

# ImFace: A Nonlinear 3D Morphable Face Model with Implicit Neural Representations

Mingwu Zheng<sup>1,2</sup>, Hongyu Yang<sup>3</sup>, Di Huang<sup>1,2,\*</sup>, Liming Chen<sup>4</sup>

<sup>1</sup> State Key Laboratory of Software Development Environment, Beihang University, China

<sup>2</sup> School of Computer Science and Engineering, Beihang University, China

<sup>3</sup> Institute of Artificial Intelligence, Beihang University, China

<sup>4</sup> LIRIS, École Centrale de Lyon, France

{zhengmingwu, hongyuyang, dhuang}@buaa.edu.cn, liming.chen@ec-lyon.fr

## Abstract

Precise representations of 3D faces are beneficial to various computer vision and graphics applications. Due to the data discretization and model linearity, however, it remains challenging to capture accurate identity and expression clues in current studies. This paper presents a novel 3D morphable face model, namely ImFace, to learn a nonlinear and continuous space with implicit neural representations. It builds two explicitly disentangled deformation fields to model complex shapes associated with identities and expressions, respectively, and designs an improved learning strategy to extend embeddings of expressions to allow more diverse changes. We further introduce a Neural Blend-Field to learn sophisticated details by adaptively blending a series of local fields. In addition to ImFace, an effective pre-processing pipeline is proposed to address the issue of watertight input requirement in implicit representations, enabling them to work with common facial surfaces for the first time. Extensive experiments are performed to demonstrate the superiority of ImFace.

## 1. Introduction

3D Morphable Face Models (3DMMs) are well-reputed statistical models, established by learning techniques upon prior distributions of facial shapes and textures from a set of samples with dense correspondence, aiming at rendering realistic faces of a high variety. Since a morphable representation is unique across different downstream tasks where the geometry and appearance are separately controllable, 3DMMs are pervasively exploited in many face analysis applications in the field of computer vision, computer graphics, biometrics, and medical imaging [2, 8, 27, 50].

In 3DMMs, the most fundamental issue lies in the way to generate latent morphable representations, and during the

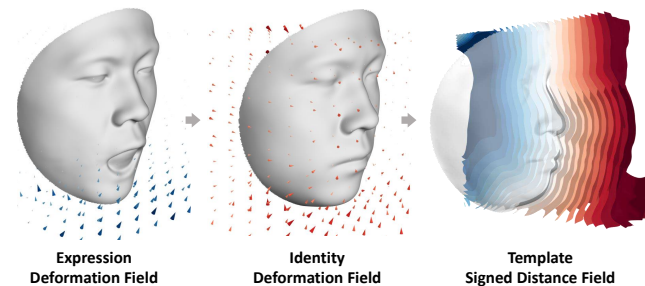


Figure 1. ImFace encodes complex face variations by two explicitly disentangled deformation fields with respect to a template face, resulting in a morphable implicit representation for 3D faces.

past two decades, along with data improvement in scale, diversity and quality [10, 13, 31, 58], remarkable progresses have been achieved. The methods are initially linear model based [7, 39, 40] and further extended to multilinear model based [9, 12, 55], where different modes are individually encoded. Unfortunately, for the relatively limited representation ability of linear models, these methods are not so competent at handling the cases with complicated variations, e.g. exaggerated expressions. In the context of deep learning, a number of nonlinear models have been investigated with the input of 2D images [53, 54] or 3D meshes [5, 11, 15, 17, 46] by using Convolutional Neural Networks or Graph Neural Networks. They indeed deliver performance gains; however, restricted by the resolution of discrete representing strategies on input data, facial priors are not sufficiently captured, incurring loss of shape details. Besides, all current methods are dependent on the preposed procedure of point-to-point correspondence [1, 6, 25, 34], but face registration itself remains challenging.

Recently, several studies on Implicit Neural Representations (INRs) [16, 26, 32, 35, 38] have shown that 3D geometries can be precisely modeled by learning continuous deep implicit functions. They describe an input observation as a low-dimensional shape embedding and estimate the Signed

\*Corresponding author.

Distance Function (SDF) or the occupancy value of a query point so that the surface of an arbitrary resolution and topology can be defined by an isocontour. Due to the continuous parameterization and consistent representation, INRs prove superior to the discrete voxels, pointclouds and meshes, and report decent results in shape reconstruction [24, 56, 57, 61] and surface registration [21, 33, 62]. Such an advantage suggests an alternative to 3DMM that can fulfill accurate correspondence and fine-grained modeling in a unified network. Nevertheless, unlike the objects with apparent shape differences and limited non-rigid variations such as indoor scenes and human bodies, all facial surfaces look very similar but include more complex deformations, where multiple identities and rich expressions deeply interweave with each other, making current INR methods problematic in face modeling, as evidenced by the preliminary attempt [59]. Another difficulty is that implicit functions primarily require watertight input, which is not friendly to facial surfaces.

This paper proposes a novel 3D face morphable model, namely ImFace, which substantially upgrades conventional 3DMMs by learning INRs. To capture nonlinear facial geometry changes, ImFace builds separate INR sub-networks to explicitly disentangle shape morphs into two deformation fields for identity and expression respectively (as Fig. 1 shows), and an improved embedding learning strategy is introduced to extend the latent space of expressions to allow more diverse details. In this way, inter-individual differences and fine-grained deformations can be accurately modeled, which simultaneously takes into account the flexibility when applied to related tasks. Furthermore, inspired by linear blend skinning [30], a Neural Blend-Field is presented to decompose the entire facial deformation or geometry into semantically meaningful regions encoded by a set of local implicit functions and adaptively blend them through a lightweight module, leading to more sophisticated representations with reduced parameters. Besides, a new preprocessing pipeline is designed, which bypasses the need of watertight face data as in existing SDF-based INR models and works well for various facial surfaces, *i.e.*, either hardware-acquired or artificially synthesized.

In summary, the main contributions of this study include:

- We propose a novel INR-based 3DMM, which encodes complex face shape variations by two explicitly disentangled deformation fields, learning powerful representations in a fine-grained and semantically meaningful manner.
- We present an effective preprocessing pipeline, which defines a general SDF for non-watertight 3D faces, enabling INRs to work with them for the first time.
- We experimentally demonstrate that ImFace has the advantage in synthesizing high-quality 3D faces with plausible details, outperforming the state-of-the-art counterparts in 3D face reconstruction.

## 2. Related Work

**3D Morphable Face Models.** 3DMMs were first proposed by Blanz and Vetter [7] as a general face representation. A known template mesh was registered to all training scans by utilizing Non-rigid Iterative Closest Point algorithm (NICP) [4], and Principal Components Analysis (PCA) was used to span prior face distributions. To model identity-dependent expressions, 3DMMs were further extended to multilinear models [9, 12, 55]. After then, great advancements have been made with data improvement [10, 13, 31, 58]. FLAME [31] was an expressive model controlling facial expressions by combining jaw articulation with linear expression blend-shapes. It was learned from a large 3D face dataset including D3DFACS [20] and delivered more impressive results than ever, whereas nonlinear facial deformations cannot be well captured.

The development of deep networks has catalyzed more powerful nonlinear 3DMMs. A number of models were learned from 2D images [44, 53, 54], but they mostly lacked high fidelity and fine details due to the low resolution of input images in this ill-posed inverse problem. To better leverage 3D face scans, Bagautdinov *et al.* [5] mapped 3D meshes to the 2D space and more studies [11, 15, 17, 46] directly learned 3DMMs from meshes by spectral or spiral convolutions. These neural networks were established upon discrete 3D representations, thus limiting the performance in the presence of complex deformations.

Please refer to [22] for a more comprehensive report on 3DMMs.

**Implicit Neural Representations.** Recently, implicit neural functions have emerged as a more effective and suitable representation of 3D geometry [16, 26, 32, 35, 38, 49] as they model shapes continuously without discretization. To preserve fine details, the structured local features divided by shape elements [23], grids [28, 42] or octrees [51, 52] were further exploited. Moreover, to well capture shape variations and correspondence relationships, an additional implicit deformation latent space was specially learned [21, 62]. However, high visual fidelity and variety could hardly be simultaneously achieved by existing techniques and they are therefore not so qualified to morphable face modeling. In addition, current studies on INRs mostly focused on the watertight input, such as the ones in ShapeNet [60], and a number of methods on watertight human heads or bodies were proposed accordingly [3, 14, 19, 43, 45, 47, 48, 59]. Among them, H3D-Net [45] learned an implicit head shape space for 2D reconstruction, which was not a generic model. i3DMM [59] was the first implicit 3D morphable model designed for human heads. However, it severely suffered from a low quality in representing facial regions. To overcome the limitation of watertightness, [18] learned an Unsigned Distance Function (UDF) to handle open surfaces, but the reconstructed results were not adequately plausible.

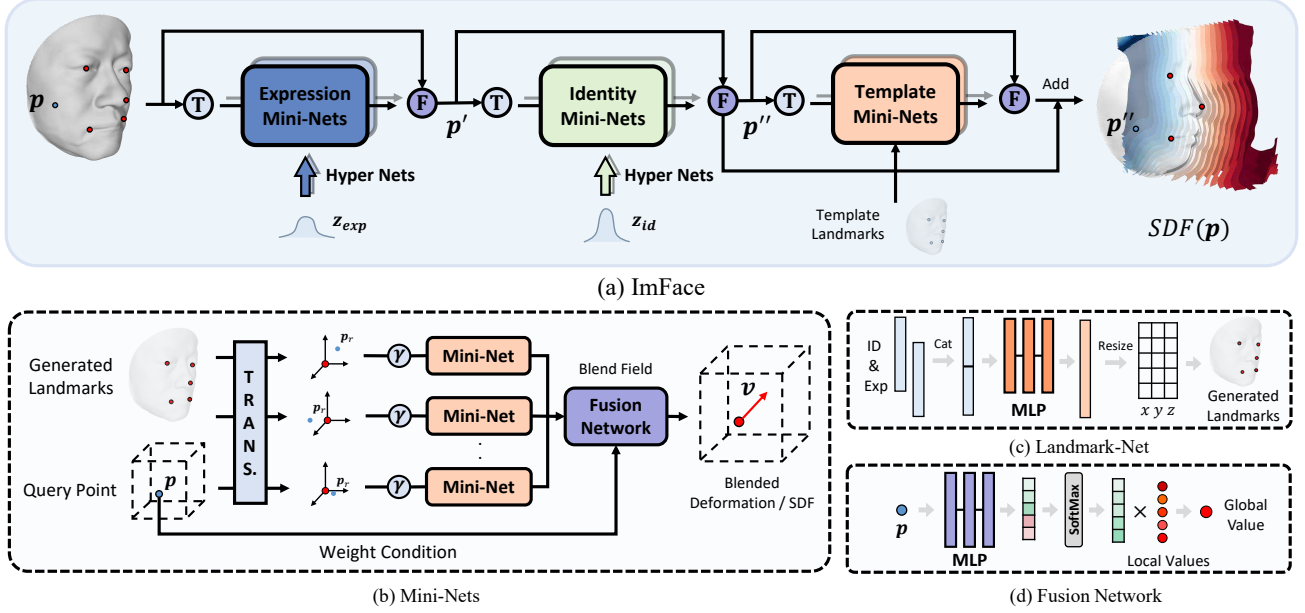


Figure 2. **ImFace overview.** (a) The proposed network consists of three Mini-Nets blocks to explicitly disentangle shape morphs into separate deformation fields, where the Expression and Identity Mini-Nets blocks are associated with expression and identity deformations, respectively, and the Template Mini-Nets block learns the SDF of a template face space. (b) The Mini-Nets block is a shared architecture, which decomposes an entire facial feature into semantically meaningful parts and encodes them by a set of local field functions. It is tailed by a Fusion Network for more comprehensive representations. (c) The Landmark-Net is introduced to softly partition the entire facial surface. (d) The Fusion Network is a lightweight module conditioned on the query point position, which adaptively blends the local field functions, resulting in an elaborate Neural Blend-Field.

### 3. Method

We take advantages of INRs to learn a nonlinear 3D morphable face model. The proposed ImFace explicitly disentangles facial shape morphs into two separate deformation fields associated with identity and expression, respectively, and a deep SDF is learned to represent the template shape. All the fields are blended with a series of local implicit functions for more detailed representations.

#### 3.1. Disentangled INRs Network

The fundamental idea of INRs is to train a neural network to fit a continuous function  $f$ , which implicitly represents surfaces through level-sets. The function can be defined in various formats, *e.g.* occupancies [35], SDF [38], or UDF [18]. We exploit a deep SDF conditioned on the latent embeddings of both expression and identity for comprehensive face representations. It outputs the signed distance  $s$  from a query point:

$$f : (\mathbf{p}, \mathbf{z}_{exp}, \mathbf{z}_{id}) \in \mathbb{R}^3 \times \mathbb{R}^{d_{exp}} \times \mathbb{R}^{d_{id}} \mapsto s \in \mathbb{R}, \quad (1)$$

where  $\mathbf{p} \in \mathbb{R}^3$  is the coordinate of the query point in the 3D space,  $\mathbf{z}_{exp}$  and  $\mathbf{z}_{id}$  denote the expression and identity embeddings, respectively.

Our goal is to learn a neural network to parameterize  $f$ , making it satisfy the genuine facial shape priors. As

shown in Fig. 2, the proposed network for Imface is composed of three Mini-Nets blocks, which explicitly disentangles the learning process of facial shape morphs, ensuring that inter-individual differences and fine-grained deformations can be accurately modeled. In particular, the first two Mini-Nets blocks learn separate deformation fields associated with expression- and identity-variation, respectively, and the Template Mini-Nets block learns a signed distance field of a template face shape.

All the fields above are implemented by a shared Mini-Nets architecture, where the entire facial deformation or geometry is further decomposed into a number of semantically meaningful parts and encoded by a set of local field functions, so that rich details can be sufficiently captured. A lightweight module conditioned on the query point position, *i.e.*, Fusion Network, is stacked at the end of the Mini-Nets block to adaptively blend the local fields. As such, an elaborate Neural Blend-Field is achieved. The three core components of ImFace work for different purposes and their structures are slightly changed accordingly. We briefly describe them as follows:

**Expression Mini-Nets (ExpNet).** The facial deformations incurred by expressions are represented by ExpNet  $\mathcal{E}$ , which learns an observation-to-canonical warping for every face scan:

$$\mathcal{E} : (\mathbf{p}, \mathbf{z}_{exp}, l) \mapsto \mathbf{p}' \in \mathbb{R}^3, \quad (2)$$

where  $l \in \mathbb{R}^{k \times 3}$  denotes  $k$  3D landmarks on an observed face generated by a Landmark-Net  $\eta : (\mathbf{z}_{exp}, \mathbf{z}_{id}) \mapsto l$ , introduced to localize the query point  $\mathbf{p}$  in the Neural Blend-Field. The point  $\mathbf{p}$  in the observation space is deformed by  $\mathcal{E}$  to a new point  $\mathbf{p}'$  in the person-specific canonical space, which represents faces with a neutral expression.

**Identity Mini-Nets (IDNet).** To model shape morphs among individuals, IDNet  $\mathcal{I}$  further warps the canonical space to a template shape space shared by all faces:

$$\mathcal{I} : (\mathbf{p}', \mathbf{z}_{id}, l') \mapsto (\mathbf{p}'', \delta) \in \mathbb{R}^3 \times \mathbb{R}, \quad (3)$$

where  $l' \in \mathbb{R}^{k \times 3}$  denotes  $k$  landmarks on the canonical face generated by another Landmark-Net conditioned only on the identity embedding  $\eta' : \mathbf{z}_{id} \mapsto l'$ , and  $\mathbf{p}''$  is the deformed point in the template space. To cope with the possible non-existent correspondences generated during preprocessing,  $\mathcal{I}$  additionally predicts a residual term  $\delta \in \mathbb{R}$  to correct the predicted SDF value  $s_0$ , which is similar to [21].

**Template Mini-Nets (TempNet).** TempNet  $\mathcal{T}$  learns a signed distance field of the shared template face:

$$\mathcal{T} : (\mathbf{p}'', l'') \mapsto s_0 \in \mathbb{R}, \quad (4)$$

where  $l'' \in \mathbb{R}^{k \times 3}$  denotes  $k$  landmarks on the template face, which is averaged on the whole training set, and  $s_0$  is the uncorrected SDF value. The final SDF value of a query point is calculated via  $s = s_0 + \delta$ , and the ImFace model can be ultimately formulated as:

$$f(\mathbf{p}) = \mathcal{T}(\mathcal{I}_{\mathbf{p}''}(\mathcal{E}(\mathbf{p}, \mathbf{z}_{exp}), \mathbf{z}_{id})) + \mathcal{I}_{\delta}(\mathcal{E}(\mathbf{p}, \mathbf{z}_{exp}), \mathbf{z}_{id}). \quad (5)$$

The proposed ImFace learns face morphs by disentangled deformation fields in a fine-grained and meaningful manner, ensuring that more diverse and sophisticated facial deformations can be accurately learned. We detailedly introduce the architectures of the main modules, learning strategy, training critics and data preprocessing pipeline in the subsequent.

### 3.2. Neural Blend-Field

The Mini-Nets block is a common architecture shared by the three sub-networks  $\mathcal{E}$ ,  $\mathcal{I}$ , and  $\mathcal{T}$ . It learns a continuous field function  $\psi : \mathbf{x} \in \mathbb{R}^3 \mapsto v$ , to produce a Neural Blend-Field for comprehensive face representations. In particular, to overcome the limited expressivity of a single network, we decompose a face space into a set of semantically meaningful local regions, and learn  $v$  (e.g. deformation or signed distance value) individually before blending. Such design is inspired by the recent INRs study [41] on human body, which introduces the linear blend skinning algorithm [30] to make the network learn from separate transformations of body parts. To better represent detailed facial surfaces, we replace the constant transformation term in the original linear blend skinning algorithm with  $\psi_n(\mathbf{x} - l_n)$ , and define the Neural Blend-Field as:

$$v = \psi(\mathbf{x}) = \sum_{n=1}^k w_n(\mathbf{x}) \psi_n(\mathbf{x} - l_n), \quad (6)$$

where  $l_n$  is a parameter that describes the  $n$ -th local region,  $w_n(\mathbf{x})$  is the  $n$ -th blend weight, and  $\psi_n(\mathbf{x} - l_n)$  is the corresponding local field. In this way, the blending is performed on a series of local fields, rather than calculating a weighted average of the output values  $v$  of some fixed positions, leading to more powerful representation capability in handling complicated local features.

Specifically, five landmarks located at the outer eye corners, mouth corners, and nose tip are utilized to describe the local regions ( $l_n \in \mathbb{R}^3$ ) $_{n=1}^5$ , and each region is assigned a tiny MLP with sinusoidal activations [49] to generate the local field, denoted as  $\psi_n$ . To capture high-frequency local variations, we leverage sinusoidal positional encoding  $\gamma$  [36] on the coordinate  $\mathbf{x} - l_n$ . At the end of a Mini-Nets block, a lightweight Fusion Network conditioned on the absolute coordinate of input  $\mathbf{x}$  is equipped, which is implemented by a 3-layer MLP with softmax to predict the blend weights ( $w_n \in \mathbb{R}^+$ ) $_{n=1}^5$ .

**Deformation Formulation.** We formulate the deformation with a SE(3) field  $(\boldsymbol{\omega}, \mathbf{v}) \in \mathbb{R}^6$ , where  $\boldsymbol{\omega} \in so(3)$  is a rotate vector representing the screw axis and the angle of rotation. The deformed coordinates  $\mathbf{x}'$  can be calculated by  $e^{\boldsymbol{\omega}} \mathbf{x} + \mathbf{t}$ , where the rotation matrix  $e^{\boldsymbol{\omega}}$  (exponential map form of Rodrigues' formula) is written as:

$$e^{\boldsymbol{\omega}} = \mathbf{I} + \frac{\sin \|\boldsymbol{\omega}\|}{\|\boldsymbol{\omega}\|} \boldsymbol{\omega}^\wedge + \frac{1 - \cos \|\boldsymbol{\omega}\|}{\|\boldsymbol{\omega}\|^2} (\boldsymbol{\omega}^\wedge)^2, \quad (7)$$

and the translation  $\mathbf{t}$  is formulated as:

$$\mathbf{t} = \left[ \mathbf{I} + \frac{1 - \cos \|\boldsymbol{\omega}\|}{\|\boldsymbol{\omega}\|^2} \boldsymbol{\omega}^\wedge + \frac{\|\boldsymbol{\omega}\| - \sin \|\boldsymbol{\omega}\|}{\|\boldsymbol{\omega}\|^3} (\boldsymbol{\omega}^\wedge)^2 \right] \mathbf{v}, \quad (8)$$

where  $\boldsymbol{\omega}^\wedge$  denotes the skew-symmetric matrix of  $\boldsymbol{\omega}$ . We exploit SE(3) to describe facial shape morphs for its superior capability in handling mandibular rotations and higher robustness to pose perturbations than the common translation deformation  $\mathbf{x}' = \mathbf{x} + \mathbf{t}$ .

**Hyper Nets.** To obtain a more compact and expressive latent space, we then introduce a meta-learning approach [49]. A Hyper Net  $\phi_n$  is implemented by an MLP and it predicts the instance-specific parameters for ExpNet  $\mathcal{E}$  and IDNet  $\mathcal{I}$ . It takes a latent code  $\mathbf{z}$  as input and generates the parameters for the neurons in a Mini-Net  $\psi_n$  so that the learned facial representations possess a higher variety.

### 3.3. Improved Expression Embedding Learning

The auto-decoder framework proposed by [38] has been widely adopted in INRs to jointly learn embeddings and network parameters. In the previous attempt [59], each expression type is encoded by one embedding for attribute disentangling. Unfortunately, such an embedding is merely able to represent the average shape morph of an expression

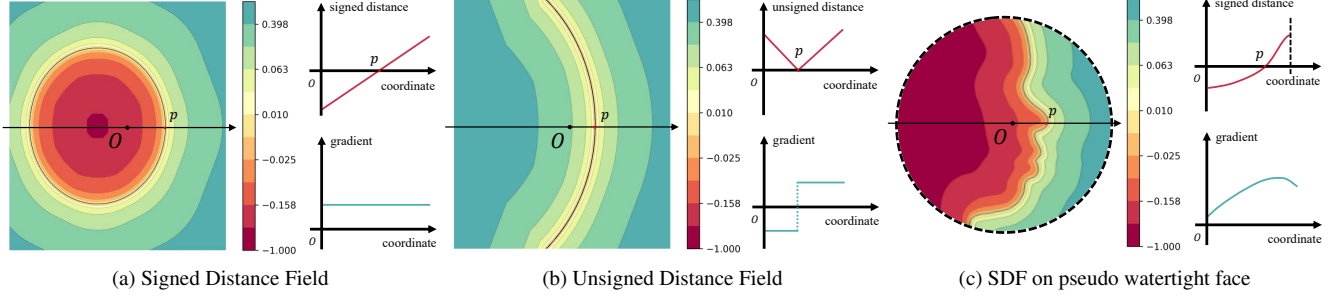


Figure 3. **(a)** SDF is able to represent closed shapes. **(b)** UDF is capable of representing an open surface but the gradient is discontinuous at the boundary, making it hard to be fitted by neural networks. **(c)** The proposed method generates pseudo watertight faces and restricts implicit functions on them, enabling implicit neural networks to learn geometry representations on 3D faces.

type, making the learned latent space fail to capture more diverse deformation details across individuals. To avoid the dilemma above, we improve the learning strategy by treating each non-neutral face scan as a unique expression and generating a specific embedding for it. In this way, the latent space is significantly extended, which enables  $\mathcal{E}$  to represent more fine-grained details. On the other side, there exists a potential failure mode that the identity properties are tangled into the expression space again, and  $\mathcal{I}$  collapses to an identity mapping. To tackle this challenge, we suppress  $\mathcal{E}$  when the current training sample is a neutral face, written as:

$$\mathcal{E}(\mathbf{p}_{nu}, \mathbf{z}_{exp}, l) \equiv \mathbf{p}_{nu}, \quad (9)$$

where  $\mathbf{p}_{nu}$  denotes a point from a neutral face. By applying such a learning strategy,  $\mathcal{I}$  and  $\mathcal{T}$  jointly learn shape representations on neutral faces, and  $\mathcal{E}$  focuses only on expression deformations. Moreover, only neutral labels are required during training, bypassing the dense expression labels.

### 3.4. Loss Functions

ImFace is trained with several loss functions to learn plausible facial shape representations and dense correspondence.

**Reconstruction Loss.** The basic SDF structure loss is applied to learn implicit fields:

$$\mathcal{L}_{sdf}^i = \lambda_1 \sum_{\mathbf{p} \in \Omega_i} |f(\mathbf{p}) - \bar{s}| + \lambda_2 \sum_{\mathbf{p} \in \Omega_i} (1 - \langle \nabla f(\mathbf{p}), \bar{\mathbf{n}} \rangle), \quad (10)$$

where  $\bar{s}$  and  $\bar{\mathbf{n}}$  denote the ground-truth SDF values and the field gradients, respectively.  $\Omega_i$  is the sampling space of the face scan  $i$ , and  $\lambda$  indicates the trade-off parameter.

**Eikonal Loss.** To obtain reasonable fields throughout the network, multiple Eikonal losses are used to enforce the L-2 norm of spatial gradients to be unit:

$$\mathcal{L}_{eik}^i = \lambda_3 \sum_{\mathbf{p} \in \Omega_i} (|\|\nabla f(\mathbf{p})\| - 1| + |\|\nabla \mathcal{T}(\mathcal{I}(\mathbf{p}'))\| - 1|), \quad (11)$$

where  $\mathcal{L}_{eik}^i$  enables the network to satisfy the Eikonal constraint [26] in the observation and canonical spaces simulta-

neously, which also contributes to a reasonable correspondence along face deformations at all the network stages.

**Embedding Loss.** It regularizes the embedding with a zero-mean Gaussian prior:

$$\mathcal{L}_{emb}^i = \lambda_4 (\|\mathbf{z}_{exp}\|^2 + \|\mathbf{z}_{id}\|^2). \quad (12)$$

**Landmark Generation Loss.** The  $l_1$ -loss is used to learn the Landmark-Nets  $\eta, \eta'$ :

$$\mathcal{L}_{lmkg}^i = \lambda_5 \sum_{n=1}^k (|l_n - \bar{l}_n| + |l'_n - \bar{l}'_n|), \quad (13)$$

where  $\bar{l}^i$  denotes the  $k$  labeled landmarks on sample  $i$ , and  $\bar{l}'$  denotes the landmarks on the corresponding neutral face.

**Landmark Consistency Loss.** We exploit this loss to guide the deformed landmarks to be located at the corresponding positions on the ground-truth neutral and template faces for better correspondence performance:

$$\mathcal{L}_{lmkc}^i = \lambda_6 \sum_{n=1}^k (|\mathcal{E}(l_n) - \bar{l}_n| + |\mathcal{I}(\mathcal{E}(l_n)) - l''_n|). \quad (14)$$

**Residual Constraint.** As in [21], to avoid that the residual item  $\delta$  learns too much template face information and downgrades the morphable model, we penalize  $\delta$  by:

$$\mathcal{L}_{res}^i = \lambda_7 \sum_{\mathbf{p} \in \Omega_i} |\delta(\mathbf{p})|. \quad (15)$$

The total training loss is calculated on all face samples indexed by  $i$ , finally formulated as:

$$\mathcal{L} = \sum_i (\mathcal{L}_{sdf}^i + \mathcal{L}_{eik}^i + \mathcal{L}_{emb}^i + \mathcal{L}_{lmkg}^i + \mathcal{L}_{lmkc}^i + \mathcal{L}_{res}^i). \quad (16)$$

At the testing phase, for each 3D face indexed by  $j$ , we minimize the following objective to obtain its latent embedding and the reconstructed 3D face:

$$\arg \min_{\mathbf{z}_{exp}, \mathbf{z}_{id}} \sum_j (\mathcal{L}_{sdf}^j + \mathcal{L}_{eik}^j + \mathcal{L}_{emb}^j). \quad (17)$$



Figure 4. Reconstruction comparison with i3DMM [59], FLAME [31], and FaceScape [58]. Each column corresponds to a person with a non-neutral expression. Visually inspected, ImFace captures much richer shape variations with more compact latent embeddings.

### 3.5. Data Preprocessing

Since neural networks excel in fitting functions that are differentiable everywhere, the current studies on implicit functions generally require watertight input. Although functions like UDF do not demand watertightness, they are nondifferentiable when crossing a surface and are not so competent at handling details (see Fig. 3 for an illustration). We present an effective preprocessing pipeline, which generates pseudo watertight faces and defines a general SDF on them, so that the geometry and correspondence can be learned as exquisitely as on watertight objects.

**Pseudo watertight face generation.** The faces are rigidly aligned to frontal using landmarks and each mesh is normalized to a unit of 10 *cm*. The coordinate origin is set at the point 4 *cm* behind the nose tip, and the sphere with a radius of 10 *cm* is then defined as the sampling area where the mesh triangles outside are cropped away. The Ray-Triangle Intersection Algorithm [37] is applied to remove the hidden surface such as nasal and oral cavity, and the Delaunay Triangulation Algorithm [29] is then performed on *x-y* coordinates for an oriented and pseudo watertight mesh.

**SDF computation on facial surfaces.** With the pseudo watertight faces generated, SDF values can then be computed through a distance transform on them. The sign of the samples is determined simply by the angle between its distance vector to the nearest surface and the *z*-axis positive direction. The values of coordinates behind the facial surface are defined to be negative. We uniformly sample 250,000 points on each facial surface and 15,000 points

in the sphere and calculate their signed distance and gradient vectors. The sampled data are eventually formulated as  $\{(\mathbf{p}, \bar{\mathbf{n}}, \bar{s})\}$  triplets (query point, gradient vectors, signed distance value) for ImFace training.

## 4. Experiments

We conduct extensive subjective and objective evaluations on ImFace, and ablation studies are performed to validate the specifically designed modules.

**Dataset.** FaceScape [58] is a large-scale high quality 3D face Dataset consisting of 938 individuals with 20 types of expressions. The data from 365 individuals are publicly available and we mainly use them for experiments. To be specific, 5,323 face scans from 355 persons with 15 expressions are sampled as the training set, and another 200 face scans from the remaining 10 persons with 20 expressions are used as the testing set.

**Network Architecture.** All Mini-Nets blocks  $\psi_n$  are implemented as MLPs with 3 hidden layers and 32 dimensional hidden features activated by sine. The Hyper Nets  $\phi_n$  are 3-layer MLPs activated by ReLU, where the hidden layer dimensionality is 64. Landmark-Nets  $\eta$  and  $\eta'$  have three 128 dimensional fully connected layers. Please refer to the supplementary material for more details of the networks.

**Implementation Details.** The model is trained with Adam in an end-to-end manner. We train the model for 1,500 epochs with an initial learning rate of 0.0001, and after 200 epochs, we decay it by a factor of 0.95 for every 10 epochs.

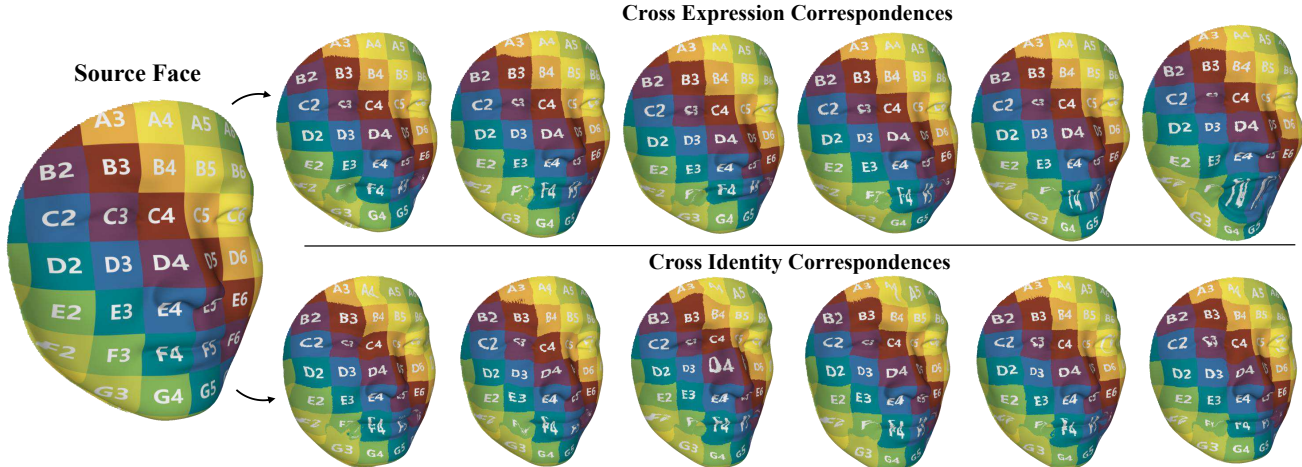


Figure 5. Correspondence results. The leftmost face is morphed into multiple expressions (upper row) and identities (bottom row).

Metrics	Dim.	Chamfer(mm) <sup>†</sup>	F-score@0.001 <sup>‡</sup>
i3DMM [59]	256	1.635	42.26
FLAME [31]	400	0.971	64.73
FaceScape [58]	352	0.929	67.09
ImFace	256	<b>0.625</b>	<b>91.11</b>

Table 1. Quantitative comparison with the state-of-the-art methods (<sup>†</sup>Lower is better; <sup>‡</sup>Higher is better).

The training phase takes around 2 days on 4 NVIDIA RTX 3090 GPUs with minibatches of size 72. During testing, it takes about 4 hours to optimize 200 samples on a single GPU.

#### 4.1. Reconstruction

We use the proposed ImFace model to fit face scans by optimizing Eq. (17) and compare the reconstruction results with FLAME [31], FaceScape [58] and the geometry model of i3DMM [59], which signify the state-of-the-art. The official code of FLAME is used to fit the full face scans in the test set, with 300 identity parameters and 100 expression parameters. For FaceScape, we use their released bilinear model built from 938 individuals for testing, where the identity and expression parameters are 300 and 52, respectively. Note that our test scans are included in the training set of FaceScape. In addition, we modify its official code to fit the full scans instead of only the landmarks for improved results. For i3DMM, since the original model is trained only on 58 individuals, we therefore re-train the model on the same training set as in ImFace for fair comparison. In both i3DMM and ImFace, the identity and expression embeddings are 128-dimensional.


**Qualitative Evaluation.** Fig. 4 visualizes the reconstruction results achieved by different models, where each column corresponds to a test person with a non-neutral expression. The results also include unseen expressions during

learning. i3DMM is the first deep implicit model for human heads, but it is less capable of capturing complicated deformations and fine-grained details under a relatively intricate circumstance, resulting in artifacts on the reconstructed faces. FLAME is able to well present identity characteristics, but is not so competent at dealing with nonlinear deformations, delivering stiff facial expressions. FaceScape performs more favourably mainly due to the high-quality training scans and the test faces are included in the training set, but it still cannot precisely present expression morphs. Comparatively, ImFace reconstructs faces with more accurate identity and expression properties, and it is able to preserve subtle and rich nonlinear facial muscle deformations such as frowns and pouts by fewer latent parameters.

**Quantitative Evaluation.** To make fair comparison, all faces are processed as described in Sec. 3.5 to remove the inner structures like eyeballs, so that the quantitative metrics can be computed in the same facial region for all the models. Specifically, the symmetric Chamfer distance and F-score are used as metrics, and the threshold of F-score is set to 0.001 as a strict standard. The results are shown in Table 1. As we can see, ImFace exceeds the counterparts by a large margin under both the metrics, which clearly validates its effectiveness.

#### 4.2. Correspondence

In contrast to existing methods that generally require accurate face registration, correspondences can be automatically learned in INRs models. We further design training critic to enhance such a feature, and this evaluation aims to confirm it. Given two 3D faces, we use ImFace to fit them and deform the dense sampled points to the template space so that point-to-point correspondences can be fulfilled by nearest neighbor search. Fig. 5 visualizes some results generated by our method, where we manually paint color patterns on the shapes to better check the quality. In gen-



(a) w/o dist. (b) w/o blend (c) w/o extend. (d) ImFace (e) GT

Figure 6. Qualitative ablation study results.

Metrics	Chamfer ( $mm$ ) <sup>†</sup>	F-score@0.001 <sup>‡</sup>
Ours w/o dist.	0.772	82.70
Ours w/o blend	0.767	82.37
Ours w/o extend.	0.705	86.98
ImFace	<b>0.625</b>	<b>91.11</b>

Table 2. Quantitative ablation study results.

eral, ImFace is able to establish pleasing correspondences across various expressions and identities. Meanwhile, it can be inspected that tiny internal texture dispersion indeed occasionally occurs around mouth corners, and this is mainly because facial shapes change drastically in these local areas under different expressions, which is also very hard case to specialized correspondence methods.

### 4.3. Ablation Study

ImFace is built on the following core components: disentangled deformation fields (dist.), Neural Blend-Field (blend), and improved expression embedding learning (extend.). We experimentally verify the credits of such designs.

**On Disentangled Deformation Fields.** To highlight the disentangled deformation learning process, we build a baseline network which contains only one deformation field to learn face shape morphs universally. Accordingly,  $\mathbf{z}_{exp}$  and  $\mathbf{z}_{id}$  are concatenated as the input of the hyper net. Fig. 6 (a) provides a demonstration. In spite of some fine-grained details brought by other designs, there exists a chaos on the reconstructed faces, especially for the ones with large expressions. The quantitative results in Table 2 also indicate the significance of decoupled deformation learning.

**On Neural Blend-Field.** We replace the Neural Blend-Field in  $\mathcal{E}$ ,  $\mathcal{I}$ ,  $\mathcal{T}$  with vanilla MLPs of the same amount of parameters, which directly predict global deformations or SDF values of an entire face. As shown in Fig. 6 (b), a visible blur appears due to the limited capability in learning

high-frequency details. The quantitative results in Table 2 confirm the necessity of Neural Blend-Field in learning sophisticated representations.

**On Improved Embedding Learning.** This strategy is introduced to learn more diverse and fine-grained facial deformations. As shown in Fig. 6 (c), when restricting the number of expression embeddings to be the same as expression categories, the generated expressions tend to be average. Moreover, for exaggerated expressions, such as mouth stretch, the counterpart models can hardly converge to a reasonable state.

## 5. Discussion

**Limitations.** Despite the significant advance in 3D facial shape representation, we mainly focus on face geometry modeling, whereas facial texture is less considered. Indeed, a basic texture model can be achieved by plugging a color field, but a more comprehensive INR-based 3DMM which presents the facial appearance with realistic diffuse and specular albedo remains to be explored.

**Societal Impact.** Our model aims at high-quality face modeling, and similar to existing models, it has the potential to be applied to downstream scenarios such as 2D reconstruction and face animation, which may result in unethical practices like privacy invasion or identity fraud. We encourage researchers and developers to consider the questions, such as how to prevent personal face data from being maliciously accessed, before applying the model to real world.

## 6. Conclusion

This paper presents a novel nonlinear 3D morphable face model with INRs, *i.e.* ImFace, which learns complex facial shape variations by two explicitly disentangled deformation fields associated with expression and identity, respectively, along with an improved embedding learning strategy to allow more fine-grained expressions. To precisely capture detailed facial deformations and geometries, it further presents a Neural Blend-Field. In addition, an effective preprocessing pipeline is presented, which enables INRs to work with non-watertight facial surfaces for the first time. Experiments show that ImFace is competent for this issue and outperforms the state-of-the-art counterparts.

## Acknowledgment

This work is partly supported by the National Natural Science Foundation of China (No. 62022011), Beijing Municipal Natural Science Foundation (No. 4222049), the Research Program of State Key Laboratory of Software Development Environment (SKLSDE-2021ZX-04), and the Fundamental Research Funds for the Central Universities.



## References

- [1] Victoria Fernández Abrevaya, Stefanie Wuhler, and Edmond Boyer. Spatiotemporal modeling for efficient registration of dynamic 3d faces. In *3DV*, 2018. 1
- [2] Oswald Aldrian and William AP Smith. Inverse rendering of faces with a 3d morphable model. *IEEE TPAMI*, 35(5):1080–1093, 2012. 1
- [3] Thiemo Alldieck, Hongyi Xu, and Cristian Sminchisescu. imghum: Implicit generative models of 3d human shape and articulated pose. In *ICCV*, 2021. 2
- [4] Brian Amberg, Sami Romdhani, and Thomas Vetter. Optimal step nonrigid icp algorithms for surface registration. In *CVPR*, 2007. 2
- [5] Timur Bagautdinov, Chenglei Wu, Jason Saragih, Pascal Fua, and Yaser Sheikh. Modeling facial geometry using compositional vaes. In *CVPR*, 2018. 1, 2
- [6] Mehdi Bahri, Eimear O’Sullivan, Shunwang Gong, Feng Liu, Xiaoming Liu, Michael M Bronstein, and Stefanos Zafeiriou. Shape my face: registering 3d face scans by surface-to-surface translation. *IJCV*, 129(9):2680–2713, 2021. 1
- [7] Volker Blanz and Thomas Vetter. A morphable model for the synthesis of 3d faces. In *SIGGRAPH*, 1999. 1, 2
- [8] Volker Blanz and Thomas Vetter. Face recognition based on fitting a 3d morphable model. *IEEE TPAMI*, 25(9):1063–1074, 2003. 1
- [9] Timo Bolkart and Stefanie Wuhler. A groupwise multilinear correspondence optimization for 3d faces. In *ICCV*, 2015. 1, 2
- [10] James Booth, Anastasios Roussos, Allan Ponniah, David Dunaway, and Stefanos Zafeiriou. Large scale 3d morphable models. *IJCV*, 126(2):233–254, 2018. 1, 2
- [11] Giorgos Bouritsas, Sergiy Bokhnyak, Stylianos Ploumpis, Michael Bronstein, and Stefanos Zafeiriou. Neural 3d morphable models: Spiral convolutional networks for 3d shape representation learning and generation. In *ICCV*, 2019. 1, 2
- [12] Alan Brunton, Timo Bolkart, and Stefanie Wuhler. Multilinear wavelets: A statistical shape space for human faces. In *ECCV*, 2014. 1, 2
- [13] Chen Cao, Yanlin Weng, Shun Zhou, Yiyong Tong, and Kun Zhou. Facewarehouse: A 3d facial expression database for visual computing. *IEEE TVCG*, 20(3):413–425, 2013. 1, 2
- [14] Xu Chen, Yufeng Zheng, Michael J. Black, Otmar Hilliges, and Andreas Geiger. Snarf: Differentiable forward skinning for animating non-rigid neural implicit shapes. In *ICCV*, 2021. 2
- [15] Zhixiang Chen and Tae-Kyun Kim. Learning feature aggregation for deep 3d morphable models. In *CVPR*, 2021. 1, 2
- [16] Zhiqin Chen and Hao Zhang. Learning implicit fields for generative shape modeling. In *CVPR*, 2019. 1, 2
- [17] Shiyang Cheng, Michael M. Bronstein, Yuxiang Zhou, Irene Kotsia, Maja Pantic, and Stefanos Zafeiriou. Meshgan: Non-linear 3d morphable models of faces. *CoRR*, abs/1903.10384, 2019. 1, 2
- [18] Julian Chibane, Aymen Mir, and Gerard Pons-Moll. Neural unsigned distance fields for implicit function learning. In *NeurIPS*, 2020. 2, 3
- [19] Enric Corona, Albert Pumarola, Guillem Alenya, Gerard Pons-Moll, and Francesc Moreno-Noguer. Smplicit: Topology-aware generative model for clothed people. In *CVPR*, 2021. 2
- [20] Darren Cosker, Eva Krumerhuber, and Adrian Hilton. A face valid 3d dynamic action unit database with applications to 3d dynamic morphable facial modeling. In *ICCV*, 2011. 2
- [21] Yu Deng, Jiaolong Yang, and Xin Tong. Deformed implicit field: Modeling 3d shapes with learned dense correspondence. In *CVPR*, 2021. 2, 4, 5
- [22] Bernhard Egger, William AP Smith, Ayush Tewari, Stefanie Wuhler, Michael Zollhoefer, Thabo Beeler, Florian Bernard, Timo Bolkart, Adam Kortylewski, Sami Romdhani, et al. 3d morphable face models—past, present, and future. *ACM TOG*, 39(5):1–38, 2020. 2
- [23] Kyle Genova, Forrester Cole, Avneesh Sud, Aaron Sarna, and Thomas Funkhouser. Local deep implicit functions for 3d shape. In *CVPR*, 2020. 2
- [24] Kyle Genova, Forrester Cole, Daniel Vlasic, Aaron Sarna, William T Freeman, and Thomas Funkhouser. Learning shape templates with structured implicit functions. In *ICCV*, 2019. 2
- [25] Syed Zulqarnain Gilani, Ajmal Mian, Faisal Shafait, and Ian Reid. Dense 3d face correspondence. *IEEE TPAMI*, 40(7):1584–1598, 2017. 1
- [26] Amos Gropp, Lior Yariv, Niv Haim, Matan Atzmon, and Yaron Lipman. Implicit geometric regularization for learning shapes. In *ICML*. 2020. 1, 2, 5
- [27] Guosheng Hu, Fei Yan, Chi-Ho Chan, Weihong Deng, William Christmas, Josef Kittler, and Neil M Robertson. Face recognition using a unified 3d morphable model. In *ECCV*, 2016. 1
- [28] Moritz Ibing, Isaak Lim, and Leif Kobbelt. 3d shape generation with grid-based implicit functions. In *CVPR*, 2021. 2
- [29] Der-Tsai Lee and Bruce J Schachter. Two algorithms for constructing a delaunay triangulation. *International Journal of Computer & Information Sciences*, 9(3):219–242, 1980. 6
- [30] John P Lewis, Matt Cordner, and Nickson Fong. Pose space deformation: a unified approach to shape interpolation and skeleton-driven deformation. In *SIGGRAPH*, 2000. 2, 4
- [31] Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4d scans. *ACM TOG*, 36(6):194–1, 2017. 1, 2, 6, 7
- [32] Yaron Lipman. Phase transitions, distance functions, and implicit neural representations. In *ICML*, 2021. 1, 2
- [33] Feng Liu and Xiaoming Liu. Learning implicit functions for topology-varying dense 3d shape correspondence. In *NeurIPS*, 2020. 2
- [34] Feng Liu, Luan Tran, and Xiaoming Liu. 3d face modeling from diverse raw scan data. In *ICCV*, 2019. 1
- [35] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *CVPR*, 2019. 1, 2, 3

- [36] 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. In *ECCV*, 2020. 4
- [37] Tomas Möller and Ben Trumbore. Fast, minimum storage ray-triangle intersection. *Journal of graphics tools*, 2(1):21–28, 1997. 6
- [38] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *CVPR*, 2019. 1, 2, 3, 4
- [39] Ankur Patel and William AP Smith. 3d morphable face models revisited. In *CVPR*, 2009. 1
- [40] Pascal Paysan, Reinhard Knothe, Brian Amberg, Sami Romdhani, and Thomas Vetter. A 3d face model for pose and illumination invariant face recognition. In *AVSS*, 2009. 1
- [41] Sida Peng, Junting Dong, Qianqian Wang, Shangzhan Zhang, Qing Shuai, Xiaowei Zhou, and Hujun Bao. Animatable neural radiance fields for modeling dynamic human bodies. In *ICCV*, 2021. 4
- [42] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *ECCV*, 2020. 2
- [43] Sida Peng, Yuanqing Zhang, Yinghao Xu, Qianqian Wang, Qing Shuai, Hujun Bao, and Xiaowei Zhou. Neural body: Implicit neural representations with structured latent codes for novel view synthesis of dynamic humans. In *CVPR*, 2021. 2
- [44] Mallikarjun B R, Ayush Tewari, Hans-Peter Seidel, Mohamed Elgharib, and Christian Theobalt. Learning complete 3d morphable face models from images and videos. In *CVPR*, 2021. 2
- [45] Eduard Ramon, Gil Triginer, Janna Escur, Albert Pumarola, Jaime Garcia, Xavier Giro-i Nieto, and Francesc Moreno-Noguer. H3d-net: Few-shot high-fidelity 3d head reconstruction. In *ICCV*, 2021. 2
- [46] Anurag Ranjan, Timo Bolkart, Soubhik Sanyal, and Michael J Black. Generating 3d faces using convolutional mesh autoencoders. In *ECCV*, 2018. 1, 2
- [47] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *ICCV*, 2019. 2
- [48] Shunsuke Saito, Tomas Simon, Jason Saragih, and Hanbyul Joo. Pifuhd: Multi-level pixel-aligned implicit function for high-resolution 3d human digitization. In *CVPR*, 2020. 2
- [49] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *NeurIPS*, 2020. 2, 4
- [50] Femke CR Staal, Allan JT Ponniah, Freida Angullia, Clifford Ruff, Maarten J Koudstaal, and David Dunaway. Describing crouzon and pfeiffer syndrome based on principal component analysis. *Journal of Cranio-Maxillofacial Surgery*, 43(4):528–536, 2015. 1
- [51] Towaki Takikawa, Joey Litalien, Kangxue Yin, Karsten Kreis, Charles Loop, Derek Nowrouzezahrai, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Neural geometric level of detail: Real-time rendering with implicit 3d shapes. In *CVPR*, 2021. 2
- [52] Jia-Heng Tang, Weikai Chen, Jie Yang, Bo Wang, Songrun Liu, Bo Yang, and Lin Gao. Octfield: Hierarchical implicit functions for 3d modeling. In *NeurIPS*, 2021. 2
- [53] Luan Tran, Feng Liu, and Xiaoming Liu. Towards high-fidelity nonlinear 3d face morphable model. In *CVPR*, 2019. 1, 2
- [54] Luan Tran and Xiaoming Liu. Nonlinear 3d face morphable model. In *CVPR*, 2018. 1, 2
- [55] Daniel Vlasic, Matthew Brand, Hanspeter Pfister, and Jovan Popovic. Face transfer with multilinear models. In *ACM SIGGRAPH Courses*. 2006. 1, 2
- [56] 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. In *NeurIPS*, 2021. 2
- [57] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. In *NeurIPS*, 2019. 2
- [58] Haotian Yang, Hao Zhu, Yanru Wang, Mingkai Huang, Qiu Shen, Ruigang Yang, and Xun Cao. Facescape: a large-scale high quality 3d face dataset and detailed riggable 3d face prediction. In *CVPR*, 2020. 1, 2, 6, 7
- [59] 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 *CVPR*, 2021. 2, 4, 6, 7
- [60] Li Yi, Vladimir G Kim, Duygu Ceylan, I-Chao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas Guibas. A scalable active framework for region annotation in 3d shape collections. *ACM TOG*, 35(6):1–12, 2016. 2
- [61] Jingyang Zhang, Yao Yao, and Long Quan. Learning signed distance field for multi-view surface reconstruction. In *ICCV*, 2021. 2
- [62] Zerong Zheng, Tao Yu, Qionghai Dai, and Yebin Liu. Deep implicit templates for 3d shape representation. In *CVPR*, 2021. 2