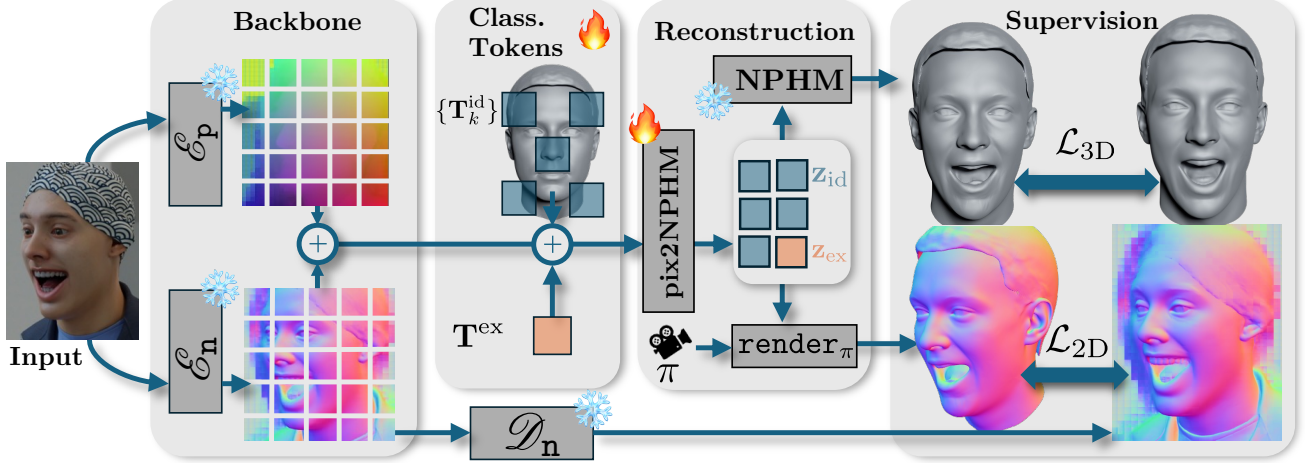


Figure 3. **Method Overview:** We use pretrained ViTs, \mathcal{E}_n and \mathcal{E}_p , as backbone, which encode the input into a token sequence. The resulting sequence is concatenated with learnable classifier token $\{T_k^{\text{id}}\}$ and T^{ex} , which are decoded into NPHM identity (z_{id}) and expression (z_{ex}) parameters using transformer network. We train using a 3D SDF loss, and a normal rendering loss against pseudo g.t. normals.

Face [56] combines 3D and 2D supervision in a ViT-based framework. Recent work highlights the role of improved differentiable rendering: SMIRK [34] employs a neural renderer, while SHeaP [40] uses 2D Gaussian splats [18]. Our method follows this self-supervised line but replaces photometric cues with surface-normal supervision and builds upon MonoNPHM rather than FLAME.

Neural 3DMMs Classical 3DMMs [2, 19, 25] rely on fixed topology and linear PCA spaces, limiting detail and expressiveness. Neural 3DMMs address these issues by using implicit neural representations such as SDFs [12, 30, 54, 57, 58] or occupancy grids [5]. NPHM [12] models head shape and appearance via local SDF experts anchored at semantic keypoints, and uses a deformation field for expressions. MonoNPHM [13] proposed the first monocular fitting method for NPHM through volumetric SDF rendering and canonical appearance modeling. Other works explore photorealistic neural head models using NeRFs [27] and 3D Gaussian splats [20], such as [3, 20, 27, 48, 53, 59]. Compared to FLAME, NPHM’s manifold has proven valuable for downstream tasks, including realistic avatars [14, 22] and high-fidelity audio-driven geometry [1]. Motivated by this, we develop a fast, robust, and accurate reconstruction approach that makes MonoNPHM practical for broader downstream use.

3. Pix2NPHM

Given a single input image I as input, it is our goal to estimate identity parameters z_{id} and expression parameters z_{ex} from which the 3D head geometry can be recovered using MonoNPHM [13] as decoder. After providing background information on NPHMs in Sec. 3.1, we begin by describing

our face-specific pre-training strategy to obtain our image encoding backbone in Sec. 3.2. Next, in Sec. 3.3 we introduce Pix2NPHM, the first method for feed-forward NPHM parameter regression based on a single image. Fig. 3 provides an overview of our network architecture, and training strategy, which is further described in Sec. 3.4. Finally, Sec. 3.5 describes our inference-time optimization, which can further improve our reconstruction fidelity, as showcased in Fig. 2.

3.1. Background: Neural 3DMMs

In this paper we adopt the neural 3DMM formulation introduced by MonoNPHM [13]. It defines a neural field

$$\mathcal{F}_{\text{NPHM}} : (x, z_{\text{id}}, z_{\text{ex}}) \mapsto \text{SDF}(x), \quad (1)$$

which predicts SDF-values for the 3D head geometry, given points $x \in \mathbb{R}^3$, and is conditioned on disentangled identity and expression latent codes. Internally, $\mathcal{F}_{\text{NPHM}}$ consists of a backward deformation field, conditioned on z_{ex} and z_{id} and a canonical SDF conditioned on z_{id} . A mesh can be extracting using marching cubes [26], once latent codes have been obtained. The remainder of the paper will focus on the inference of ideal latent codes from a single image. Throughout all our experiments we use the public checkpoint from MonoNPHM [13].

3.2. Robust Encoder via Geometric Pre-Training

We tackle the inherent ambiguity of the single-image reconstruction task using a heavily data-driven approach. Our method starts with pre-training a ViT-based backbone on a per-pixel geometric reconstruction task. We follow the

training strategy and encoder-decoder architecture

$$\mathcal{E}_{\text{geo}} : \mathbb{R}^{H \times W \times 3} \rightarrow \mathbb{R}^{L \times D} \quad (2)$$

$$\mathcal{D}_{\text{geo}} : \mathbb{R}^{L \times D} \rightarrow \mathbb{R}^{H \times W \times 3} \quad (3)$$

from Pixel3DMM [15], where $\text{geo} \in \{\text{n}, \text{p}\}$ represent surface normal and canonical point cloud estimators, respectively. While point cloud prediction has to be handled using relative coordinate systems for arbitrary scenes, such as in DUST3R [51] and VGGT [49], we exploit the possibility to define a unique coordinate system for 3D heads. Once pre-training is completed, we discard \mathcal{D}_{geo} , and use the face-specific geometrically aware encoders \mathcal{E}_{geo} as backbone for our regression network. Later we will show that this face-specific pre-training significantly outperforms DINOv2 image encodings.

3.3. Feed-Forward NPHM Parameter Prediction

Following the recent success of transformer networks, we employ learnable classifier tokens, similar to ViT [9] and TokenFace [56], which read out relevant identity and expression information after several transformer layers. To this end we concatenate our encoded token sequence $[\mathcal{E}_{\text{n}}(I), \mathcal{E}_{\text{p}}(I)]$ with learnable classifier tokens

$$\mathbf{T}_{\text{CLS}} = [\mathbf{T}^{\text{ex}}, \{\mathbf{T}_k^{\text{id}}\}_{k=1}^{66}] \quad (4)$$

for expression and identity. Note that \mathbf{T}^{id} is composed of 65 tokens, one for each local identity codes from MonoNPHM [13], and one for the global identity part. Finally, our regressor

$$\text{pix2NPHM} : ([\mathcal{E}_{\text{n}}(I), \mathcal{E}_{\text{p}}(I)], \mathbf{T}_{\text{CLS}}) \mapsto (\mathbf{z}_{\text{id}}, \mathbf{z}_{\text{ex}}) \quad (5)$$

consists of several transformer layers, and MLP prediction heads which map the resulting classification tokens \mathbf{T}'_{CLS} from the transformer to \mathbf{z}_{id} and \mathbf{z}_{ex} , respectively. Finally, the predicted latent codes can be decoded into a 3D shape using the MonoNPHM model, which represents the 3D head geometry implicitly by conditioning an SDF, see Eq. (1)

3.4. Training Strategy

3.4.1. 3D supervision in SDF Space

On 3D datasets we aim to formulate the supervision signal as unambiguous as possible, and exploit the fact that ground truth NPHM codes $(\mathbf{z}_{\text{id}}^{\text{gt}}, \mathbf{z}_{\text{ex}}^{\text{gt}})$ can be obtained as a pre-processing step using our registration procedure, as described in Sec. 3.6. Given an image I , we predict NPHM latents $\text{pix2NPHM}(I, \mathbf{T}_{\text{CLS}}) = (\hat{\mathbf{z}}_{\text{id}}, \hat{\mathbf{z}}_{\text{ex}})$, and compute the loss

$$\mathcal{L}_{3\text{D}} = \sum_{x \in \mathcal{X}} \|\text{NPHM}(x; \hat{\mathbf{z}}_{\text{id}}, \hat{\mathbf{z}}_{\text{ex}}) - \text{NPHM}(x; \mathbf{z}_{\text{id}}^{\text{gt}}, \mathbf{z}_{\text{ex}}^{\text{gt}})\|_1, \quad (6)$$

between the induced SDFs, where \mathcal{X} is a point cloud randomly sampled near the 3D surface. Note, that directly supervising $\|\hat{\mathbf{z}} - \mathbf{z}\|$ did not lead to training convergence.

3.4.2. 2D Self-Supervision: Adding Diversity

Since 3D datasets barely cover all relevant modes of the valid input distribution, we add data diversity by training on large scale 2D video datasets using estimated surface normals $I^{\text{n}} = \mathcal{D}_{\text{n}}(\mathcal{E}_{\text{n}}(I))$ as pseudo ground truth. To this end, we modify the NeuS-based [50] rendering approach from MonoNPHM [13] to render surface normals, and supervise via the cosine similarity

$$\mathcal{L}_{2\text{D}}^{\text{n}} = \sum_{\mathbf{p} \in \mathcal{P}} \langle \text{render}_{\pi}(\text{NPHM}; \hat{\mathbf{z}}_{\text{id}}, \hat{\mathbf{z}}_{\text{ex}})_{\mathbf{p}}, I_{\mathbf{p}}^{\text{n}} \rangle \quad (7)$$

between rendered and estimated normals, where camera parameters π are provided by our data registration. Note that due to the memory-intensive nature of MLP-based volumetric rendering we only supervise a random subset of pixels $\mathbf{p} \in \mathcal{P}$ sampled in the facial area. Overall, we find this supervision to provide much more meaningful and stable gradients, compared to computing a photometric loss using spherical harmonics, which is the most established loss function in existing FLAME-based regressors [10, 40].

3.4.3. Full Training Objective

Overall, we train our feed-forward predictor using

$$\mathcal{L}_{\text{total}} = \lambda_{3\text{D}} \mathcal{L}_{3\text{D}} + \lambda_{2\text{D}} \mathcal{L}_{2\text{D}} + \lambda_{\text{reg}} \mathcal{R} \quad (8)$$

as our complete training objective, which combines 3D and 2D losses with regularization terms, where $\lambda_{3\text{D}}$ is set to 0 for 2D video datasets. Our regularization term

$$\mathcal{R}(\mathbf{z}_{\text{id}}, \mathbf{z}_{\text{ex}}) = \lambda_{\text{id}}^{\mathcal{R}} \|\mathbf{z}_{\text{id}}\|_2 + \lambda_{\text{ex}}^{\mathcal{R}} \|\mathbf{z}_{\text{ex}}\|_2 \quad (9)$$

simply punishes the norm of the predicted latents.

3.5. Test-Time Optimization: Increasing Fidelity

For even more accurate 3D reconstructions our feed-forward estimates can be naturally combined with inference-time optimization against per-pixel geometric predictions, similar to Pixel3DMM [15]. We optimize for

$$\arg \min_{\mathbf{z}_{\text{id}}, \mathbf{z}_{\text{ex}}, \pi} \lambda_{\text{n}} \mathcal{L}_{2\text{D}}^{\text{n}} + \lambda_{\text{p}} \mathcal{L}_{2\text{D}}^{\text{p}} + \lambda_{\text{reg}} \mathcal{R}(\mathbf{z}_{\text{id}} - \hat{\mathbf{z}}_{\text{id}}, \mathbf{z}_{\text{ex}} - \hat{\mathbf{z}}_{\text{ex}}), \quad (10)$$

where our feed-forward predictions serve as initialization and regularization target, and $\mathcal{L}_{2\text{D}}^{\text{p}}$ is defined similar to Eq. (7) but with an L_1 -Loss. Note that we first need to estimate camera parameters π , which includes the head pose, using dense landmark predictions from Pixel3DMM, which are fine-tuned later alongside \mathbf{z}_{id} and \mathbf{z}_{ex} .

3.6. High-Quality Training Dataset Curation

As a data-driven approach, a core requirement for success is sufficient high-quality training data. Therefore, we spend a considerable amount of effort on curating 3D and 2D training datasets. We briefly describe our approach for different data types below, and provide an overview of the used training data in table Tab. 1.

	#IDs	#Expr./#Views	#Images	#3D shapes
NPHM [12]	450	23 / 40	414K	10K
FaceScape [62]	300	20 / 50	300K	6K
DAViD [38]	65K	1 / 1	65K	65K
MimicMe [31]	2000	10 / 5	100K	20K
LYHM [7]	1200	1 / 2	2.4K	1.2K
Videos [6, 55, 61]	<50K	5 / 1	250K	0
Total 3D	69K	-/-	880K	102K
Total	119K	-/-	1.13M	102K

Table 1. **Training Data:** We list number of identities, facial expressions, camera view-points, total images and total 3D shapes.

3D datasets For 3D datasets, we estimate the canonical coordinate frame using FLAME registration similar to [12]. After transforming the ground truth 3D shapes into a unified coordinate system, we sample points $x \in \mathcal{X}$ on the ground truth surface and optimize for NPHM parameters

$$\arg \min_{\mathbf{z}_{\text{id}}, \mathbf{z}_{\text{ex}}} \sum_{x \in \mathcal{X}} \|\text{NPHM}(x; \mathbf{z}_{\text{id}}, \mathbf{z}_{\text{ex}})\|_1 + \lambda_{\text{reg}} \mathcal{R}(\mathbf{z}_{\text{id}}, \mathbf{z}_{\text{ex}}) \quad (11)$$

2D datasets For large-scale, in-the-wild 2D datasets, obtaining reliable and 3D accurate NPHM fittings was a previously unsolved problem. Therefore, we simply estimate camera poses using the video tracker from Pixel3DMM [15], and purely rely on our rendering loss \mathcal{L}_{2D}^n .

Data Cleaning To ensure high-quality training data, we devise a simple outlier filtering strategy, for both, 3D NPHM fittings and 2D FLAME video fittings. We define outlier thresholds, by analyzing the histograms of the norms of certain attributes, such as shape, expressions, neck and jaw parameters.

4. Experimental Results

To verify our proposed method experimentally, we compare against recent SotA methods on the NeRSemble single-view face reconstruction (SVFR) benchmark [15] (see Sec. 4.3), and on the established Now benchmark [39] (see Sec. 4.4). In Sec. 4.6 we conduct thorough ablations experiments to quantify the significance of our individual technical contributions. We highly encourage the reviewers to watch our supplementary video for more results. We will release our training data, model and full code base for research purposes.

4.1. Implementation Details

We implement our pipeline using PyTorch, and train our regression model using the Adam [21] optimizer, a batch

Method	Neutral			Posed		
	L1↓	L2↓	NC↑	L1↓	L2↓	NC↑
MICA [63]	1.68	1.14	0.883	-	-	-
TokenFace [56]	-	-	-	2.62	1.78	0.865
DECA [10]	2.07	1.40	0.876	2.38	1.61	0.870
EMOCAv2[8]	2.21	1.49	0.873	2.63	1.78	0.860
SHeaP [40]	1.86	1.26	0.882	2.08	1.41	0.876
Ours (ffwd. only)	1.57	1.06	0.896	1.55	1.05	0.894
Metr. Tracker[63]	-	-	-	2.03	1.37	0.878
FlowFace [41]	1.93	1.31	0.878	1.96	1.33	0.879
Pixel3DMM [15]	1.66	1.12	0.883	1.66	1.11	0.884
MonoNPHM [13]	2.32	1.56	0.878	2.50	1.68	0.870
Ours	1.54	1.04	0.897	1.37	0.92	0.897

Table 2. **NeRSemble-SVFR Benchmark [15]:** Single image posed and neutral 3D face reconstruction. Methods above the line represent feed-forward networks, methods below require optimization.

Method	Error (mm)			Avail.
	Median	Mean	Std	
DECA [10]	1.09	1.38	1.18	✓
MICA [63]	0.90	1.11	0.92	✓
SHeaP [40]	0.95	1.18	0.99	✓
TokenFace [56]	0.76	0.95	0.82	✗
Ours (ffwd. only)	0.83	1.03	0.88	✓*
FlowFace [41]	0.87	1.07	0.88	✗
Pixel3DMM [15]	0.87	1.07	0.89	✓
Ours	0.81	1.01	0.85	✓*

Table 3. **NoW Benchmark [39]:** Single-image neutral 3D face reconstruction. Methods above the line are feed-forward networks.

size of 32 and learning rate of $1e^{-4}$ on a single A100-80GB GPU, which takes 4 days until convergence. We set $\lambda_{3D}=10.0$, $\lambda_{2D}=1.0$ and $\lambda_{\text{reg}}=1e^{-4}$. We use 8 transformer layers, each consisting of pre-norm LayerNorm, self-attention with 8 heads, a 2-layer MLP with GeLU activation, and a width of 1024.

We pre-train \mathcal{E}_n and \mathcal{E}_p separately for 3 days on 2 A6000 GPUs, and follow the network architecture and training strategy from Pixel3DMM [15].

During inference our complete feed-forward estimator runs at 8fps on an RTX3080-10GB. For optimization we perform 100 steps, which takes 85s on an RTX3080. For more details on the optimization procedure, we refer to our supplementary material. To register our training dataset, consisting of 102K 3D shapes, with the NPHM model, we invested roughly 2.500 GPU-hours. We will release our training data to the community to drive future research.

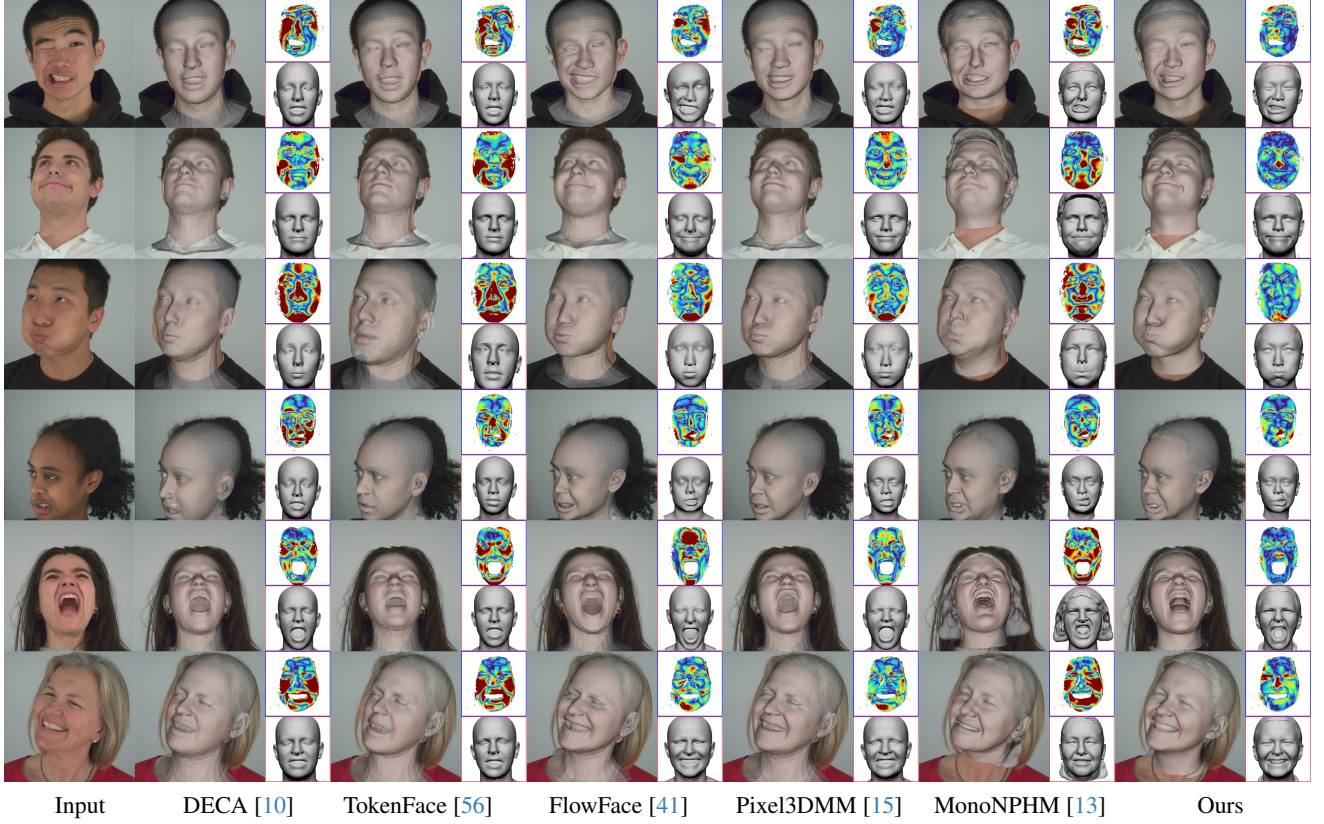


Figure 4. **Posed Reconstruction:** We show overlays of the reconstructed meshes to judge the reconstruction alignment. Insets with a blue border depict L_2 -Chamfer distance as an error map, rendered from a frontal camera. Red insets show the reconstructed mesh from the same camera. All our figures are best viewed digitally and zoomed-in.




Figure 5. **Neutral Reconstruction, NeRsemble:** Comparison against available SotA methods on top of neutral reference image.

4.2. Baselines

Feed-Forward FLAME Regressors The first class of baselines are existing regressors for FLAME [25] parameters. We select DECA [10], MICA [63], EMOCaV2 [8], TokenFace [56] and SHeaP [40].

FLAME Optimization The next class combines feed-forward predictions as a prior with inference-time optimization. As such, this category is closely related to our full approach, but uses a classical 3DMM representa-




Figure 6. **Neutral Reconstruction, NoW:** We show frontal mesh renderings in comparison to the ground truth mesh, as well as, error maps and zoom ins.

tion. We select MetricalTracker [63], FlowFace [41] and Pixel3DMM [15].

NPHM Optimization Finally, MonoNPHM [13] proposes a photometric fitting pipeline to obtain NPHM pa-

rameters. Since we also use the MonoNPHM model to decode latent parameters into 3D shapes, this comparison highlights the added robustness of our approach compared to photometric fitting.

Baseline Availability Finally, we point out the poor public availability of recent SotA methods, *i.e.* TokenFace [56] and FlowFace [41], two of the top performing methods on NoW, remain unreleased, and haven’t been replicated as of yet.

4.3. NeRSemble SVFR Benchmark

The recent NeRSemble SVFR benchmark [15] features diverse and challenging facial expressions, and is the first to simultaneously allow the evaluation of *posed* and *neutral* face reconstruction. Given a single image as input, the *posed* task measures the 3D reconstruction accuracy, while the *neutral* task measures how well the face under neutral expression can be reconstructed. We provide quantitative and qualitative results in Tab. 2 and Figs. 4 and 5, respectively. We significantly outperform all baselines across all metrics, even without test-time optimization.

4.4. NoW Benchmark

The NoW benchmark [39] can only measure the *neutral* reconstruction task, but, compared to the NeRSemble SVFR benchmark, it covers a wider variety of identities, lighting conditions, hair styles, head accessories and other types of occlusion. Our method significantly outperforms all baselines, except for TokenFace [56], which is, however, not publicly available and performs poorly on the NeRSemble SVFR benchmark, as presented in Tab. 3. Qualitative results from the validation set are presented in Fig. 6.

4.5. AffectNet Benchmark

Although emotion recognition is not directly measuring any sort of 3D understanding, we further assess the degree to which the semantics of facial expressions are accurately captured by our feedforward network on AffectNet [28]. We extract shape and expression parameters using Pix2NPHM for all train and test images in the AffectNet dataset (note that this is done in a feedforward-only manner, without per-image optimization). Following EMOCA [8], we then fit a 4-layer MLP to the the parameters extracted from the train set, and evaluate on the test set. Pix2NPHM outperforms all competing methods on this benchmark (Table 4), showcasing our trained feedforward network’s ability to extract detailed expression information.

4.6. Ablation Studies

We demonstrate the effectiveness of our technical contributions by removing them one-by-one from our complete method. Quantitative and qualitative results are presented in

Model	Arousal		Valence		Emo
	CCC \uparrow	RMSE \downarrow	CCC \uparrow	RMSE \downarrow	Acc.
SMIRK [34]	0.560	0.288	0.681	0.313	0.653
EMOCA [8]	0.577	0.282	0.70	0.307	0.676
SHeaP [40]	0.615	0.274	0.735	0.301	0.695
Ours-FLAME	0.553	0.289	0.702	0.305	0.660
Ours-NPHM	0.621	0.273	0.739	0.291	0.711

Table 4. **Emotion recognition:** We report the concordance correlation coefficient (CCC) and root mean squared error (RMSE) on predicting Valence and Arousal, and, most importantly, we report accuracy in 8-way emotion classification on AffectNet [28].

Method	2D	3D	Input	Opt.	3DMM	Error
3D only		\checkmark	$\mathcal{E}_n + \mathcal{E}_p$		NPHM	1.074
2D only	\checkmark		$\mathcal{E}_n + \mathcal{E}_p$		NPHM	1.238
RGB Input		\checkmark	DINO(I)		NPHM	1.165
Normals Inp.	\checkmark	\checkmark	DINO(I^n)		NPHM	1.168
no \mathcal{E}_p	\checkmark	\checkmark	\mathcal{E}_n		NPHM	1.053
ffwd NPHM	\checkmark	\checkmark	$\mathcal{E}_n + \mathcal{E}_p$		NPHM	1.016
ffwd. FLAME	\checkmark	\checkmark	$\mathcal{E}_n + \mathcal{E}_p$		FLAME	1.073
opt. only w/o MICA	\checkmark	\checkmark	$\mathcal{E}_n + \mathcal{E}_p$	\checkmark	NPHM	1.137
opt. only w/ MICA	\checkmark	\checkmark	$\mathcal{E}_n + \mathcal{E}_p$	\checkmark	NPHM	1.029
Ours	\checkmark	\checkmark	$\mathcal{E}_n + \mathcal{E}_p$	\checkmark	NPHM	0.988

Table 5. **NoW Benchmark [39]:** Ablations are performed on the validation subset of the NoW benchmark.

Method	Neutral			Posed		
	L1 \downarrow	L2 \downarrow	NC \uparrow	L1 \downarrow	L2 \downarrow	NC \uparrow
3D only	1.64	1.11	0.894	1.65	1.11	0.890
2D only	1.79	1.21	0.893	1.67	1.13	0.893
RGB input	1.89	1.28	0.891	2.10	1.42	0.883
Normals inp.	1.76	1.19	0.893	1.87	1.26	0.886
no \mathcal{E}_p	1.67	1.13	0.895	1.65	1.11	0.891
ffwd. NPHM	1.57	1.06	0.896	1.55	1.05	0.894
ffwd. FLAME	1.71	1.15	0.883	1.81	1.22	0.881
opt. only w/o MICA	1.76	1.18	0.894	1.61	1.09	0.893
opt. only w/ MICA	1.60	1.08	0.890	1.50	1.01	0.895
Ours	1.54	1.04	0.897	1.37	0.92	0.897

Table 6. **Ablations on NeRSemble-SVFR Benchmark [15].**


Figs. 7 and 8, and Tabs. 5 and 6, respectively. Additionally, Tab. 5 has a verbose description of the ablated components.

What effect does the underlying 3DMM have? The main motivation for Pix2NPHM is the increased representational capacity of neural 3DMMs, *i.e.* MonoNPHM in our case. To ablate the effect of the underlying 3DMM, we train a version of our approach with FLAME instead of MonoNPHM. Quantitative and qualitative results on the SVFR-Benchmark and NoW support our claim, that leveraging a neural 3DMM increases reconstruction fidelity. Note that our FLAME-based feed-forward predictor outperforms



Input ffwd. FLAME ffwd. NPHM Opt. only w/ MICA Ours

Figure 7. **Ablations, Posed:** NPHM feed-forward predictions exhibit more details compared to FLAME. Without the feed-forward initialization our optimization sometimes fails to reconstruct extreme expressions (*e.g.* see rows 1 and 3).



Input ffwd. FLAME ffwd. NPHM Opt. only w/ MICA Ours Reference

Figure 8. **Ablations, Neutral:** Without the feed-forward prior, our optimization cannot properly disentangle identity and expression.

the best available FLAME-based predictor MICA on the NoW validation set (1.073 vs. 1.109 mean error), and outperforms SHeaP, the best FLAME-based feed-forward method on the posed NeRSemble SVFR task. On Affect-Net, MonoNPHM also provides significantly better scores than FLAME, as shown in Tab. 4.

Inference-Time Optimization While our feed-forward prediction already reaches SotA performance, reconstruction fidelity can be improved further, *e.g.* see Fig. 2. Especially posed reconstructions benefit from optimization, while neutral reconstructions improve slightly (see Tab. 2). Importantly, we note that optimization without using our feed-forward prior sometimes fails to reconstruct complicated expressions, *e.g.* see rows one and three of Fig. 7.

Network Architecture Furthermore, we ablate several aspects of our method’s architecture, especially w.r.t. to how the input image is encoded. Using DINOv2 [29] encodings of the RGB input I , similar to Pixel3DMM [15], performs the worst, and did not properly converged on 2D train-

ing data. Using DINOv2 encodings of estimated normals $I^n = \mathcal{D}_n(\mathcal{E}_n(I))$ resulted in similar performance on NoW, but increased performance on NeRSemble. Directly using the tokens from \mathcal{E}_n gives significantly better performance and adding tokens from \mathcal{E}_p further boosts results.

Training Strategy/Data Finally, we show the importance of training on a mixture of 2D and 3D data. We hypothesize that 2D training data is essential to handle the appearance diversity of input images, while 3D data provides an unambiguous training signal.

4.7. In-the-Wild Results

For results on in-the-wild images we refer to Fig. 1 and our supplementary material. Additionally, Fig. 2 compares our FLAME-based and NPHM-based regressors, and our full approach.

5. Limitations and Future Work

We believe that our scalable approach to 3D face reconstructions shows great potential, however, performance is currently limited by several aspects of MonoNPHM [13]. For example, its latent space prevents reliable registration of 3D hair styles, which leads to suboptimal geometry, even in the facial region. Moreover, the MLP-based volumetric rendering and marching cubes extraction introduce substantial computational overhead. Future work could develop an improved NPHM variant, *e.g.* leveraging 2DGS [18] for faster rendering and mesh extraction, and training on more diverse 3D hairstyles such as Difflocks [37]. Jointly fine-tuning NPHM with the feed-forward regressor is another promising direction. Finally, as single-image reconstruction remains inherently ambiguous, incorporating uncertainty estimation, probabilistic reconstruction, *e.g.* via conditional generation, or multi-image extensions, *e.g.* by leveraging sequential information as in VGGT [49], could resolve any remaining ambiguities.

6. Conclusion

We introduced Pix2NPHM, the first feed-forward framework for predicting NPHM parameters, enabling fast and high-fidelity 3D face reconstruction from a single input image. Despite considerable research effort on FLAME parameter regression, our initial attempt on NPHM reconstruction demonstrates significant benefits. Key to this success are our large-scale 3D data registration and the self-supervised training strategy on 2D data using surface normal estimators. Moreover, we show that ViTs pretrained on face-specific geometric tasks capture facial structure far more effectively than DINOv2. Together, these insights establish a new path toward scalable, high-quality monocular 3D face reconstruction.

Acknowledgements

This work was supported by Toyota Motor Europe and Woven by Toyota. This work was also supported by the ERC Consolidator Grant Gen3D (101171131). We would also like to thank Angela Dai for the video voice-over.

References

- [1] Shivangi Aneja, Justus Thies, Angela Dai, and Matthias Nießner. Facetalk: Audio-driven motion diffusion for neural parametric head models. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3
- [2] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194, 1999. 1, 2, 3
- [3] Marcel C. Buehler, Gengyan Li, Erroll Wood, Leonhard Helminger, Xu Chen, Tanmay Shah, Daoye Wang, Stephan Garbin, Sergio Orts-Escolano, Otmar Hilliges, Dmitry Lagun, Jérémy Riviere, Paulo Gotardo, Thabo Beeler, Abhimitra Meka, and Kripasindhu Sarkar. Cafca: High-quality novel view synthesis of expressive faces from casual few-shot captures. In *ACM SIGGRAPH Asia 2024 Conference Paper*. 2024. 3
- [4] Chen Cao, Tomas Simon, Jin Kyu Kim, Gabe Schwartz, Michael Zollhoefer, Shunsuke Saito, Stephen Lombardi, Shih-En Wei, Danielle Belko, Shou-I Yu, Yaser Sheikh, and Jason Saragih. Authentic volumetric avatars from a phone scan. *ACM Trans. Graph.*, 41(4), 2022. 1
- [5] Xu Chen, Tianjian Jiang, Jie Song, Jinlong Yang, Michael J Black, Andreas Geiger, and Otmar Hilliges. gdna: Towards generative detailed neural avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20427–20437, 2022. 3
- [6] Jiahao Cui, Hui Li, Yun Zhan, Hanlin Shang, Kaihui Cheng, Yuqi Ma, Shan Mu, Hang Zhou, Jingdong Wang, and Siyu Zhu. Hallo3: Highly dynamic and realistic portrait image animation with video diffusion transformer, 2024. 5
- [7] Hang Dai, Nick Pears, William A. P. Smith, and Christian Duncan. Statistical modeling of craniofacial shape and texture. *International Journal of Computer Vision*, 128(2):547–571, 2020. 5
- [8] Radek Daněček, Michael J Black, and Timo Bolkart. Emoca: Emotion driven monocular face capture and animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20311–20322, 2022. 1, 2, 5, 6, 7
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2, 4
- [10] Yao Feng, Haiwen Feng, Michael J Black, and Timo Bolkart. Learning an animatable detailed 3d face model from in-the-wild images. *ACM Transactions on Graphics (ToG)*, 40(4): 1–13, 2021. 1, 2, 4, 5, 6
- [11] Panagiotis P. Filntisis, George Retsinas, Foivos Paraperas-Papantoniou, Athanasios Katsamanis, Anastasios Roussos, and Petros Maragos. Visual speech-aware perceptual 3d facial expression reconstruction from videos, 2022. 2
- [12] Simon Giebenhain, Tobias Kirschstein, Markos Georgopoulos, Martin Rünz, Lourdes Agapito, and Matthias Nießner. Learning neural parametric head models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21003–21012, 2023. 1, 2, 3, 5
- [13] Simon Giebenhain, Tobias Kirschstein, Markos Georgopoulos, Martin Rünz, Lourdes Agapito, and Matthias Nießner. Mononphm: Dynamic head reconstruction from monocular videos. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 3, 4, 5, 6, 8
- [14] Simon Giebenhain, Tobias Kirschstein, Martin Rünz, Lourdes Agapito, and Matthias Nießner. Npga: Neural parametric gaussian avatars. In *SIGGRAPH Asia 2024 Conference Papers (SA Conference Papers '24), December 3-6, Tokyo, Japan*, 2024. 1, 3
- [15] Simon Giebenhain, Tobias Kirschstein, Martin Rünz, Lourdes Agapito, and Matthias Nießner. Pixel3dmm: Versatile screen-space priors for single-image 3d face reconstruction, 2025. 1, 2, 4, 5, 6, 7, 8
- [16] Philip-William Grassal, Malte Prinzler, Titus Leistner, Carsten Rother, Matthias Nießner, and Justus Thies. Neural head avatars from monocular rgb videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18653–18664, 2022. 2
- [17] Jianzhu Guo, Xiangyu Zhu, Yang Yang, Fan Yang, Zhen Lei, and Stan Z Li. Towards fast, accurate and stable 3d dense face alignment. In *European Conference on Computer Vision*, pages 152–168. Springer, 2020. 2
- [18] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *SIGGRAPH 2024 Conference Papers*. Association for Computing Machinery, 2024. 3, 8
- [19] *A 3D Face Model for Pose and Illumination Invariant Face Recognition*, Genova, Italy, 2009. IEEE. 2, 3
- [20] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. 3
- [21] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2014. cite arxiv:1412.6980Comment: Published as a conference paper at the 3rd International Conference for Learning Representations, San Diego, 2015. 5
- [22] Tobias Kirschstein, Simon Giebenhain, and Matthias Nießner. Diffusionavatars: Deferred diffusion for high-fidelity 3d head avatars. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5481–5492, 2024. 2, 3
- [23] Tobias Kirschstein, Javier Romero, Artem Sevastopolsky, Matthias Nießner, and Shunsuke Saito. Avat3r: Large animatable gaussian reconstruction model for high-fidelity 3d head avatars. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12089–12100, 2025. 1

- [24] Junxuan Li, Chen Cao, Gabriel Schwartz, Rawal Khirodkar, Christian Richardt, Tomas Simon, Yaser Sheikh, and Shunsuke Saito. Uravatar: Universal relightable gaussian codec avatars. In *ACM SIGGRAPH 2024 Conference Papers*, 2024. 1
- [25] Tianye Li, Timo Bolkart, Michael J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 36(6):194:1–194:17, 2017. 2, 3, 6
- [26] William E. Lorensen and Harvey E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, page 163–169, New York, NY, USA, 1987. Association for Computing Machinery. 3
- [27] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 3
- [28] Ali Mollahosseini, Behzad Hasani, and Mohammad H Mahoor. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *IEEE Transactions on Affective Computing*, 10(1):18–31, 2017. 7
- [29] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 2, 8
- [30] Pablo Palafox, Aljaž Božič, Justus Thies, Matthias Nießner, and Angela Dai. Npms: Neural parametric models for 3d deformable shapes. *arXiv preprint arXiv:2104.00702*, 2021. 3
- [31] Gecer Papaioannou, Athanasios an Baris, Shiyang Cheng, Grigorios G. Chrysos, Jiankang Deng, Eftychia Fotiadou, Christos Kampouris, Dimitrios Kollias, Stylianos Moschoglou, Kritaphat Songsri-In, Stylianos Ploumpis, George Trigeorgis, Panagiotis Tzirakis, Evangelos Ververas, Yuxiang Zhou, Allan Ponniah, Anastasios Roussos, and Stefanos Zafeiriou. Mimicme: A large scale diverse 4d database for facial expression analysis. In *Proceedings of the European Conference on Computer Vision*, 2022. 5
- [32] Malte Prinzler, Egor Zakharov, Vanessa Sklyarova, Berna Kabadayi, and Justus Thies. Joker: Conditional 3d head synthesis with extreme facial expressions. *arXiv preprint arXiv:2410.16395*, 2024. 2
- [33] Shenhan Qian, Tobias Kirschstein, Liam Schoneveld, Davide Davoli, Simon Giebenhain, and Matthias Nießner. Gaussianavatars: Photorealistic head avatars with rigged 3d gaussians. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20299–20309, 2024. 1, 2
- [34] George Retsinas, Panagiotis P. Filntisis, Radek Danecsek, Victoria F. Abrevaya, Anastasios Roussos, Timo Bolkart, and Petros Maragos. 3d facial expressions through analysis-by-neural-synthesis. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 2, 3, 7
- [35] Elad Richardson, Matan Sela, and Ron Kimmel. 3d face reconstruction by learning from synthetic data. In *2016 fourth international conference on 3D vision (3DV)*, pages 460–469. IEEE, 2016. 2
- [36] Elad Richardson, Matan Sela, Roy Or-El, and Ron Kimmel. Learning detailed face reconstruction from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1259–1268, 2017. 2
- [37] Radu Alexandru Rosu, Keyu Wu, Yao Feng, Youyi Zheng, and Michael J. Black. DiffLocks: Generating 3d hair from a single image using diffusion models. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2025. 8
- [38] Fatemeh Saleh, Sadeh Aliakbarian, Charlie Hewitt, Lohit Petikam, Xiao-Xian, Antonio Criminisi, Thomas J. Cashman, and Tadas Baltrušaitis. DAVID: Data-efficient and accurate vision models from synthetic data, 2025. 5
- [39] Soubhik Sanyal, Timo Bolkart, Haiwen Feng, and Michael Black. Learning to regress 3D face shape and expression from an image without 3D supervision. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 7763–7772, 2019. 1, 2, 5, 7
- [40] Liam Schoneveld, Zhe Chen, Davide Davoli, Jiapeng Tang, Saimon Terazawa, Ko Nishino, and Matthias Nießner. Sheap: Self-supervised head geometry predictor learned via 2d gaussians. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14162–14172, 2025. 2, 3, 4, 5, 6, 7
- [41] Felix Taubner, Prashant Raina, Mathieu Tuli, Eu Wern Teh, Chul Lee, and Jinmiao Huang. 3D face tracking from 2D video through iterative dense UV to image flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1227–1237, 2024. 1, 2, 5, 6, 7
- [42] Felix Taubner, Ruihang Zhang, Mathieu Tuli, and David B Lindell. Cap4d: Creating animatable 4d portrait avatars with morphable multi-view diffusion models. *arXiv preprint arXiv:2412.12093*, 2024. 2
- [43] Felix Taubner, Ruihang Zhang, Mathieu Tuli, Sherwin Bahmani, and David B. Lindell. MVP4D: Multi-view portrait video diffusion for animatable 4D avatars, 2025. 2
- [44] Ayush Tewari, Michael Zollhofer, Hyeonwoo Kim, Pablo Garrido, Florian Bernard, Patrick Perez, and Christian Theobalt. Mofa: Model-based deep convolutional face auto-encoder for unsupervised monocular reconstruction. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 1274–1283, 2017. 2
- [45] Justus Thies, Michael Zollhofer, Marc Stamminger, Christian Theobalt, and Matthias Nießner. Face2face: Real-time face capture and reenactment of rgb videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2387–2395, 2016. 1, 2
- [46] Anh Tuan Tran, Tal Hassner, Iacopo Masi, Eran Paz, Yuval Nirkin, and Gérard Medioni. Extreme 3d face reconstruction: Seeing through occlusions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3935–3944, 2018. 2
- [47] Anh Tuan Tran, Tal Hassner, Iacopo Masi, and Gérard Medioni. Regressing robust and discriminative 3d mor-

- phable models with a very deep neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5163–5172, 2017. [2](#)
- [48] Daoye Wang, Prashanth Chandran, Gaspard Zoss, Derek Bradley, and Paulo Gotardo. Morf: Morphable radiance fields for multiview neural head modeling. In *ACM SIGGRAPH 2022 Conference Proceedings*, New York, NY, USA, 2022. Association for Computing Machinery. [3](#)
- [49] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. Vggt: Visual geometry grounded transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025. [4](#), [8](#)
- [50] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *NeurIPS*, 2021. [4](#)
- [51] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024. [4](#)
- [52] Erroll Wood, Tadas Baltrušaitis, Charlie Hewitt, Matthew Johnson, Jingjing Shen, Nikola Milosavljević, Daniel Wilde, Stephan Garbin, Toby Sharp, Ivan Stojiljković, et al. 3d face reconstruction with dense landmarks. In *European Conference on Computer Vision*, pages 160–177. Springer, 2022. [1](#), [2](#)
- [53] Yuelang Xu, Lizhen Wang, Zerong Zheng, Zhaoqi Su, and Yebin Liu. 3d gaussian parametric head model. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2024. [3](#)
- [54] Tarun Yenamandra, Ayush Tewari, Florian Bernard, Hans-Peter Seidel, Mohamed Elgharib, Daniel Cremers, and Christian Theobalt. i3dmm: Deep implicit 3d morphable model of human heads. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12803–12813, 2021. [2](#), [3](#)
- [55] Jianhui Yu, Hao Zhu, Liming Jiang, Chen Change Loy, Weidong Cai, and Wayne Wu. CelebV-Text: A large-scale facial text-video dataset. In *CVPR*, 2023. [5](#)
- [56] Tianke Zhang, Xuangeng Chu, Yunfei Liu, Lijian Lin, Zhendong Yang, Zhengzhuo Xu, Chengkun Cao, Fei Yu, Changyin Zhou, Chun Yuan, et al. Accurate 3d face reconstruction with facial component tokens. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9033–9042, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [57] Mingwu Zheng, Hongyu Yang, Di Huang, and Liming Chen. Imface: A nonlinear 3d morphable face model with implicit neural representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20343–20352, 2022. [2](#), [3](#)
- [58] Mingwu Zheng, Haiyu Zhang, Hongyu Yang, Liming Chen, and Di Huang. Imface++: A sophisticated nonlinear 3d morphable face model with implicit neural representations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. [3](#)
- [59] Xiaozheng Zheng, Chao Wen, Zhaohu Li, Weiyi Zhang, Zhuo Su, Xu Chang, Yang Zhao, Zheng Lv, Xiaoyuan Zhang, Yongjie Zhang, Guidong Wang, and Xu Lan. Headgap: Few-shot 3d head avatar via generalizable gaussian priors. *arXiv preprint arXiv:2408.06019*, 2024. [3](#)
- [60] and Cecilia Zhang Zheng Ding, Zhihao Xia, Lars Jebe, Zhuowen Tu, and Xiuming Zhang. Diffusionrig: Learning personalized priors for facial appearance editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. [2](#)
- [61] Hao Zhu, Wayne Wu, Wentao Zhu, Liming Jiang, Siwei Tang, Li Zhang, Ziwei Liu, and Chen Change Loy. CelebV-HQ: A large-scale video facial attributes dataset. In *ECCV*, 2022. [5](#)
- [62] Hao Zhu, Haotian Yang, Longwei Guo, Yidi Zhang, Yanru Wang, Mingkai Huang, Menghua Wu, Qiu Shen, Ruigang Yang, and Xun Cao. Facescape: 3d facial dataset and benchmark for single-view 3d face reconstruction. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2023. [5](#)
- [63] Wojciech Zielonka, Timo Bolkart, and Justus Thies. Towards metrical reconstruction of human faces. In *European conference on computer vision*, pages 250–269. Springer, 2022. [1](#), [2](#), [5](#), [6](#)