Discovering Geometric Biases in 3D Face Reconstruction: A Curvature-Aware Spectral Framework for Fairness Evaluation

Veronika Shilova[A,B], Emmanuel Malherbe[A], Giovanni Palma[C], Panagiotis-Alexandros Bokaris[C], Laurent Risser[B] and Jean-Michel Loubes[B,D,E]

[A]Artefact Research Center, Paris, France , [B]Institut de Mathématique de Toulouse, CNRS, Toulouse, France, [C]L'Oréal Recherche et Innovation, Paris, France, [D]Université de Toulouse, ANITI, Toulouse, France, [E]INRIA Regalia, Bordeaux, France

3D Morphable Models (3DMMs) remain the standard parametric shape priors for many state-of-the-art 3D face reconstruction algorithms. However, as these models are derived from a finite number of 3D face samples, they inherit the morphological biases of their training data, potentially limiting their generalizability across diverse global populations. In this paper, we propose a novel framework to analyze 3DMM reconstructions through the lens of surface curvature, with the objective to discover, quantify and visualize biases. While standard evaluation metrics often rely on Euclidean distances, our reconstruction error captures subtle surface nuances such as local topology or undulations. To do so, we leverage the Laplace-Beltrami Operator (LBO) to generate high-resolution curvature error maps, providing a localized and geometrically meaningful visualization of discrepancies between ground truth faces and reconstructed meshes. We derive from it an error metric that we validated through a user study, observing a significantly higher correlation to human perception compared to traditional methods. Furthermore, we conduct extensive experiments across several 3DMM bases and fitting algorithms, uncovering systematic age-related biases and providing preliminary evidence of biases associated with gender and ethnicity. Our findings highlight the necessity of adopting curvature-aware evaluation protocols to ensure demographic fairness and geometric precision in future 3D face reconstruction research. The code and annotation data are available at https://github.com/artefactory/3dface-fairness.

1. Introduction

The field of 3D face reconstruction from images has become important in computer vision, as it makes it possible to compute accurate representations of the human facial morphology, with many applications such as in the field of cinema [17], virtual reality [24, 36], cosmetics [39], biometric security [21] or medicine [27]. An underlying component of any face reconstruction is the parametrization of the 3D shapes, for which 3D Morphable Models (3DMM) [5] have established themselves as fundamental parametric shape priors for 3D face reconstruction [13].

The 3DMM formalism is based on two fundamental assumptions. First, it establishes a dense, point-to-point correspondence across registered face scans in its database. This shared topology allows for the generation of anatomically plausible identities through linear combinations of basis shapes. Second, the model explicitly disentangles facial shape and texture, and separates them from extrinsic scene parameters, such as illumination and camera parameters. The classic 3DMM formulation decomposes facial geometry and texture using Principal Component Analysis (PCA) bases derived on a finite set of 3D face scans.

In this sense, they are prone to machine-learning bias issues [4, 34, 10], meaning that they can systematically offer poor performances on specific categories of faces, due to, *e.g.*, age, gender, or geographical origin. These limitations can have a negative impact on the reputation of companies using these tools and can even be considered as discriminatory in certain cases. Strategies for detecting

Published at European Conference on Computer Vision (ECCV) 2026
Artefact Research Center, Paris, France

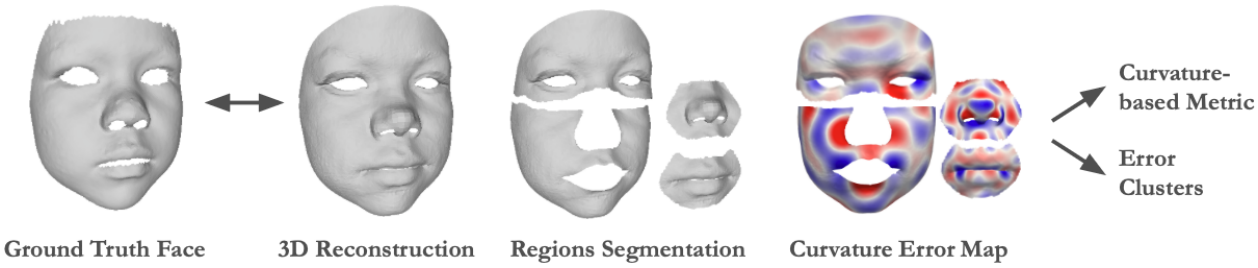


Figure 1 | Overview of our curvature-aware evaluation framework for bias discovery. Given a ground truth scan and its corresponding 3D reconstruction, we perform an anatomical segmentation of the reconstructed mesh. We compute Curvature-Based Error Maps using the Laplace-Beltrami Operator to isolate localized geometric discrepancies that Euclidean metrics often overlook. By performing spectral clustering on these error maps, we identify distinct failure modes and correlate them with demographic factors.

and quantifying these biases are therefore necessary to ensure full confidence in these tools. We want to emphasize that regulatory texts such as the *AI act* [1], explicitly mention the need for accuracy metrics, as *e.g.*, in its Article 15: *The levels of accuracy and the relevant accuracy metrics of high-risk AI systems shall be declared in the accompanying instructions of use.* To the best of our knowledge, this work represents the first comprehensive evaluation of demographic bias in 3D face reconstruction, addressing a critical gap in a field where facial analysis is increasingly categorized as high-risk.

In our paper, we explore new solutions to quantify and detect algorithmic biases in 3D face reconstruction algorithms (fig. 1). Our main contributions are:

- *Curvature-Aware Evaluation Framework:* We introduce a new framework for 3D face reconstruction error analysis that focuses on local differential geometry using the Laplace-Beltrami Operator (LBO), to quantify and visualize surface discrepancies that are often invisible to standard Euclidean metrics.
- *Perceptual Validation:* We demonstrate through a user study that our curvature-based reconstruction error aligns more closely with human qualitative judgment than traditional point-wise distances.
- *Systematic Bias Discovery:* To our knowledge, we conduct the first study to expose morphological biases in 3D face reconstruction algorithms. Our analysis reveals substantial fidelity gaps across demographic factors, specifically identifying systematic age-related biases and providing preliminary evidence of gender- and ethnicity-related biases.

While demonstrated on 3DMMs, our curvature-aware evaluation framework is agnostic to the reconstruction method, and our bias discovery strategy generalizes to any template-based 3D surface reconstruction.

2. Related Work

2.1. 3D Morphable Model (3DMM) and 3D Face Reconstruction

The objective of 3D face reconstruction is to recover the dense geometry of a face from a single RGB(-D) image or an image collection of the same person. To do so, most methods use a 3DMM as its shape prior. Consequently, the reconstruction task is formulated as a regression or optimization problem to estimate the latent shape parameters. Current approaches either use deep networks to

predict them [12, 40, 37, 14, 32, 30] or optimize a combination of energy terms [6, 20, 35, 38], such as landmark-based alignment and photometric consistency, often subject to various statistical regularizations.

Basel Face Model (BFM) [29] was the first publicly available 3DMM trained on 200 individuals. More recently, FLAME [23] has become the standard for its database diversity (3800 face scans) and its ability to model the head, neck, and jaw articulation more realistically than previous vertex-based models. However, since these models were trained on limited amount of data, they often fail to capture the fine-grained geometric details that differentiate individuals across diverse demographics.

2.2. 3D Face Reconstruction Metrics

Standard evaluation in the field typically relies on global geometric fidelity. The most prevalent metric is point-to-point or point-to-surface Root Mean Square Error (RMSE) or Normalized Mean Square Error (NMSE) [15], often calculated using the Iterative Closest Point (ICP) [3, 2] algorithm for alignment. While these metrics assess the global "shell" of the face, they are insensitive to regional anatomical properties.

To try to bridge the gap, [8] proposed a new benchmark called REALY consisting of 100 globally aligned high-quality face scans with 4 anatomical region masks (including nose, mouth, cheek and forehead) and a bi-directional evaluation pipeline with Non-Rigid Iterative Closest Point [3, 2] for better alignment. Moreover, [31] introduced Modularized 3D Face reconstruction Benchmark (M3DFB) where they separate the fundamental components of error computation such as rigid alignment or point correspondence, allowing one to quantify the effect of each. Nevertheless, they still focus only on Euclidean distances between meshes in their evaluation pipelines.

While bias in facial analysis is well-documented in 2D recognition [18, 9], its quantification in 3D reconstruction has not been widely studied in the literature. Recent attempts to quantify bias [10] were limited by reliance on global Euclidean distances, resulting in marginal and ultimately inconclusive performance gaps across demographics. This highlights the need for curvature-aware metrics capable of uncovering localized reconstruction failures.

3. Methodology

3.1. Background and Notations

3.1.1. Face Mesh.

Let $\mathcal{M} = \{\mathcal{V}, \mathcal{F}\}$ be a triangle mesh with vertex positions $\mathcal{V}$ in the 3D ambient space, and a list of triangular faces $\mathcal{F}$. In this paper, $\mathcal{M}$ generally represents the surface of someone's face, however we can equivalently consider a subpart of the mesh $R_{\mathcal{M}} \subsetneq \mathcal{M}$ constituting a meaningful morphological region of the face (*e.g.*, nose), also called region-of-interest. Let $\mathcal{M}'$ be another mesh that aims at representing the same face, possibly with a different number of vertices $|\mathcal{V}'| \neq |\mathcal{V}|$. In the context of our paper, dedicated to 3D surface reconstruction evaluation, the ground-truth mesh is $\mathcal{M}$ and its reconstruction is $\mathcal{M}'$. In all the following, we assume that meshes $\mathcal{M}$ and $\mathcal{M}'$ are first rigidly aligned using Iterative Closest Point algorithm [3].

3.1.2. Mesh Mapping.

We denote by $T_{\mathcal{M}' \to \mathcal{M}}$ a mapping between $\mathcal{M}'$ and $\mathcal{M}$, where the subscript represents the direction of the map and $T_{\mathcal{M}' \to \mathcal{M}}(v)$ gives the coordinates of the mapped position of a vertex $v \in \mathcal{M}'$ on shape $\mathcal{M}$.

We consider two different types of maps: the point-to-surface map $T^{pts}_{\mathcal{M}'\to\mathcal{M}}$ (also called vertex-to-plane or vertex-to-point), which is the most commonly used, and the vertex-to-vertex map $T^{vtx}_{\mathcal{M}'\to\mathcal{M}}$, which is less accurate but remains considered in some applications. Specifically, $T^{vtx}_{\mathcal{M}'\to\mathcal{M}}$ maps each vertex in shape $\mathcal{M}'$ to the closest vertex in $\mathcal{V}$, and $T^{pts}_{\mathcal{M}'\to\mathcal{M}}$ maps each vertex in shape $\mathcal{M}'$ to the closest point in a face of $\mathcal{F}$. We will only use the superscript to disambiguate the two maps when necessary.

3.1.3. Normalized Mean Square Error (NMSE).

The NMSE metric, denoted as $\mathcal{E}^{pts}_{NMSE}$, computes the distance between two meshes $\mathcal{M}$ and $\mathcal{M}'$ using the vertex-to-point map $T^{pts}_{\mathcal{M}\to\mathcal{M}'}$:

$$\mathcal{E}^{pts}_{NMSE}(\mathcal{M},\mathcal{M}') = \frac{1}{|\mathcal{V}'|}\sum_{v\in\mathcal{M}'} ||v - T^{pts}_{\mathcal{M}'\to\mathcal{M}}(v)||^2, \tag{1}$$

where $|\mathcal{V}'|$ is the number of vertices in shape $\mathcal{M}'$. We similarly define $\mathcal{E}^{vtx}_{NMSE}$, substituting point-to-surface mapping for the vertex-to-vertex correspondence. We recall that such similarity metrics are routinely used to assess the accuracy of 3D surface reconstructions [3, 2, 15].

3.1.4. 3DMM.

A face $\mathcal{M} = (\mathcal{V},\mathcal{F})$ is expressed in the 3DMM base if it verifies $\mathcal{V} = \vec{\mu} + \sum_{j=1}^{n}\beta_j\vec{\psi}_j$, with $\vec{\mu}$ the mean shape, $\vec{\psi}_j$ the principal components and β_j the shape parameters of the 3DMM. In this paper, we assume that $\mathcal{M}'$ is expressed in a given 3DMM base. In such base, all vertices systematically represent the same points semantically. Note that all our methods would equivalently apply if we use a standard base that do not originates from a 3DMM, or if it is the mesh $\mathcal{M}$ that is expressed in the 3DMM.

3.2. Quantifying and Visualizing the Geometric Reconstruction Accuracy Using Curvature Reconstruction Errors (CRE)

To quantify the preservation of structural identity and identify systematic geometric biases, we propose to use the **Mean Curvature** (H), a second-order differential property characterizing the local geometry of the manifold. Unlike point-wise Euclidean metrics, H is sensitive to the local surface "character"—such as ridges, valleys, and high-frequency morphological details—making it a robust indicator of reconstruction fidelity.

3.2.1. Mathematical Definition of the Mean Curvature.

For a smooth surface $\mathcal{S}\subset\mathbb{R}^3$, the mean curvature H at a surface point (x,y,z) is defined as the arithmetic mean of the two principal curvatures, κ_1 and κ_2:

$$H = \frac{\kappa_1+\kappa_2}{2}\,. \tag{2}$$

In the context of differential geometry, H can also be derived using the Laplace-Beltrami operator, which is the manifold generalization of the Laplacian to curved surfaces [7]. By writing this operator Δ and the parametrized coordinate function $\vec{p}:(\xi,\eta)\in\mathcal{S}\to\big(x(\xi,\eta),y(\xi,\eta),z(\xi,\eta)\big)\in\mathbb{R}^3$ defined on surface $\mathcal{S}$, we have

$$H = -\frac{1}{2}\langle\vec{n},\Delta\vec{p}\rangle\,, \tag{3}$$

where $\vec{n}$ is the unit normal vector at the considered point. In the following, we will describe a numerical estimation of H following this equation.

3.2.2. Laplace-Beltrami Estimation using Local Quadratic Regressions.

Since triangular meshes are piecewise linear, they lack well-defined analytic derivatives at the vertices. Existing methodologies to estimate the curvature [19, 26, 16] starts by fitting a quadratic bivariate surface to each vertex $\vec{v}_i \in \mathcal{V}$ and its neighbor vertices. We propose to differ by locally estimating the Laplace-Beltrami operator, in order to capture the underlying differential properties and apply the Equation 3. Indeed, such operator will be independent from the distances to tangent plane, and thus less prone to noise in our data, such as the one related to the tangent plane and its normal. In addition, we propose to use a numerical estimation of this estimator, performed on local neighborhoods that are adapted to user-defined parameters.

For each vertex $\vec{v}_i$, we consider its R-ring neighborhood $\mathcal{N}_R(\vec{v}_i) = \{\vec{v} \in \mathcal{V} \mid d_G(\vec{v}_i, \vec{v}) \leq R\}$, where $d_G(\vec{v}_i, \vec{v})$ denotes the graph distance (the shortest path in terms of number of edges) between $\vec{v}_i$ and $\vec{v}$. We compute the corresponding surface normal $\vec{n}_i$ and the tangent plane, written in coordinates (ξ, η). We consider a scalar function defined on this plane $f : (\xi, \eta) \in \mathbb{R}^2 \to \mathbb{R}$. For instance, the coordinate function $\vec{p}$ (Equation 3) is locally a concatenation of 3 such functions. We write $\hat{f}_{\vec{a}}$ the second-order polynomial approximation of f:

$$\hat{f}_{\vec{a}}(\xi, \eta) = a_1\xi^2 + a_2\xi\eta + a_3\eta^2 + a_4\xi + a_5\eta + a_6 \,. \tag{4}$$

where the vector of coefficients $\vec{a} = [a_1, a_2, a_3, a_4, a_5, a_6]^\top \in \mathbb{R}^6$ are obtained by solving the least squares problem on the local neighborhood of v_i:

$$\min_{\vec{a}} \sum_{v_j \in \mathcal{N}_R(v_i)} \left(\hat{f}_{\vec{a}} \left(\xi_{v_j}, \eta_{v_j} \right) - f_{v_j} \right)^2 , \tag{5}$$

where we write (ξ_{v_j}, η_{v_j}) the projected coordinates of the neighbor v_j in the tangent plane, and $f_{v_j} = f(\xi_{v_j}, \eta_{v_j})$. This minimization can be rewritten in the form $\min_{\vec{a}} \|\mathbf{X}\vec{a} - \vec{f}\|^2$. It has an analytical solution $\vec{a} = \mathbf{C}(v_i)\vec{f}$, where $\mathbf{C}(v_i) = (\mathbf{X}^\top\mathbf{X})^{-1}\mathbf{X}^\top$ is calculated from the neighbors v_j of vertex v_i. This matrix $\mathbf{C}(v_i)$ acts as the pseudoinverse that maps a function's values f_{v_j} directly to the coefficients a_m of its polynomial approximation $f_{\vec{a}}$, and can be used to approximate the Laplace operator at the origin $(0, 0)$:

$$\Delta\vec{f} = \frac{\partial^2 f}{\partial\xi^2} + \frac{\partial^2 f}{\partial\eta^2} \approx 2a_1 + 2a_3 = \sum_{v_j \in \mathcal{N}_R(v_i)} 2(C_{1j} + C_{3j}) f_{v_j}. \tag{6}$$

Thus, we can approximate Laplace operator Δ with the sparse matrix $\tilde{\Delta} \in \mathbb{R}^{|\mathcal{V}|\times|\mathcal{V}|}$ with entries $\tilde{\Delta}_{ij} = 2(C(v_i)_{1j} + C(v_i)_{3j})$ when i, j verifies $v_j \in \mathcal{N}_R(v_i)$ ($\tilde{\Delta}_{ij} = 0$ otherwise). Following the fundamental identity in eq. (3) from differential geometry, we calculate the mean curvature at each vertex v_i as follows

$$H^{\mathcal{M}}(v_i) = -\frac{1}{2}\langle \vec{n}_i, (\tilde{\Delta}v)_i \rangle \,, \tag{7}$$

where $\vec{n}_i \in \mathbb{R}^3$ is the unit normal vector at the vertex v_i, $v \in \mathbb{R}^{|\mathcal{V}|\times 3}$ is the matrix of the Cartesian coordinates of the vertices $\mathcal{V}$, and $(\tilde{\Delta}v)_i \in \mathbb{R}^3$ is the local Laplace-Beltrami estimation of each $\vec{p}$ dimension.

Please note that the sparse matrix $\tilde{\Delta}$ is nonsymetric and can be efficiently computed since the matrices $\mathbf{C}(v_i)$ are of small size. The user-defined parameter R provides a trade-off between local precision and numerical stability when estimating the curvature. While a larger R-ring neighborhood enhances the robustness of the operator against discretization noise, it introduces a smoothing effect that may attenuate high-frequency geometric details.

3.2.3. Curvature-Based Error Metric.

Following this formulation, we introduce the **Curvature Reconstruction Error (CRE)**. Given a ground-truth mesh $\mathcal{M}$ and its reconstruction $\mathcal{M}'$, we define the curvature-based error map $\epsilon_H(v_i) \in \mathbb{R}^{|\mathcal{V}'|}$ as:

$$\epsilon_H(v_i) = H^{\mathcal{M}'}(v_i) - H^{\mathcal{M}}\left(T^{pts}_{\mathcal{M}'\to\mathcal{M}}(v_i)\right), \tag{8}$$

where $H^{\mathcal{M}}(.)$ is computed on the vertices of $\mathcal{M}$ using the methodology described in the previous section, and linearly interpolated on the triangular faces $\mathcal{F}$ for the positions given by the mapping $T^{pts}_{\mathcal{M}'\to\mathcal{M}}$.

We derive from ϵ_H a global error metric as a scalar quantity

$$\mathcal{E}_{CRE}(\mathcal{M}, \mathcal{M}') = \frac{1}{Tr(M)} \|M\epsilon_H\|_{L_1}, \tag{9}$$

where M is the diagonal mass matrix whose entries $M_{i,i}$ correspond to the barycentric area for each vertex $v_i \in \mathcal{V}'$, $\|\cdot\|_{L_1}$ denotes the L_1 norm and $Tr(\cdot)$ the trace.

3.3. Unknown Geometric Bias Detection Using the Laplace-Beltrami Operator

In this section, we leverage the spectral decomposition induced from Laplace-Beltrami operator and use it to identify various failure modes of 3D face reconstruction methods.

3.3.1. Spectral Decomposition.

Manifold harmonics (also called shape harmonics) are defined as the eigenfunctions ϕ_k of the Laplace-Beltrami operator Δ introduced in the previous section:

$$\Delta\phi_k = \lambda_k\phi_k, \;\; k = 1, \ldots, |\mathcal{V}|. \tag{10}$$

These eigenfunctions form a spectral basis tailored to the specific geometry of the manifold, where lower eigenvalues λ_k correspond to smooth, global deformations while higher ones correspond to rapidly oscillating, local details. Please note that this spectral decomposition would be difficult to be computed using the estimated $\tilde{\Delta}$ matrix introduced in the previous section, since it is not a symmetric matrix. Moreover, the Laplace-Beltrami operator was previously estimated mesh-wise, while the operator below will be the same across meshes of a 3DMM.

For a discrete triangular mesh $\mathcal{M} = (\mathcal{V}, \mathcal{F})$, the eigenvectors and eigenvalues are obtained by solving the generalized eigenvalue problem

$$L\vec{\phi}_k = \lambda_k M\vec{\phi}_k, \tag{11}$$

where L is the cotangent weight (stiffness) matrix representing the weak Laplacian, and M is the diagonal mass matrix defined above. $M^{-1}L$ corresponds to the analytical expression of the Laplace-Beltrami operator for 1-ring neighborhoods, and is symmetrical. The resulting eigenvectors $\vec{\phi}_k \in \mathbb{R}^{|\mathcal{V}|}$ are orthonormal with respect to the mass matrix, meaning $\vec{\phi}_k^{\top} M\vec{\phi}_{k'} = \delta_{kk'}$. These eigenvectors can be easily precomputed for a mesh $\mathcal{M}$ using standard numerical packages [22, 33].

3.3.2. Spectral Error Projection for a 3DMM.

We propose to leverage the shape harmonics in order to re-express the reconstruction error into geometric frequencies, instead of scalars defined at the vertices level. To do so, we solve the eigenvalues problem Eq. (11) for the mesh $\mathcal{M}$ – or region mesh $R_{\mathcal{M}}$ – from the 3DMM template $\vec{\mu}$, and its corresponding M and L matrices. Please note that, contrary to curvature-based error computation that is mesh-wise, we now use a general Laplace-Beltrami operator, whose the corresponding spectral projection applies to any face expressed in the 3DMM.

Given a user-defined parameter K, let $\Phi := [\vec{\phi}_k] \in \mathbb{R}^{|\mathcal{V}|\times K}$ be the orthonormal basis with respect to the mass matrix M, with eigenvalues in increasing order $0 \leq \lambda_1 < \cdots < \lambda_K$. Given a curvature-based reconstruction error $\epsilon_H \in \mathbb{R}^{|\mathcal{V}|}$, we can compute its spectral coefficients $\vec{c}_{\epsilon_H} \in \mathbb{R}^K$ via the spectral projection

$$\vec{c}_{\epsilon_H} = \Phi^\top M \epsilon_H. \tag{12}$$

3.3.3. Unknown Bias Discovery using Clustering in the Spectral Space.

We now consider a set of N pairs of ground truth faces with their reconstruction $\{(\mathcal{M}^{(n)}, \mathcal{M}'^{(n)})\}_{n=1}^N$, assuming that they sufficiently represent the possible biases of the underlying reconstruction algorithm. We propose a clustering method that aims at discovering systematic failure modes. To do so, this method identifies recurring patterns of geometric bias by analyzing the curvature-based reconstruction error maps ϵ_H of these N pairs.

With this objective, a practical challenge is that the number of ground truth faces N is generally much lower than the dimension of ϵ_H, which is $|\mathcal{V}'|$. As described in Algorithm 1, we project the error maps into the spectral space (Equation 12), and perform a K-means clustering [25] on the spectral errors $\{\mathbf{c}_{\epsilon_H}^{(n)}\}_{n=1}^N$. Indeed, this spectral space is of much lower dimension $K \ll |\mathcal{V}'|$ and still well represents the geometric patterns of the errors.

Given a resulting cluster, its centroid can be easily visualized by reconstructing the signal in the spatial domain through the inverse transform:

$$\hat{\epsilon}_H = \Phi \vec{c}_{\epsilon_H}. \tag{13}$$

which can be used for visualization purposes to illustrate the failure mode of the cluster.

4. Experiments

4.1. Protocol

4.1.1. Ground Truth Faces.

For a high-fidelity evaluation of morphological preservation, we choose REALY benchmark [8], a Region-aware benchmark based on the LYHM [11] dataset. Unlike standard 3D datasets, REALY provides 100 globally aligned 3D face scans with accurate facial keypoints and high-quality region masks (nose, mouth, cheek and forehead). Besides, these faces are assigned with demographic labels: age, gender and ethnicity (see fig. 2 for details).

4.1.2. 3DMM Base Representativeness.

In the following sections, we evaluate the representation capacity of the Basel Face Model (BFM) [29] and FLAME [23], the two most prevalent 3DMMs underpinning current state-of-the-art methods.

Algorithm 1 Clustering of Error Maps

Input: Template mesh $\vec{\mu}$, set of error maps $\{\epsilon_H^{(n)}\}_{n=1}^N$, maximum frequency $K \in \mathbb{N}$

Calculate cotangent matrix L from template mesh $\vec{\mu}$
Calculate mass matrix M from template mesh $\vec{\mu}$
Find $\Phi = [\vec{\phi}_k]$, λ_k, $k = 1, \ldots, K$, from eq. (11)
for $n = 1$ to N **do**
    Project error map to spectral space of template: $c_{\epsilon_H}^{\vec{(n)}} \leftarrow \Phi^\top M \epsilon_H^{(n)}$
end for
Perform K-means on the set $\{c_{\epsilon_H}^{\vec{(n)}}\}_{n=1}^N$
for each centroid $\vec{\bar{c}}_{\epsilon_H}$ **do**
    Project back to vertex space: $\bar{\epsilon}_H \leftarrow \Phi \vec{\bar{c}}_{\epsilon_H}$
end for
return $\bar{\epsilon}_H$, centroids in the vertex space

To evaluate this representativeness, given a 3D face scan $\mathcal{M}$, we want to find the best face $\mathcal{M}' = (\mathcal{V}', \mathcal{F}')$ that can be expressed in the 3DMM base, written $\mathcal{V}' = \vec{\mu} + \sum_{j=1}^n \beta_j \vec{\psi}_j$ with $\vec{\mu}$ the mean face and $\vec{\psi}_j$ the PCs of the 3DMM basis. Note that the reference 3DMM mean shape $\vec{\mu}$ has been preliminarily rigidly aligned and scaled to $\mathcal{M}$ using the ICP. We obtain the β_j values by solving using gradient descent the following optimization problem

$$\min_{\{\beta_j\}_{j=1}^n} \sum_{\vec{v} \in \mathcal{V}} ||\vec{v} - T^{vtx}_{\mathcal{M} \to \mathcal{M}'}(\vec{v})||^2, \tag{14}$$

By minimizing this 3D-to-3D distance, we obtain the optimal approximation of the target geometry within the span of the 3DMM basis, allowing us to conclude directly on the model's intrinsic expressive capacity, agnostic of any reconstruction method (e.g. from a 2D picture).

4.1.3. Reconstruction Algorithms from 2D Image.

We extended our evaluation to state-of-the-art 3D face reconstruction algorithms from 2D face picture that use 3DMMs as their shape priors: Deep3D[12], 3DDFA-v3[37], MGCNet[32], DECA[14] and MICA[40]. To do so, we used as input the frontal face image of each panelist as provided by REALY dataset. Please note that in this case, the ground truth mesh $\mathcal{M}$ is only used to calculate the reconstruction error.

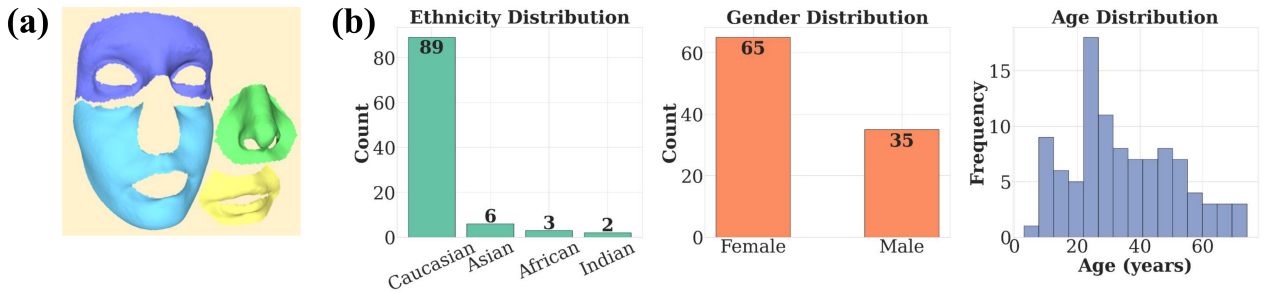


Figure 2 | Overview of REALY benchmark: (a) REALY region masks, (b) REALY Demographic statistics

4.1.4. Anatomical Zones Segmentation.

In order to perform our evaluation region by region, we can leverage the anatomical zones defined in REALY, which are high-quality masks for the nose, mouth, cheeks, and forehead. However, we need to extend these zones into the 3DMM base, that is to define a list of indices for vertices v_i corresponding to each region. To do so, we propose a *Topological Voting Procedure* to assign regions to the vertices of the 3DMM mean shape $\vec{\mu}$.

To do so, given a 3DMM, we consider its template $\vec{\mu}$ and aim at assigning a region label to each of its vertex $\vec{v}_i$, that can be later used for any face expressed in this base. Given a ground truth scan $\mathcal{M} = (\mathcal{V}, \mathcal{F})$ in the REALY dataset, we first align it with the mesh of $\vec{\mu}$ via the ICP algorithm [3, 2]. Then, for every vertex $\vec{v}_i$ in the mesh of $\vec{\mu}$, we identify its regional assignment $l(\vec{v}_i, \mathcal{M})$ by finding its closest vertex in $\mathcal{M}$:

$$l(\vec{v}_i, \mathcal{M}) = \text{Label}\left(\arg\min_{\vec{v}\in\mathcal{V}} \|\vec{v}_i - \vec{v}\|^2\right), \tag{15}$$

where Label$(\cdot)$ returns the anatomical zone (*e.g.*, {Nose, Mouth, ...}) as labeled vertex by vertex in REALY dataset. We compute similarly this region mask $l(\vec{v}_i, \mathcal{M})$ for every face $\mathcal{M}$ of REALY, and finally assign to v_i the region label that has been the most frequently assigned. It provides a more robust region label to v_i vertex and solves the potential inconsistency of Equation 15 from one ground-truth face to another. This consensus-based approach results in a standardized partition of the 3DMM topology that is anatomically consistent with the REALY definition of morphological zones.

4.2. Evaluation of the Curvature-Based Metric

4.2.1. Perceptual Validation.

In order to assess whether our proposed **Curvature Reconstruction Error (CRE)** aligns more closely with human judgment than traditional Euclidean metrics, we conducted a user study. 19 annotators from diverse demographic backgrounds evaluated the reconstruction fidelity in a pairwise comparison.

In detail, for each ground truth scans $\mathcal{M}$ of nose regions from REALY, we reconstructed $\mathcal{M}'$ in Basel Face Model (BFM) using Equation 14. The users were presented with the following pair at a time, $((\mathcal{M}^A, \mathcal{M}'^A), (\mathcal{M}^B, \mathcal{M}'^B))$, with A and B drawn randomly among a set of 2,414 pairs. They voted which reconstruction between A and B exhibited higher morphological fidelity relative to its respective ground truth. This pairwise comparative setup allows us to determine whether the reconstruction error comparison provided by a metric $\mathcal{E}$ (i.e., whether $\mathcal{E}(\mathcal{M}^A, \mathcal{M}'^A) < \mathcal{E}(\mathcal{M}^B, \mathcal{M}'^B)$) aligns with human perceptual preference. For each pair A, B, we obtained $n_{(A,B)}$ annotations ($n_{(A,B)} \geq 5$), on which we computed the ratio of agreement. We retained as labeled pairs the 1,111 pairs for which this ratio exceeds 75%.

Given a metric $\mathcal{E}$, we computed its accuracy as the probability that its choice matches the human majority consensus

$$\text{Pairwise-Accuracy}(\mathcal{E}) = \frac{1}{N_{\text{pairs}}} \sum_{j=1}^{N_{\text{pairs}}} \mathbb{1}_{\left(\mathcal{E}(\mathcal{M}^A, \mathcal{M}'^A) - \mathcal{E}(\mathcal{M}^B, \mathcal{M}'^B)\right) P_j \geq 0}$$

where $P_j \in \{-1, 1\}$ is the perceptual ground truth for j-th pair.

4.2.2. Comparison to Existing Methods and Ablation Study.

table 1 summarizes results of our user experiment. We compared CRE with NMSE based on vertex-to-vertex and point-to-surface alignments, and on bi-directional alignment protocol proposed by REALY

[8] authors. Our curvature-based metric achieves a perceptual accuracy of **73.63%** (for $R = 3$ and $K = 128$), significantly surpassing the performance of NMSE-based benchmarks, which hover near the level of random chance (≈ 0.5 for binary choices). This significant margin confirms that local surface character, as captured by CRE, is a much more salient indicator of reconstruction quality than global point-wise distances, making it an ideal candidate for detecting subtle demographic biases in 3DMMs.

To further motivate CRE and justify our design choices, we conduct an extensive ablation study varying R and K, and comparing L_1, L_2, and Huber norms in eq. (9). As shown in the lower portion of table 1, CRE is robust to its key hyperparameters; remaining configurations with similar trends are reported in the Supplementary Material. fig. 3 qualitatively contrasts CRE and NMSE on three representative pairs. CRE aligns with human judgments in the first two and diverge on the third one.

4.3. Biases and Fairness Analysis

In this section, we use the proposed Curvature Reconstruction Error and regional spectral analysis to quantify demographic biases in 3DMM bases. By analyzing the relationship between reconstruction fidelity and sensitive attributes, we uncover systematic disparities in how these models represent diverse populations.

4.3.1. Age Bias with respect to CRE.

We first investigate the correlation between subject age and the fidelity of the 3DMM representation. We hypothesize that as facial morphology becomes more complex with age (*e.g.*, due to the appearance of folds, creases, and sagging), the linear basis of standard 3DMMs becomes increasingly insufficient. To quantify this, we compute Pearson r correlation coefficient [28] between subject age and the regional $\mathcal{E}_{CRE}$.

table 2 summarizes these findings for the BFM and FLAME models, and the 3D reconstruction models from images. For 3DMMs, we observe a strong positive correlation between age and reconstruction error across nearly all anatomical regions. Notably, in the cheek region, an area highly susceptible to age, related morphological changes – the Pearson correlation reaches **0.77** for BFM and **0.82** for FLAME. We also see for 3D reconstruction models that the age-related bias is consistently propagated from the 3D prior to the final output. Notably, however, this bias do not seem exacerbated by these predictive algorithms.

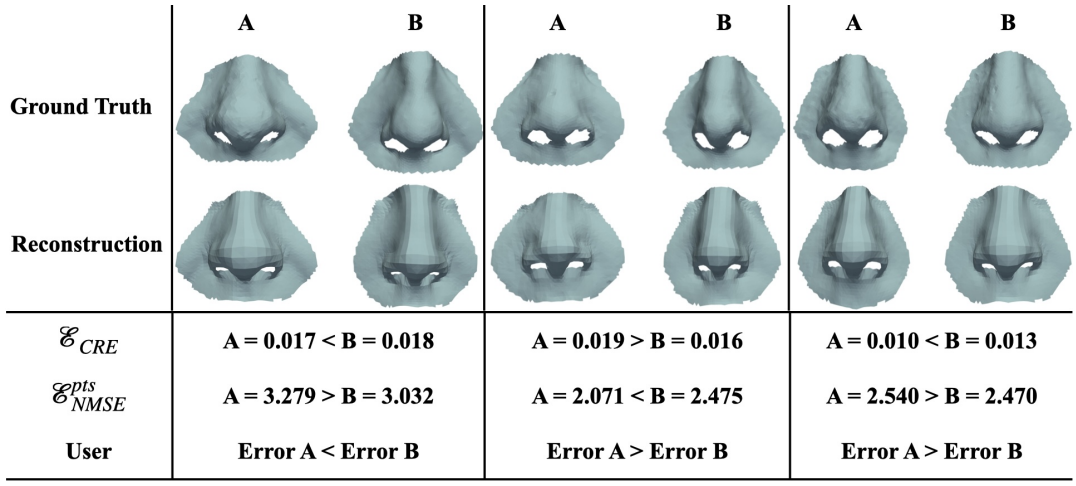


Figure 3 | Example of three pairs with their metrics and the user choice.

Table 1 | Correspondence of reconstruction metrics with human judgments. Top: comparison with existing metrics. Bottom: ablation over $\mathcal{E}_{CRE}$ hyperparameters.

Metric	Mapping	R	K	Norm	Pairwise-Accuracy, %	Pairwise-ROC-AUC, %
$\mathcal{E}_{NMSE}^{REALY}$	REALY	–	–	–	53.92	56.53
$\mathcal{E}_{NMSE}^{pts}$	Points-to-surface	–	–	–	49.32	55.02
$\mathcal{E}_{NMSE}^{vtx}$	Vertex-to-vertex	–	–	–	42.03	42.47
$\mathcal{E}_{CRE}$	Points-to-surface	**3**	**128**	L_1	**73.63**	**82.66**
$\mathcal{E}_{CRE}$	Points-to-surface	1	128	L_1	71.38	80.63
$\mathcal{E}_{CRE}$	Points-to-surface	1	512	L_1	68.32	78.17
$\mathcal{E}_{CRE}$	Points-to-surface	1	1024	L_1	66.52	74.35
$\mathcal{E}_{CRE}$	Points-to-surface	2	128	L_1	72.46	81.86
$\mathcal{E}_{CRE}$	Points-to-surface	2	512	L_1	68.77	77.46
$\mathcal{E}_{CRE}$	Points-to-surface	2	1024	L_1	66.79	77.46
$\mathcal{E}_{CRE}$	Points-to-surface	3	512	L_1	70.66	79.36
$\mathcal{E}_{CRE}$	Points-to-surface	3	1024	L_1	69.85	78.83
$\mathcal{E}_{CRE}$	Points-to-surface	4	128	L_1	73.00	81.85
$\mathcal{E}_{CRE}$	Points-to-surface	4	512	L_1	70.93	79.40
$\mathcal{E}_{CRE}$	Points-to-surface	4	1024	L_1	70.66	79.17
$\mathcal{E}_{CRE}$	Points-to-surface	3	128	L_2	72.28	81.17
$\mathcal{E}_{CRE}$	Points-to-surface	3	512	L_2	69.67	77.81
$\mathcal{E}_{CRE}$	Points-to-surface	3	1024	L_2	68.14	76.99
$\mathcal{E}_{CRE}$	Points-to-surface	3	128	*Huber*	67.60	75.82
$\mathcal{E}_{CRE}$	Points-to-surface	3	512	*Huber*	62.74	70.40
$\mathcal{E}_{CRE}$	Points-to-surface	3	1024	*Huber*	61.57	69.51

These results reveal a significant representation bias: as age increases, the 3DMM basis and reconstruction models fails to capture high-frequency geometric details, effectively "smoothing out" the structural markers characteristic of older individuals. fig. 4 illustrates such effects. Moreover, gender-stratified analysis shows that this age-related degradation is prevalent in both populations, though it is notably intensified in males within the BFM basis (*e.g.*, $r = 0.88$ for the cheek region). To further confirm the benefits of a curvature-based approach, we repeated this analysis using the standard Euclidean-based NMSE and observed no significant bias (see the Supplementary Material).

4.3.2. Ethnicity with respect to Error Clusters.

Given the very limited number of samples per ethnicity (except Caucasians), we moved beyond global error analysis and clustered regional error maps ϵ_H as described in algorithm 1.

Table 2 | Pearson correlation coefficients between age and CRE across different regions. Strong correlations (≥ 0.5) are highlighted in bold.

		Pearson correlation: $\mathcal{E}_{CRE}$ *vs.* 3D Reconstruction											
		@nose			@mouth			@cheek			@forehead		
		All	Female	Male	All	Female	Male	All	Female	Male	All	Female	Male
3DMM Repr.	BFM	0.39	0.32	**0.59**	0.33	0.21	**0.67**	**0.77**	**0.72**	**0.88**	**0.52**	0.49	**0.59**
	FLAME	**0.78**	**0.79**	**0.78**	-0.03	-0.16	0.18	**0.82**	**0.84**	**0.82**	0.41	0.48	0.25
Model from 2D image	Deep3D	0.24	0.18	0.41	-0.23	-0.19	-0.38	**0.69**	**0.81**	**0.62**	**0.58**	**0.54**	**0.66**
	3DDFA-v3	**0.67**	**0.65**	**0.73**	0.31	0.30	0.37	**0.79**	**0.78**	**0.88**	**0.53**	**0.55**	0.48
	MGCNet	**0.54**	**0.50**	**0.67**	0.30	0.30	0.32	**0.80**	**0.79**	**0.89**	**0.51**	**0.54**	0.47
	DECA	**0.73**	**0.75**	**0.70**	-0.32	-0.30	-0.35	**0.76**	**0.80**	**0.85**	0.40	0.46	0.28
	MICA	**0.69**	**0.72**	**0.65**	-0.31	-0.31	-0.30	**0.79**	**0.76**	**0.88**	0.43	0.48	0.31

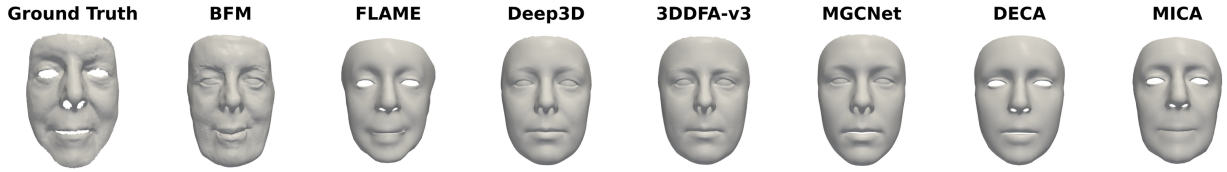


Figure 4 | Example of age bias on a senior person from REALY benchmark.

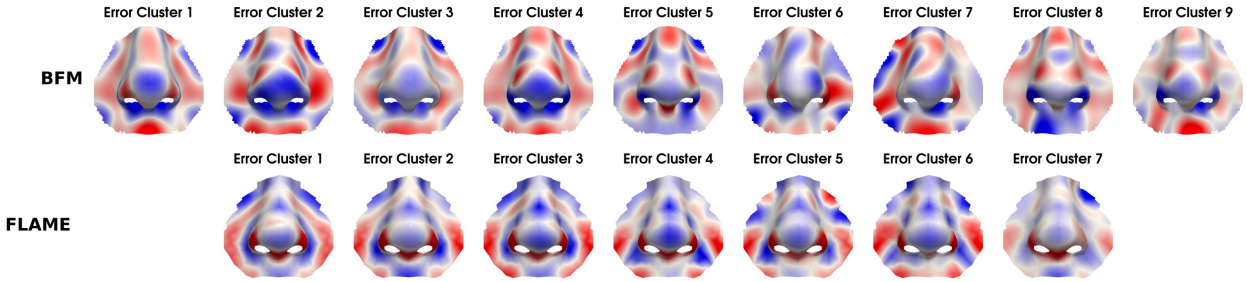
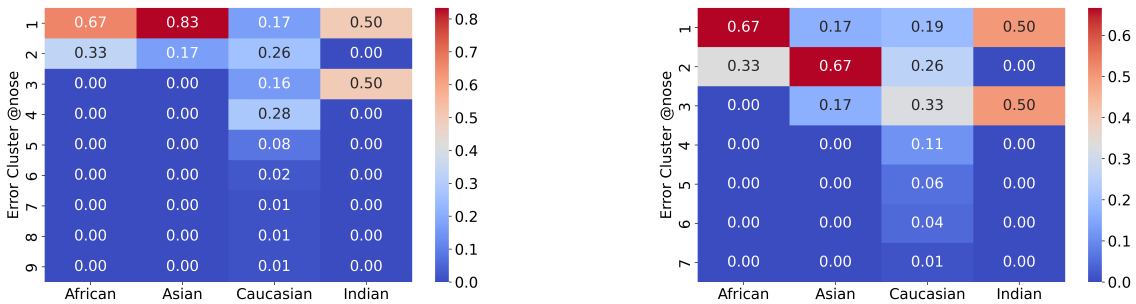


Figure 5 | Error cluster centroids found for BFM (top row) and FLAME (bottom row).

fig. 5 illustrates 9 (resp. 7) unique failure modes that were identified for BFM (resp. FLAME) reconstructions for the nose zone (see Supplementary Material for other anatomical zones). fig. 6 summarizes the distributions of these failure modes for each ethnicity. Whereas Caucasian subjects are distributed across different error clusters, African and Asian populations are primarily concentrated at the first error cluster for BFM and at the first and second error clusters for FLAME. This might suggest that these populations suffer from a distinct systematic error pattern produced by these 3DMMs. fig. 7 illustrates this pattern for BFM reconstructions. We observe that the nose wings of reconstructions have unnatural creases and the tip of the nose is drawn up and more round. The linear subspaces of both these 3DMMs lack the requisite expressive capacity to capture the distinct geometric nuances found in globally diverse populations, leading to a systematic "flattening" of ethnic-specific features toward the model mean.

4.3.3. Morphology Traits with respect to Error Clusters.

To demonstrate that 3DMM failure modes are systematic rather than stochastic, we investigate the intrinsic relationship between specific morphologies and curvature-based reconstruction errors. To visualize the morphologies associated to each error cluster, we computed a PCA on the REALY ground truth faces. Such decomposition mirrors the statistical formulation used to construct a 3DMM basis,



(a) BFM [29] (b) FLAME [23]

Figure 6 | Distribution of failure modes @nose for each ethnicity in REALY benchmark.

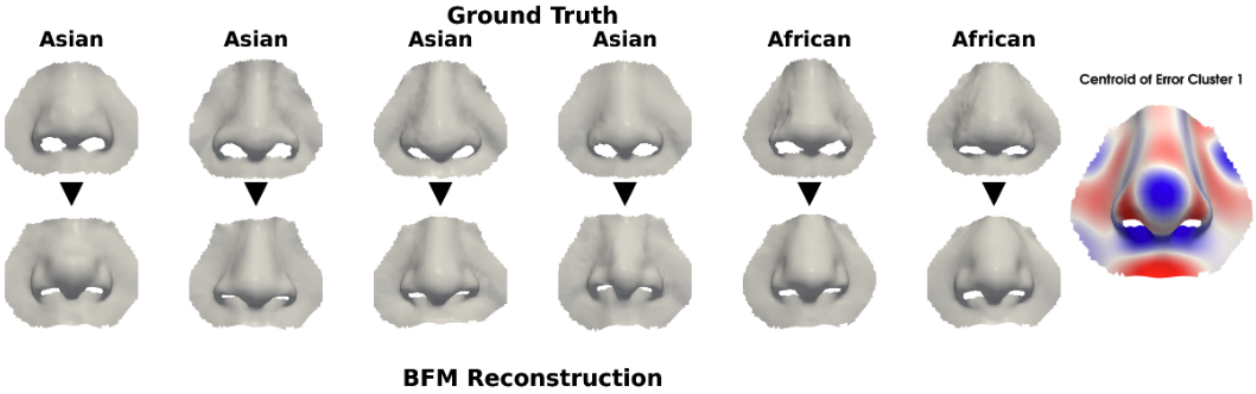


Figure 7 | Example of African and Asian subjects affected by a failure mode 1 of BFM [29].

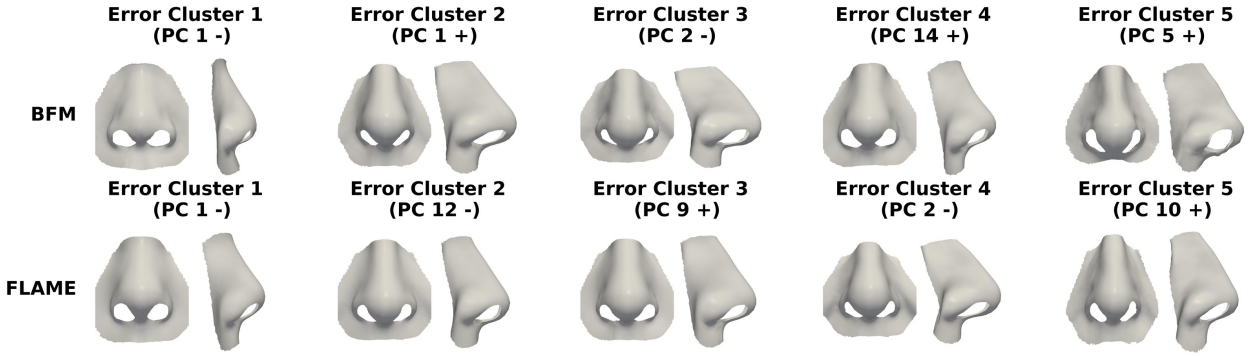


Figure 8 | Centroids for error clusters identified for BFM [29] and FLAME [23]. In $(\cdot)$ we put the *PC* with maximal mean absolute value for the error cluster and its sign.

with the difference that we focused on a region of interest (more details are given in the Supplementary Material). For each error cluster, we identify the dominant principal component – the dimension that primarily governs the cluster's morphology – defined as the component with the maximal mean absolute coefficient value. Furthermore, we analyze the sign of these mean coefficients to determine the direction of morphological deviation relative to the mean shape.

The results for five representative error clusters in both BFM and FLAME are illustrated in fig. 8 (refer to the Supplementary Material for an exhaustive analysis). Our observations confirm that each error pattern maps to a distinct, well-defined morphological phenotype. For example, within the BFM framework, Error Cluster 1 is driven by PC_1, which encodes variations in nasal length. A negative coefficient sign indicates that this cluster is predominantly composed of individuals with significantly shorter nasal structures than the population mean. This direct mapping suggests that 3DMM biases are not random artifacts of the fitting process, but are fundamentally rooted in the model's inability to span specific regions of the morphological manifold.

5. Discussion and Conclusion

We proposed a new curvature-aware paradigm for analyzing 3D face reconstruction errors, enabling us to define: (1) an error map for localized geometric visualization; (2) a scalar metric that is significantly better aligned with human perception than standard Euclidean-based metrics; and (3) a clustering of errors to discover unknown biases.

Our fairness analysis reveals that both evaluated 3DMMs exhibit pronounced age-related biases, even in the idealized case of projecting the full 3D geometric ground truth into the 3DMM basis. We further find preliminary evidence of gender, ethnic, and morphological biases. This suggests that the reconstruction models' representational bottlenecks are inherent to their linear prediction subspaces rather than a failure of 2D-to-3D inference, perpetuating demographic disparities regardless of the fitting methodology.

As future work, we plan to investigate how our error map can be used as a differentiable loss function for a neural network for 3D reconstruction. This curvature-aware objective would drive a secondary refinement module, in order to go beyond the raw 3DMM basis limitations.

References

[1] Regulation (eu) 2024/1689 of the european parliament and of the council of 13 june 2024 laying down harmonised rules on artificial intelligence and amending regulations (ec) no 300/2008, (eu) no 167/2013, (eu) no 168/2013, (eu) 2018/858, (eu) 2018/1139 and (eu) 2019/2144 and directives 2014/90/eu, (eu) 2016/797 and (eu) 2020/1828 (artificial intelligence act) (text with eea relevance) (Jul 2024)

[2] Amberg, B., Romdhani, S., Vetter, T.: Optimal step nonrigid icp algorithms for surface registration. In: 2007 IEEE conference on computer vision and pattern recognition. pp. 1–8. IEEE (2007)

[3] Besl, P.J., McKay, N.D.: Method for registration of 3-d shapes. In: Sensor fusion IV: control paradigms and data structures. vol. 1611, pp. 586–606. Spie (1992)

[4] Besse, P., del Barrio, E., Gordaliza, P., Loubes, J.M., Risser, L.: A survey of bias in machine learning through the prism of statistical parity. The American Statistician **76**(2), 188–198 (2022). https://doi.org/10.1080/00031305.2021.1952897

[5] Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques. p. 187–194. SIGGRAPH '99, ACM Press/Addison-Wesley Publishing Co., USA (1999)

[6] Blanz, V., Vetter, T.: Face recognition based on fitting a 3d morphable model. IEEE Transactions on pattern analysis and machine intelligence **25**(9), 1063–1074 (2003)

[7] Botsch, M., Kobbelt, L., Pauly, M., Alliez, P., Lévy, B.: Polygon mesh processing. CRC press (2010)

[8] Chai, Z., Zhang, H., Ren, J., Kang, D., Xu, Z., Zhe, X., Yuan, C., Bao, L.: Realy: Rethinking the evaluation of 3d face reconstruction. In: European conference on computer vision. pp. 74–92. Springer (2022)

[9] Cook, C.M., Howard, J.J., Sirotin, Y.B., Tipton, J.L., Vemury, A.R.: Demographic effects in facial recognition and their dependence on image acquisition: An evaluation of eleven commercial systems. IEEE Transactions on Biometrics, Behavior, and Identity Science **1**(1), 32–41 (2019)

[10] Cui, Z., Yu, J.: Boosting fairness for 3d face reconstruction. In: 2024 International Joint Conference on Neural Networks (IJCNN). pp. 1–9. IEEE (2024)

[11] Dai, H., Pears, N., Smith, W., Duncan, C.: Statistical modeling of craniofacial shape and texture. International Journal of Computer Vision **128**(2), 547–571 (2020)

[12] Deng, Y., Yang, J., Xu, S., Chen, D., Jia, Y., Tong, X.: Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops. pp. 0–0 (2019)

[13] Egger, B., Smith, W.A., Tewari, A., Wuhrer, S., Zollhoefer, M., Beeler, T., Bernard, F., Bolkart, T., Kortylewski, A., Romdhani, S., et al.: 3d morphable face models—past, present, and future. ACM Transactions on Graphics (ToG) **39**(5), 1–38 (2020)

[14] Feng, Y., Feng, H., Black, M.J., Bolkart, T.: Learning an animatable detailed 3d face model from in-the-wild images. ACM Transactions on Graphics (ToG) **40**(4), 1–13 (2021)

[15] Feng, Z.H., Huber, P., Kittler, J., Hancock, P., Wu, X.J., Zhao, Q., Koppen, P., Rätsch, M.: Evaluation of dense 3d reconstruction from 2d face images in the wild. In: 2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018). pp. 780–786. IEEE (2018)

[16] Gatzke, T., Grimm, C.: Assessing curvature metrics on triangular meshes. Washington University, St. Louis Missouri, Tech. Rep (2003)

[17] Gerogiannis, D., Papantoniou, F.P., Potamias, R.A., Lattas, A., Zafeiriou, S.: Arc2avatar: Generating expressive 3d avatars from a single image via id guidance. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 10770–10782 (2025)

[18] Grother, P., Ngan, M., Hanaoka, K.: Face recognition vendor test (fvrt): Part 3, demographic effects. National Institute of Standards and Technology Gaithersburg, MD (2019)

[19] Hamann, B.: Curvature approximation for triangulated surfaces. In: Geometric modelling, pp. 139–153. Springer (1992)

[20] Hu, L., Saito, S., Wei, L., Nagano, K., Seo, J., Fursund, J., Sadeghi, I., Sun, C., Chen, Y.C., Li, H.: Avatar digitization from a single image for real-time rendering. ACM Transactions on Graphics (ToG) **36**(6), 1–14 (2017)

[21] Leelavathi, M., Kannan, D.: Enhanced 3d face anti-spoofing for secure biometric authentication using spatio-temporal deep learning networks (stdl-net) with robust feature representation and temporal consistency modeling. Lex Localis: Journal of Local Self-Government **23** (2025)

[22] Lehoucq, R.B., Sorensen, D.C., Yang, C.: ARPACK users' guide: solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods. SIAM (1998)

[23] Li, T., Bolkart, T., Black, M.J., Li, H., Romero, J.: Learning a model of facial shape and expression from 4d scans. ACM Trans. Graph. **36**(6), 194–1 (2017)

[24] Lombardi, S., Saragih, J., Simon, T., Sheikh, Y.: Deep appearance models for face rendering. ACM Transactions on Graphics (ToG) **37**(4), 1–13 (2018)

[25] MacQueen, J.: Some methods for classification and analysis of multivariate observations. In: Proceedings of Berkeley Symposium on Mathematical Statistics and Probability/University of California Press (1967)

[26] Meek, D.S., Walton, D.J.: On surface normal and gaussian curvature approximations given data sampled from a smooth surface. Computer aided geometric design **17**(6), 521–543 (2000)

[27] Mueller, A., Paysan, P., Schumacher, R., Zeilhofer, H.F., Berg-Boerner, B.I., Maurer, J., Vetter, T., Schkommodau, E., Juergens, P., Schwenzer-Zimmerer, K.: Missing facial parts computed by a morphable model and transferred directly to a polyamide laser-sintered prosthesis: an innovation study. British journal of oral and maxillofacial surgery **49**(8), e67–e71 (2011)

[28] Myers, J.L., Well, A.D., Lorch Jr, R.F.: Research design and statistical analysis. Routledge (2013)

[29] Paysan, P., Knothe, R., Amberg, B., Romdhani, S., Vetter, T.: A 3d face model for pose and illumination invariant face recognition. In: 2009 sixth IEEE international conference on advanced video and signal based surveillance. pp. 296–301. Ieee (2009)

[30] Sanyal, S., Bolkart, T., Feng, H., Black, M.J.: Learning to regress 3d face shape and expression from an image without 3d supervision. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 7763–7772 (2019)

[31] Sariyanidi, E., Ferrari, C., Nocentini, F., Berretti, S., Cavallaro, A., Tunc, B.: 3d face reconstruction error decomposed: A modular benchmark for fair and fast method evaluation. In: 2025 IEEE 19th International Conference on Automatic Face and Gesture Recognition (FG). pp. 1–10. IEEE (2025)

[32] Shang, J., Shen, T., Li, S., Zhou, L., Zhen, M., Fang, T., Quan, L.: Self-supervised monocular 3d face reconstruction by occlusion-aware multi-view geometry consistency. In: European Conference on Computer Vision. pp. 53–70. Springer (2020)

[33] Sharp, N., Crane, K.: A laplacian for nonmanifold triangle meshes. In: Computer Graphics Forum. vol. 39, pp. 69–80. Wiley Online Library (2020)

[34] Shilova, V., Malherbe, E., Palma, G., Risser, L., Loubes, J.M.: Fairness-aware grouping for continuous sensitive variables: Application for debiasing face analysis with respect to skin tone. arXiv preprint arXiv:2507.11247 (2025)

[35] Thies, J., Zollhofer, M., Stamminger, M., Theobalt, C., Nießner, M.: Face2face: Real-time face capture and reenactment of rgb videos. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2387–2395 (2016)

[36] Thies, J., Zollhöfer, M., Stamminger, M., Theobalt, C., Nießner, M.: Facevr: Real-time facial reenactment and eye gaze control in virtual reality. arXiv preprint arXiv:1610.03151 (2016)

[37] Wang, Z., Zhu, X., Zhang, T., Wang, B., Lei, Z.: 3d face reconstruction with the geometric guidance of facial part segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1672–1682 (2024)

[38] Yamaguchi, S., Saito, S., Nagano, K., Zhao, Y., Chen, W., Olszewski, K., Morishima, S., Li, H.: High-fidelity facial reflectance and geometry inference from an unconstrained image. ACM Transactions on Graphics (TOG) **37**(4), 1–14 (2018)

[39] Yang, X., Taketomi, T., Endo, Y., Kanamori, Y.: Makeup prior models for 3d facial makeup estimation and applications. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2165–2176 (2024)

[40] Zielonka, W., Bolkart, T., Thies, J.: Towards metrical reconstruction of human faces. In: European conference on computer vision. pp. 250–269. Springer (2022)