

# Statistical Gradient Filtering for Geometry Optimization Under Limited Observations

WONJONG JANG, POSTECH, South Korea  
GWANGJIN JU, POSTECH, South Korea  
SEUNGYONG LEE, POSTECH, South Korea

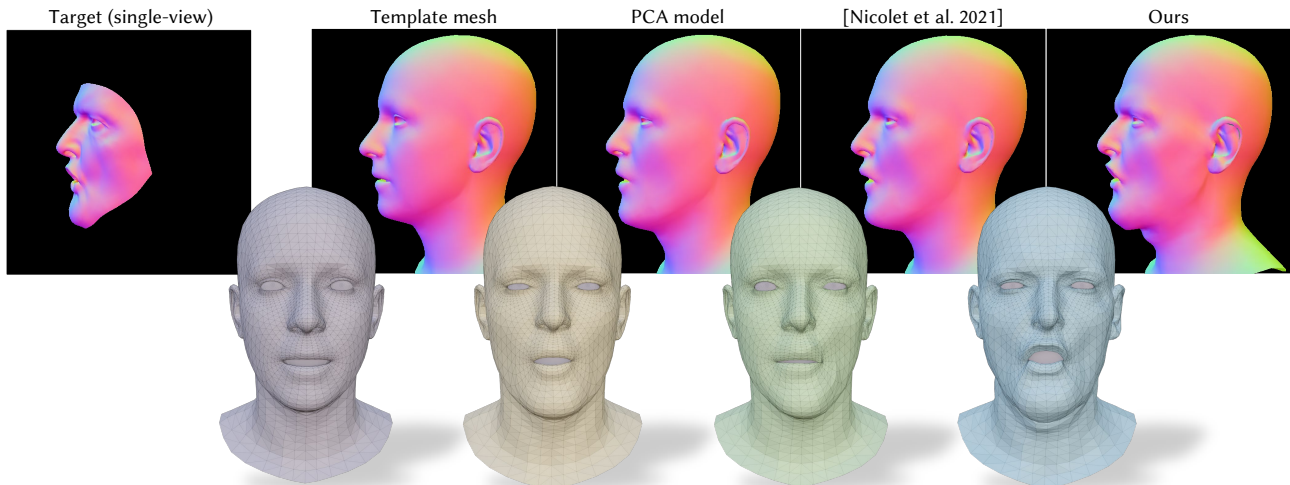


Fig. 1. Statistical gradient filtering incorporates statistical shape priors into gradient descent for template-to-target mesh deformation. PCA-based methods produce plausible shapes but restrict solutions to the linear subspace, limiting their ability to capture target geometry. Laplacian-based gradient filtering [Nicolet et al. 2021] allows unrestricted deformation but relies solely on spatial smoothness, providing limited guidance for unobserved regions. Our statistical gradient filtering combines the strengths of both, offering statistical guidance for unobserved regions while allowing out-of-distribution solutions. target: ©3DScanStore

Geometry optimization often encounters sparse gradients from limited observations, which can cause optimization to drift toward unnatural shapes. Prior work stabilizes this process by exploiting spatial structure encoded by the Laplacian operator *within a single shape*, enforcing spatial smoothness on gradients. However, such smoothness alone propagates gradients to unobserved regions without any prior knowledge of how shapes typically deform in a given domain. We introduce *statistical gradient filtering*, which leverages statistical structure *across a shape collection* by learning shape variations via principal component analysis (PCA) and guiding geometry updates along directions consistent with this learned prior. Unlike previous PCA-based methods that constrain solutions to a linear subspace or modify the objective with regularization terms, we filter gradients at each iteration to steer the optimization path toward plausible shapes, without restricting the solution space or altering the original objective. We validate our approach across a range of shape optimization tasks, demonstrating robust convergence even under challenging conditions.

CCS Concepts: • **Computing methodologies** → **Shape modeling**; • **Mathematics of computing** → **Continuous optimization**.

Authors' Contact Information: Wonjong Jang, POSTECH, South Korea, wonjong@postech.ac.kr; Gwangjin Ju, POSTECH, South Korea, gwangjin@postech.ac.kr; Seungyong Lee, POSTECH, South Korea, leesy@postech.ac.kr.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

SIGGRAPH Conference Papers '26, Los Angeles, CA, USA

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2554-8/2026/07

<https://doi.org/10.1145/3799902.3811084>

Additional Key Words and Phrases: Geometry Deformation, Gradient Filtering

## ACM Reference Format:

Wonjong Jang, Gwangjin Ju, and Seungyong Lee. 2026. Statistical Gradient Filtering for Geometry Optimization Under Limited Observations. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '26)*, July 19–23, 2026, Los Angeles, CA, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3799902.3811084>

## 1 Introduction

Geometry optimization toward a target shape is challenging when observations are limited. In single-view inverse rendering, for instance, occluded regions receive no gradients for updates, leaving much of the surface without guidance and causing optimization to drift toward an unnatural shape. To address this challenge, additional structural information needs to be integrated into the optimization process to propagate the observed gradients to unobserved regions. Consequently, the specific choice of this structural information dictates how these regions are updated during optimization.

A prominent choice is the mesh Laplacian [Dziuk 1988; Pinkall and Polthier 1993], which encodes local geometric relationships among neighboring vertices within a single shape. Building on this structure, prior work [Jung et al. 2023; Nicolet et al. 2021] smoothly propagates observed gradients to neighboring vertices, thereby improving both convergence and output mesh quality. However, since

these approaches guide only spatial smoothness, they extrapolate gradients to unobserved regions without any prior knowledge of how shapes typically deform in a given domain (Fig. 1).

In this work, we propose *statistical gradient filtering*, a method that incorporates prior knowledge learned from a collection of shapes. Unlike Laplacian-based filtering, which exploits structure within a single shape, our approach captures statistical structure across a shape collection. Specifically, we construct an eigenspace via principal component analysis (PCA) that encodes how shapes vary within the collection. At each optimization step, we modulate the gradient through this eigenspace, steering the descent direction to favor geometry updates consistent with learned shape variations. Rather than relying on spatial smoothness, this approach extrapolates geometry updates to unobserved regions based on statistical correlations learned from diverse shapes.

Shape priors from statistical models have long been used to guide geometry deformation, but typically through parametric or regularization approaches. Parametric methods [Blanz and Vetter 1999; Li et al. 2017; Paysan et al. 2009] restrict solutions to the linear span of principal components, sacrificing expressiveness for solution plausibility. Regularization methods [Li et al. 2017] retain the full solution space but alter the objective function by penalizing deviations from the statistical subspace, which biases solutions and requires careful parameter tuning. Our statistical gradient filtering avoids both limitations; it allows solutions to deviate from the linear span and preserves the original objective while guiding optimization along paths informed by shape statistics. Statistical gradient filtering thus offers a different paradigm that neither restricts nor changes where optimization can reach, but guides *the path optimization takes* toward plausible shapes.

We validate statistical gradient filtering<sup>1</sup> across diverse applications, including non-rigid registration on partial scans, inverse rendering from a single-view surface normal map, texture map inpainting, and Gaussian splatting on meshes. Across all settings, our method achieves higher-quality results compared to existing alternatives.

## 2 Related Work

### 2.1 Gradient Filtering in Geometry Deformation

Gradient filtering stabilizes and accelerates optimization by transforming sparse or noisy gradients into smoother forms. A classical approach uses the inverse Hessian as a gradient filter to achieve second-order convergence, as in Newton’s method [Polyak 2007], but computing this inverse is often expensive or intractable for geometry problems such as inverse rendering. Instead, alternative gradient filtering strategies have been developed in the form of preconditioning to address specific geometric challenges. These include isometry-aware preconditioning [Claici et al. 2017], multi-scale methods [Chen et al. 2021; Martin et al. 2013], preconditioned conjugate gradient for physical simulation [Shen et al. 2024], and stochastic preconditioning for neural fields [Ling et al. 2025].

A widely adopted approach is Laplacian-based gradient filtering, introduced by Nicolet et al. [2021] for inverse rendering. Their

method derives a gradient filter from a Laplacian smoothing regularizer, a special case of Sobolev preconditioning [Arnold et al. 1997; Neuberger 1985], diffusing sparse gradients through the mesh Laplacian to improve both convergence and output mesh quality. Jung et al. [2023] later extended this framework to non-rigid registration, and similar Laplacian-based filters have been adopted for a wide range of problems, including mesh optimization [Kovalsky et al. 2016; Zhu et al. 2018], self-intersection handling [Jang et al. 2025; Yu et al. 2021] with shape-space augmentation [Sassen et al. 2024], inverse rendering via an adaptively weighted Laplacian [An et al. 2023], 3D reconstruction [Kalthheuner et al. 2025; Slavcheva et al. 2018], and non-rigid registration [Jung et al. 2025], as well as image processing and partial differential equation solvers [Karatson and Loczi 2005; Osher et al. 2022; Park et al. 2021]. Despite this widespread adoption, Laplacian-based gradient filtering relies solely on spatial information within a single shape; when observations are limited, the Laplacian cannot provide meaningful guidance for unobserved regions. Our work addresses this limitation by introducing statistical gradient filtering, which guides optimization toward configurations consistent with shape variations learned from a shape collection.

Our statistical gradient filtering is built upon PCA, which has also been used for gradient filtering in prior work, but in a different manner. Existing approaches typically construct gradient filters online from solution snapshots during iterative solves [Astrid et al. 2011; Cortes et al. 2018; Luo and Cai 2023; Nabben and Vuik 2006]. In contrast, our method leverages pre-computed PCA bases from a shape collection as shape priors, enabling statistical guidance for geometry optimization without requiring accumulated solution history.

### 2.2 Statistical Shape Modeling

Statistical shape models learn compact representations of shape variations from training data. Beginning with Active Shape Models [Cootes et al. 1995] and Active Appearance Models [Cootes et al. 1998], this approach was extended to the 3D Morphable Model (3DMM) [Blanz and Vetter 1999], which enabled photorealistic face synthesis by modeling geometry and texture in a PCA space. Subsequent efforts have focused on building larger and more expressive face models [Booth et al. 2016; Cao et al. 2013; Gerig et al. 2018; Huber et al. 2016; Paysan et al. 2009], with FLAME [Li et al. 2017] achieving state-of-the-art expressiveness by learning pose-dependent deformations from 4D sequences. For body modeling, SCAPE [Angelov et al. 2005] introduced data-driven pose and shape deformation. Subsequently, SMPL [Loper et al. 2015] and its extension SMPL-X [Pavlakos et al. 2019] have become widely adopted for full-body capture.

While these PCA-based statistical models serve as powerful shape priors, they are typically applied through parametric methods that optimize within the constructed PCA space. This approach presents a distinct limitation, as restricting solutions to the learned subspace prevents the recovery of novel geometric features. In contrast, our method leverages the statistical structure to filter sparse gradients, enabling stable convergence to natural shapes while allowing solutions beyond the learned subspace.

<sup>1</sup>Our code is publicly available at <https://github.com/wonjongg/stat-grad>

### 3 Background

Assume we are given a PCA model [Jolliffe 1986] that represents the primary geometric variations of a shape collection  $\{s_i\}_{i=1}^N$ . This model comprises a mean shape  $\bar{\mathbf{x}}$  and the top  $k$  principal components  $\mathbf{U}_k$ , obtained by eigen-decomposition of the covariance matrix  $\mathbf{C} = \frac{1}{N-1} \sum_i (s_i - \bar{\mathbf{x}})(s_i - \bar{\mathbf{x}})^T$ . We then consider the problem of deforming a 3D mesh from the mean shape (or a template mesh) to minimize an objective function:

$$\underset{\mathbf{x} \in \mathbb{R}^{n \times 3}}{\text{minimize}} \Phi(\mathbf{x}), \quad (1)$$

where  $\mathbf{x}$  represents the deformed vertex positions and  $\Phi(\cdot)$  is an objective function that typically incorporates sparse observations, such as a partial scan with occlusions. When observations are limited, many vertices receive no gradient signal, yet must still deform plausibly. We review existing approaches to this problem: PCA-based parametric and regularization methods, and Laplacian-based gradient filtering.

#### 3.1 Parametric Method

To guarantee a valid solution, parametric approaches [Blanz and Vetter 1999; Li et al. 2017] restrict the solution space to the PCA subspace by optimizing coefficients  $\boldsymbol{\alpha}$  rather than the full shape  $\mathbf{x}$ . In the optimization process, each gradient descent iteration updates  $\boldsymbol{\alpha}$  using

$$\boldsymbol{\alpha} \leftarrow \boldsymbol{\alpha} - \eta \mathbf{U}_k^T \frac{\partial \Phi}{\partial \mathbf{x}}, \quad \text{where } \mathbf{x} = \bar{\mathbf{x}} + \mathbf{U}_k \boldsymbol{\alpha}, \quad (2)$$

where  $\eta$  is a step size. While this ensures solutions remain within the learned PCA subspace, it fundamentally limits expressiveness as the optimizer cannot represent novel deformations outside the span of the training data.

#### 3.2 PCA-based Regularization

Regularization methods [Li et al. 2017; Zhang and Schneider 2010] preserve full solution expressiveness by adding a penalty term [Moghadam and Pentland 1997; Pentland et al. 1994] that discourages deviations from the PCA subspace. The regularized objective is:

$$\Phi_{\text{reg}}(\mathbf{x}) = \Phi(\mathbf{x}) + \frac{\beta}{2} \|(\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^T)(\mathbf{x} - \bar{\mathbf{x}})\|^2, \quad (3)$$

where  $\beta$  controls the regularization strength and  $(\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^T)$  projects  $(\mathbf{x} - \bar{\mathbf{x}})$  onto the null space of  $\mathbf{U}_k \mathbf{U}_k^T$  to measure deviations from the PCA subspace. While this approach enables diverse solutions, it changes the original objective function.

Moreover, the gradient of this objective is:

$$\frac{\partial \Phi_{\text{reg}}}{\partial \mathbf{x}} = \frac{\partial \Phi}{\partial \mathbf{x}} + \beta(\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^T)(\mathbf{x} - \bar{\mathbf{x}}), \quad (4)$$

where the gradient  $\frac{\partial \Phi}{\partial \mathbf{x}}$  can be noisy due to sparse or limited observations. Early in optimization, when the geometry remains close to the mean, the regularization term is weak, allowing unreliable gradients  $\frac{\partial \Phi}{\partial \mathbf{x}}$  to directly corrupt the geometry. The regularization only grows as the geometry deviates, attempting to correct distortions after they have already occurred. However, once the geometry has been corrupted, the regularization may project the solution into an undesirable local minimum that does not correspond to a plausible shape




Fig. 2. Effect of gradient descent, regularization, and gradient filtering. Naive gradient descent applies noisy updates that corrupt geometry. Regularization tries to repair such distortions but only after geometry has already been corrupted. Before noisy gradients corrupt geometry, our method filters them to ensure only meaningful updates. target: ©3DScanStore

(Fig. 2). To prevent such corruption, unreliable gradients should be filtered before they are applied to the geometry.

#### 3.3 Laplacian-based Gradient Filtering

Laplacian-based gradient filtering [Jung et al. 2023; Nicolet et al. 2021] improves mesh quality by diffusing gradients through the mesh Laplacian  $\mathbf{L}$ . This gradient filter arises from a regularized objective with Laplacian smoothing [Desbrun et al. 1999; Taubin 1995]:

$$\underset{\mathbf{x} \in \mathbb{R}^{n \times 3}}{\text{minimize}} \Phi(\mathbf{x}) + \frac{\lambda}{2} \mathbf{x}^T \mathbf{L} \mathbf{x}, \quad (5)$$

where  $\lambda$  controls the amount of smoothing.

The approximate Hessian of this objective becomes  $(\mathbf{I} + \lambda \mathbf{L})$  when the Hessian of  $\Phi$  is approximated as identity. Using the inverse of this approximate Hessian as a gradient filter yields the geometry update:

$$\mathbf{x} \leftarrow \mathbf{x} - \eta(\mathbf{I} + \lambda \mathbf{L})^{-1} \frac{\partial \Phi}{\partial \mathbf{x}}. \quad (6)$$

This operation can be understood through the eigen-decomposition  $\mathbf{L} = \mathbf{V} \boldsymbol{\Lambda} \mathbf{V}^T$ , where  $\mathbf{V}$  contains eigenvectors representing spatial frequency modes and  $\boldsymbol{\Lambda} = \text{diag}(\mu_1, \dots, \mu_n)$  contains the corresponding eigenvalues. The operator can be expressed as:

$$(\mathbf{I} + \lambda \mathbf{L})^{-1} = \mathbf{V}(\mathbf{I} + \lambda \boldsymbol{\Lambda})^{-1} \mathbf{V}^T = \mathbf{V} \mathbf{D}_\lambda \mathbf{V}^T, \quad (7)$$

where  $\mathbf{D}_\lambda = \text{diag}\left(\frac{1}{1 + \lambda \mu_i}\right)_{i=1}^n$  applies different weights to each eigenmode. This operation filters the gradient in the frequency domain; high frequency components with large eigenvalues  $\mu_i$  are strongly attenuated, while low frequency components pass through with minimal attenuation. This approach produces spatially smooth deformations, without leveraging prior knowledge of how shapes typically deform in a given domain.

### 4 Statistical Gradient Filtering

We now introduce our method, *statistical gradient filtering*, for prior-guided geometry deformation under limited observations. Unlike parametric or regularization approaches that constrain solutions or modify objectives, we transform gradient descent trajectories while allowing out-of-distribution solutions and preserving the original objective. Similar to Eq. (6), we derive our gradient filter from the Hessian of the regularized objective, then simplify it through approximation and reparameterization to obtain a practical form

that can be intuitively integrated in automatic differentiation frameworks [Paszke et al. 2019]. We highlight our main gradient filtering formulations in the emerald-colored blocks.

#### 4.1 Gradient Filtering Derivation

Our starting point is the Hessian of the regularized objective from Section 3.2. Differentiating the gradient of  $\Phi_{\text{reg}}(\mathbf{x})$  yields:

$$\frac{\partial^2 \Phi_{\text{reg}}}{\partial \mathbf{x}^2} = \frac{\partial^2 \Phi}{\partial \mathbf{x}^2} + \beta(\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^T). \quad (8)$$

The first term captures the second-order information of the original objective and the second encodes statistical structure. Since our goal is to leverage the statistical structure, and computing  $\frac{\partial^2 \Phi}{\partial \mathbf{x}^2}$  is often expensive or intractable, we approximate the exact second-order term with the identity matrix  $\mathbf{I}$  and retain only the statistical term, yielding:

$$\mathbf{H}_{\text{approx}} = \mathbf{I} + \beta(\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^T) = (1 + \beta)\mathbf{I} - \beta \mathbf{U}_k \mathbf{U}_k^T. \quad (9)$$

Following an approximate Newton approach, we filter gradients by the inverse of this approximate Hessian. To compute the inverse, our key observation is that  $\mathbf{H}_{\text{approx}}$  is a linear combination of  $\mathbf{I}$  and the projection  $\mathbf{U}_k \mathbf{U}_k^T$ . Since the projection satisfies the idempotency property  $(\mathbf{U}_k \mathbf{U}_k^T)^2 = \mathbf{U}_k \mathbf{U}_k^T$ , matrix inversion preserves this algebraic structure, as shown in Sec. 1.1 of the supplementary document. We therefore seek an inverse of the same form:

$$\mathbf{H}_{\text{approx}}^{-1} = a\mathbf{I} + b\mathbf{U}_k \mathbf{U}_k^T, \quad (10)$$

where  $a$  and  $b$  are scalars to be determined.

To find  $a$  and  $b$ , we require that multiplying by  $\mathbf{H}_{\text{approx}}$  yields the identity matrix  $\mathbf{I}$ . After expanding the equation, we obtain:

$$a(1 + \beta)\mathbf{I} + [b(1 + \beta) - \beta(a + b)]\mathbf{U}_k \mathbf{U}_k^T = \mathbf{I}. \quad (11)$$

Matching coefficients gives  $a(1 + \beta) = 1$  and  $b(1 + \beta) - \beta(a + b) = 0$ , which yield:

$$a = \frac{1}{1 + \beta}, \quad b = \frac{\beta}{1 + \beta}. \quad (12)$$

Thus, the statistical gradient filtering can be defined as:

#### STATISTICAL GRADIENT FILTERING

$$\mathbf{P}_{\text{stat}} = \frac{\beta}{1 + \beta} \mathbf{U}_k \mathbf{U}_k^T + \frac{1}{1 + \beta} \mathbf{I}. \quad (13)$$

The filtered gradient descent update becomes:

$$\mathbf{x} \leftarrow \mathbf{x} - \eta \left( \frac{\beta}{1 + \beta} \mathbf{U}_k \mathbf{U}_k^T + \frac{1}{1 + \beta} \mathbf{I} \right)^p \frac{\partial \Phi}{\partial \mathbf{x}}, \quad (14)$$

where  $p$  determines the degree of gradient filtering.

As shown in the supplementary document,  $p$  does not constitute an additional hyper-parameter; adjusting  $p$  is equivalent to rescaling  $\beta$ .

This form of statistical gradient filtering reveals an intuitive interpolation: when  $\beta \rightarrow \infty$ , the gradient filter reduces to  $\mathbf{U}_k \mathbf{U}_k^T$ , projecting gradients entirely onto the statistical subspace. When  $\beta \rightarrow 0$ , it reduces to  $\mathbf{I}$ , recovering standard gradient descent. The parameter  $\beta$  thus controls the balance between statistical guidance and flexibility.

#### 4.2 Eigenmode Rebalancing

The gradient filtering derived above treats all principal components equally within the retained subspace  $\mathbf{U}_k$ . However, as observed in Eq. (7), applying different weights to eigenmodes can selectively emphasize specific deformation modes. In PCA, eigenmodes with larger variance  $\sigma_i^2$  generally capture more dominant deformation patterns and should be prioritized. We modify the gradient filter in Eq. (13) to incorporate variance-based weighting:

#### EIGENMODE REBALANCING

$$\mathbf{P}_{\text{balanced}} = \frac{\beta}{1 + \beta} \mathbf{U}_k \Sigma_\epsilon^2 \mathbf{U}_k^T + \frac{1}{1 + \beta} \mathbf{I}, \quad (15)$$

where  $\Sigma_\epsilon^2 = \text{diag}(\sigma_1^2 + \epsilon, \dots, \sigma_k^2 + \epsilon)$  applies weights based on the variance of each eigenmode.

The hyperparameter  $\epsilon > 0$  ensures that low-variance modes are not excessively suppressed, maintaining expressiveness for subtle deformations. This rebalancing mirrors frequency-domain filtering of Eq. (7) but operates in the statistical shape domain: eigenmodes representing dominant variation directions receive stronger influence, while minor modes remain accessible. In practice,  $\epsilon$  can be set as a fixed constant; we use  $\epsilon = 0.3$  in all experiments.

#### 4.3 Statistical Reparameterization

Inspired by diffusion reparameterization [Nicolet et al. 2021], we recast the filtered update as the solution of an auxiliary optimization problem. Let  $\boldsymbol{\delta}$  denote the displacement from the mean shape, so that the deformed vertices are  $\mathbf{x} = \bar{\mathbf{x}} + \boldsymbol{\delta}$ . Then, consider the problem:

$$\arg \min_{\boldsymbol{\delta}} \frac{1}{2} \|\boldsymbol{\delta} - \mathbf{u}\|^2 + \frac{\beta}{2} \|(\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^T) \boldsymbol{\delta}\|^2, \quad (16)$$

where  $\mathbf{u}$  is an unconstrained auxiliary variable. The first term encourages proximity to  $\mathbf{u}$ , while the second penalizes deviations from the statistical subspace.

Solving the Euler-Lagrange equation for this problem yields the optimal displacement  $\boldsymbol{\delta}$ :

$$\boldsymbol{\delta} = \left( (1 + \beta)\mathbf{I} - \beta \mathbf{U}_k \mathbf{U}_k^T \right)^{-1} \mathbf{u} = \left( \frac{\beta}{1 + \beta} \mathbf{U}_k \mathbf{U}_k^T + \frac{1}{1 + \beta} \mathbf{I} \right) \mathbf{u} = \mathbf{P}_{\text{stat}} \mathbf{u}.$$

Replacing  $\mathbf{P}_{\text{stat}}$  with  $\mathbf{P}_{\text{balanced}}$  via the eigenmode rebalancing (Eq. (15)), we obtain the reparameterization  $\boldsymbol{\delta} = \mathbf{P}_{\text{balanced}} \mathbf{u}$ . The Jacobian of this transformation is:

$$\frac{\partial \boldsymbol{\delta}}{\partial \mathbf{u}} = \left( \frac{\beta}{1 + \beta} \mathbf{U}_k \Sigma_\epsilon^2 \mathbf{U}_k^T + \frac{1}{1 + \beta} \mathbf{I} \right) = \mathbf{P}_{\text{balanced}}. \quad (17)$$

To minimize the original objective  $\Phi(\mathbf{x})$  with respect to  $\boldsymbol{\delta}$ , we perform gradient descent on the unconstrained variable  $\mathbf{u}$  using the chain rule:

$$\mathbf{u} \leftarrow \mathbf{u} - \eta \frac{\partial \boldsymbol{\delta}}{\partial \mathbf{u}} \frac{\partial \Phi}{\partial \boldsymbol{\delta}}. \quad (18)$$

After updating  $\mathbf{u}$ , we recompute  $\boldsymbol{\delta}$  by substituting Eq. (18) into  $\boldsymbol{\delta} = \mathbf{P}_{\text{balanced}} \mathbf{u}$ . Then, we obtain:

$$\boldsymbol{\delta} \leftarrow \mathbf{P}_{\text{balanced}} \left( \mathbf{u} - \eta \frac{\partial \boldsymbol{\delta}}{\partial \mathbf{u}} \frac{\partial \Phi}{\partial \boldsymbol{\delta}} \right) = \boldsymbol{\delta} - \eta \mathbf{P}_{\text{balanced}}^2 \frac{\partial \Phi}{\partial \boldsymbol{\delta}}. \quad (19)$$




Fig. 3. Non-rigid template mesh registration on partial scan. Compared to parametric methods constrained by the PCA subspace, regularization approaches that suffer from instability, and Laplacian-based gradient filtering [Jung et al. 2023] that fails to faithfully handle occlusions, our method achieves accurate reconstruction on both visible and occluded regions. target: ©3DScanStore

This reveals that the reparameterization implicitly applies the squared statistical filter  $\mathbf{P}_{\text{balanced}}^2$  to the gradient, naturally incorporating statistical guidance into each optimization step. This reparameterization through  $\mathbf{u}$  enables straightforward integration with automatic differentiation frameworks:

```

1 def optimize_step():
2     # Apply statistical reparameterization
3     x_deformed = x_mean + P_balanced @ u
4     loss = objective(x_deformed)
5     loss.backward()

```

The forward pass multiplication  $\mathbf{P}_{\text{balanced}}\mathbf{u}$  enables automatic differentiation to implicitly apply the squared gradient filter during backpropagation.

Our statistical gradient filtering provides a principled way to incorporate shape priors into gradient-based mesh optimization, extending the framework of Nicolet et al. [2021] to the statistical shape domain. In addition, the reparameterization enables seamless integration of statistical gradient filtering with automatic differentiation, requiring only a single matrix multiplication in the forward pass.

## 5 Experiments

We validate our method primarily on 3D face reconstruction, where PCA-based statistical shape models are well-established and publicly available. We adopt the FLAME face model [Li et al. 2017] for our main experiments, utilizing its 300 shape and 100 expression principal components as the basis for statistical gradient filtering. Importantly, our framework is compatible with any PCA-based model; it can provide a domain-specific gradient filtering scheme for an optimization problem, as long as PCA can be defined. We

further demonstrate applications to body models and texture maps in the supplementary document.

We primarily compare against three approaches: (1) PCA parametric, which directly optimizes in the PCA coefficient space; (2) PCA regularization in Sec. 3.2, which optimizes vertex positions with an additional term penalizing deviation from the PCA subspace; and (3) Laplacian-based gradient filtering [Jung et al. 2023; Nicolet et al. 2021], which filters gradients through the Laplacian operator. Note that all methods, including our method, can benefit from additional supervision or regularization, e.g., landmark loss and Laplacian smoothing. To purely evaluate the effect of each method, we exclude such terms in all experiments. We evaluate these methods on non-rigid template mesh registration and inverse rendering from a single-view surface normal map, with additional results on different tasks in the supplementary document.

### 5.1 Non-rigid Template Mesh Registration on Partial Scan

We formulate non-rigid registration as minimizing the distance between a deformable source mesh  $\mathcal{M}_s = (\mathcal{V}_s, \mathcal{F})$  with vertices  $\mathcal{V}_s \in \mathbb{R}^{N \times 3}$  and a target point cloud  $\mathbf{P}_t \in \mathbb{R}^{M \times 3}$  with normals  $\mathbf{N}_t \in \mathbb{R}^{M \times 3}$ . After initial rigid alignment, we establish correspondences via cycle-consistent nearest neighbor matching: a pair  $(\mathbf{p}_i, \mathbf{s}_i)$  is included in the correspondence set  $\mathcal{C}$  only if the matching is bidirectionally consistent. The objective combines position and normal alignments [Jung et al. 2023]:

$$\Phi(\mathcal{V}_s) = \lambda_{\text{dist}} \frac{1}{|\mathcal{C}|} \sum_{i \in \mathcal{C}} \|\mathbf{s}_i - \mathbf{p}_i\|^2 + \frac{1}{|\mathcal{C}|} \sum_{i \in \mathcal{C}} (1 - \cos(\mathbf{n}_i^s, \mathbf{n}_i^p)), \quad (20)$$

where  $\lambda_{\text{dist}} = 10^5$ , and  $\mathbf{n}_i^s$  and  $\mathbf{n}_i^p$  are the computed and given normals at the matched points  $\mathbf{s}_i$  and  $\mathbf{p}_i$ , respectively.

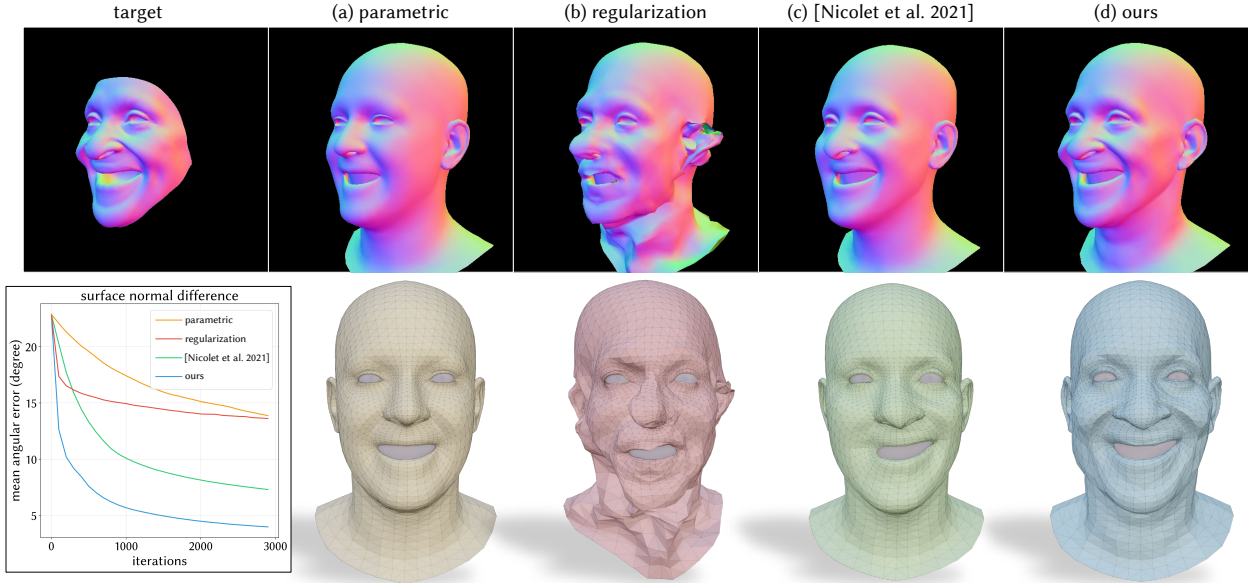


Fig. 4. Inverse rendering from a single-view surface normal map. Parametric methods cannot reconstruct the target shape, regularization exhibits instability, and Laplacian-based gradient filtering [Nicolet et al. 2021] lacks shape prior knowledge on occluded regions. Our method addresses all three limitations, accurately reconstructing the target geometry using learned shape correlations. target: ©3DScanStore

Table 1. Registration quality evaluated via chamfer distance to the ground truth mesh. For five camera viewpoints centered on the face, our method shows consistent accuracy across varying occlusion patterns.

Method	Chamfer Distance (mm)					Avg.
	Front	Top	Bottom	Left	Right	
PCA parametric	1.104	1.123	1.115	1.112	1.113	1.113
PCA regularization	1.216	1.312	1.316	1.306	1.331	1.296
Jung et al. [2023]	0.957	1.078	1.058	1.074	1.072	1.047
Statistical (ours)	<b>0.832</b>	<b>0.875</b>	<b>0.868</b>	<b>0.870</b>	<b>0.873</b>	<b>0.851</b>

We evaluate our method on 20 3D full head meshes from 3DScanStore [3DScanStore 2025], each with 29 facial expressions. To enable quantitative evaluation of occluded region reconstruction, we simulate occlusions on these meshes through ray casting from 5 virtual camera viewpoints (front, top, bottom, left, right). This yields 2,900 test cases, where the complete scan serves as ground truth for measuring chamfer distance including occluded regions. As shown in Table 1, our method consistently achieves the most accurate reconstructions across all viewpoints. By leveraging learned shape correlations, statistical gradient filtering accurately reconstructs both visible and occluded regions (Fig. 3).

## 5.2 Inverse Rendering from One-Shot Surface Normal Map

We formulate inverse rendering as recovering a 3D mesh  $\mathcal{M} = (\mathcal{V}, \mathcal{F})$  from a ground-truth normal map  $\mathbf{N}_{\text{gt}} \in \mathbb{R}^{H \times W \times 3}$  and silhouette mask  $\mathbf{S}_{\text{gt}} \in \{0, 1\}^{H \times W}$  through differentiable rendering [Laine et al. 2020]. The objective combines normal alignment [Giebenhain et al. 2025] and silhouette matching [Nicolet et al. 2021]:

$$\mathcal{L}(\mathcal{V}) = \frac{1}{|\mathcal{P}_\tau|} \sum_{i \in \mathcal{P}_\tau} \|\mathbf{n}_i - \hat{\mathbf{n}}_i\|_2^2 + \lambda_{\text{sil}} \|\mathbf{S}_{\text{render}} - \mathbf{S}_{\text{gt}}\|_2^2, \quad (21)$$

Table 2. Reconstruction accuracy on inverse rendering from a single-view normal map. Our method demonstrates the benefit of statistical guidance for recovering occluded surface normals.

Method	Mean Angular Error (degrees)	Chamfer Distance (mm)
PCA parametric	11.347	2.109
PCA regularization	9.398	2.663
Nicolet et al. [2021]	6.438	1.924
Statistical gradient filtering (Ours)	<b>3.477</b>	<b>1.404</b>

where  $\mathcal{P}_\tau$  denotes pixels with normal error below threshold  $\tau = 0.3$  to prevent outliers from dominating gradients,  $\mathbf{S}_{\text{render}}$  is a rendered binary mask on skin regions, and  $\lambda_{\text{sil}} = 5$ .

We evaluate inverse rendering using synthetic normal maps rendered from the same face meshes as in Sec. 5.1. Unlike the registration task where correspondences are relatively well-established through nearest neighbor search in 3D space, this task relies on pixel-to-pixel correspondences between rendered and target normal maps, which often produces false matches when facial features are misaligned. Under this challenging setting, as shown in Fig. 4 and Table 2, our method achieves the most accurate reconstruction in terms of both angular error and chamfer distance by filtering out spurious gradients from such false matches. Furthermore, unlike Nicolet et al. [2021] which enforces smoothness and limits detailed deformation, our method enables finer reconstruction in visible regions, as shown in Fig. 6.

We further evaluate our method on in-the-wild facial images, using Pixel3DMM [Giebenhain et al. 2025] to extract target normal maps from photographs. This scenario is more challenging than synthetic data due to noise and estimation artifacts in the predicted normals. We compare against both neural network methods that directly

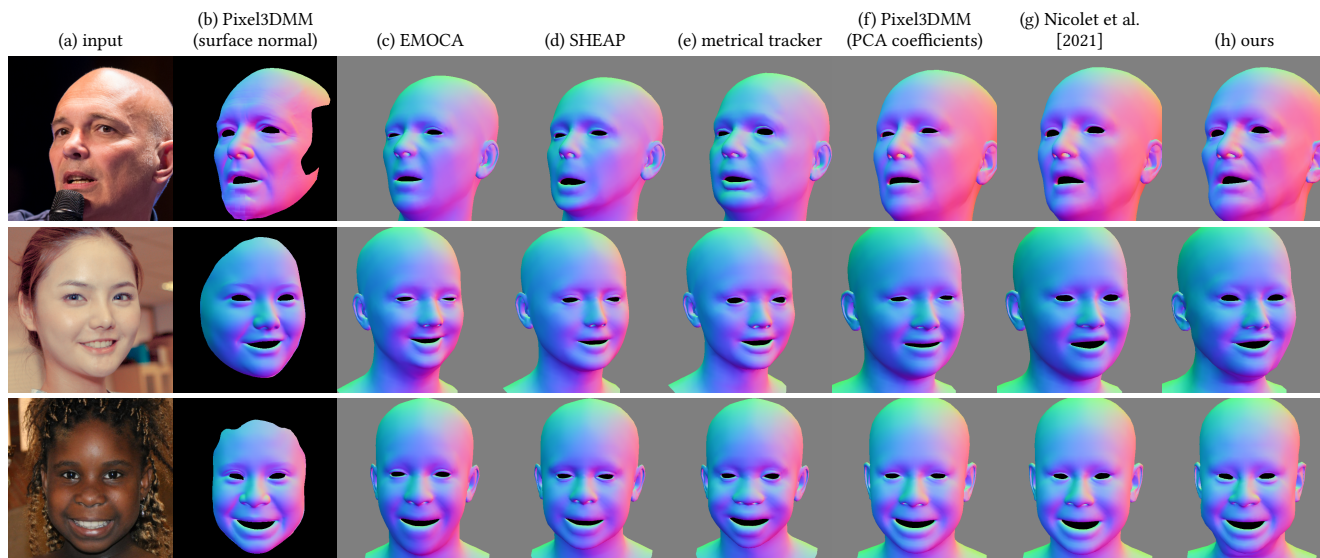


Fig. 5. 3D face reconstruction on in-the-wild images. (a) Input image and (b) its estimated surface normal from Pixel3DMM. *Neural network methods* (c)-(d) directly regress PCA coefficients from input images, limiting their ability to capture subject-specific geometry. *Optimization methods* (e)-(h): (e) and (f) fit PCA coefficients to the input image and estimated surface normals, respectively, constrained to the learned subspace; (g) and (h) directly optimize vertex positions using gradient filtering to enable better shape reconstruction without constraints on the solution space. While (g) and (h) look similar, Nicolet et al. [2021] promotes spatial smoothness and ours emphasizes a shape prior. Further comparisons between (g) and (h) on more challenging cases are provided in Fig. 6. Input photos are from the FFHQ dataset [Karras et al. 2019].

regress PCA coefficients from input facial images [Daněček et al. 2022; Schoneveld et al. 2025], and optimization methods that fit PCA coefficients to either the input image via photometric loss [Zielonka et al. 2022] or the estimated surface normal [Giebenhain et al. 2025]. As shown in Figs. 5 and 11, neural network methods are limited in capturing subject-specific geometry, while PCA-based optimization methods remain constrained to the learned subspace. Our method optimizes vertex positions with statistical gradient filtering, achieving more accurate reconstructions by leveraging the full solution space with statistical guidance.

### 5.3 Analysis

**5.3.1 Effect of eigenmode rebalancing.** Eigenmode rebalancing emphasizes high-variance modes capturing dominant, meaningful shape variations, while suppressing low-variance modes associated with minor or less reliable variations that are susceptible to noise. As shown in Fig. 7, without rebalancing, the geometry exhibits artifacts from noise propagated through low-variance modes, while applying eigenmode rebalancing produces stable and natural reconstruction.

**5.3.2 Robustness to parameter choice.** Our method introduces a single hyperparameter  $\beta$  in Eq. 15 that controls the balance between statistical structure and gradient information. As shown in Fig. 8, results become more stable as  $\beta$  increases, and remain consistent across a wide range of values once  $\beta$  is sufficiently large. This robustness to parameter choice makes our method practical, allowing a user to set  $\beta$  without extensive tuning.

**5.3.3 Mesh connectivity independence.** Our statistical gradient filtering naturally handles meshes with topologically disconnected

components, such as eyeballs separated from the facial skin surface. Laplacian-based method [Nicolet et al. 2021] relies on mesh connectivity to propagate gradients, so disconnected components receive no updates when only the skin surface is optimized (Fig. 9). In contrast, our method propagates updates through learned correlations, automatically adjusting eyeball positions when skin vertices deform, even without direct gradients computed on the eyeball surface.

**5.3.4 Limitations.** The effectiveness of statistical gradient filtering depends on the underlying shape prior. PCA-based models capture dominant shape variations but lack fine-grained details such as wrinkles, limiting our method’s guidance for deformations at fine scales (Fig. 10).

## 6 Conclusion

We introduced *statistical gradient filtering*, a method that guides geometry optimization via statistical structure learned from a shape collection. Unlike parametric methods that constrain solution spaces or regularization approaches that modify objectives, our method guides optimization paths while discovering novel shapes and preserving the original objective. Through reparameterization, we achieve seamless integration with automatic differentiation frameworks, enabling statistical guidance readily deployable in existing frameworks. Experiments on geometry deformation toward a target shape such as non-rigid registration and inverse rendering demonstrate that our method consistently improves robustness and solution quality compared to existing approaches. This work introduces a novel paradigm for geometry deformation, incorporating learned priors into the optimization trajectory rather than the solution space or objective.

## Acknowledgments

We thank anonymous reviewers for their valuable feedback. This work was supported by NRF grants (RS-2026-25485382, RS-2025-02216257) and IITP grant (RS-2022-II220290) funded by the Korean government (MSIT).

## References

- 3DScanStore. 2025. *3DScanStore*. <https://www.3dscanstore.com>
- Hyeonjang An, Wonjun Lee, and Bochang Moon. 2023. Adaptively weighted discrete Laplacian for inverse rendering. *The Visual Computer* 39, 8 (2023).
- Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. 2005. SCAPE: shape completion and animation of people. In *ACM Trans. Graph.*
- Douglas Arnold, Richard Falk, and Ragnar Winther. 1997. Preconditioning in H(div) and applications. *Math. Comp.* 66, 219 (1997).
- P Astrid, GV Papaioannou, JC Vink, and JD Jansen. 2011. Pressure preconditioning using proper orthogonal decomposition. In *SPE Reservoir Simulation Conference*.
- Volker Blanz and Thomas Vetter. 1999. A morphable model for the synthesis of 3D faces. In *Proc. ACM SIGGRAPH*.
- James Booth, Anastasios Roussos, Stefanos Zafeiriou, Allan Ponniah, and David Dunaway. 2016. A 3d morphable model learnt from 10,000 faces. In *Proc. CVPR*.
- Chen Cao, Yanlin Weng, Shun Zhou, Yiyang Tong, and Kun Zhou. 2013. Facewarehouse: A 3D facial expression database for visual computing. *IEEE TVCG* 20, 3 (2013).
- Jiong Chen, Florian Schäfer, Jin Huang, and Mathieu Desbrun. 2021. Multiscale Cholesky preconditioning for ill-conditioned problems. *ACM Trans. Graph.* 40, 4 (2021).
- Sebastian Claiici, Mikhail Bessmeltsev, Scott Schaefer, and Justin Solomon. 2017. Isometry-aware preconditioning for mesh parameterization. In *Comput. Graph. Forum*, Vol. 36. Wiley Online Library.
- Timothy F Cootes, Gareth J Edwards, and Christopher J Taylor. 1998. Active appearance models. In *Proc. ECCV*.
- Timothy F Cootes, Christopher J Taylor, David H Cooper, and Jim Graham. 1995. Active shape models—their training and application. *Computer Vision and Image Understanding* 61, 1 (1995).
- GB Diaz Cortes, Cornelis Vuik, and Jan Dirk Jansen. 2018. On POD-based deflation vectors for DPCG applied to porous media problems. *J. Comput. Appl. Math.* 330 (2018).
- Radek Daněček, Michael J Black, and Timo Bolkart. 2022. Emoca: Emotion driven monocular face capture and animation. In *Proc. CVPR*.
- Mathieu Desbrun, Mark Meyer, Peter Schröder, and Alan H Barr. 1999. Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proc. ACM SIGGRAPH*.
- Gerhard Dziuk. 1988. *Finite elements for the Beltrami operator on arbitrary surfaces*. Springer.
- Thomas Gerig, Andreas Morel-Forster, Clemens Blumer, Bernhard Egger, Marcel Luthi, Sandro Schönborn, and Thomas Vetter. 2018. Morphable face models—an open framework. In *IEEE International Conference on Automatic Face & Gesture Recognition*.
- Simon Giebenhain, Tobias Kirschstein, Martin Rünz, Lourdes Agapito, and Matthias Nießner. 2025. Pixel3DMM: Versatile Screen-Space Priors for Single-Image 3D Face Reconstruction. *arXiv preprint arXiv:2505.00615* (2025).
- Patrick Huber, Guosheng Hu, Rafael Tena, Pouria Mortazavian, Willem P Koppen, William Christmas, Matthias Rätzsch, and Josef Kittler. 2016. A multiresolution 3D morphable face model and fitting framework. In *Proc. Joint conference on computer vision, imaging and computer graphics theory and applications*.
- Wonjong Jang, Yuchoel Jung, Gyeongmin Lee, and Seungyong Lee. 2025. Instant Self-Intersection Repair for 3D Meshes. *ACM Trans. Graph.* 44, 4 (2025).
- Ian T Jolliffe. 1986. Principal components in regression analysis. In *Principal component analysis*. Springer.
- Yuchoel Jung, Hyomin Kim, Gyeongha Hwang, Seung-Hwan Baek, and Seungyong Lee. 2023. Mesh density adaptation for template-based shape reconstruction. In *Proc. ACM SIGGRAPH*.
- Yuchoel Jung, Hyomin Kim, Hyejeong Yoon, and Seungyong Lee. 2025. Preconditioned Single-step Transforms for Non-rigid ICP. In *Comput. Graph. Forum*, Vol. 44. Wiley Online Library.
- Julian Kaltheuner, Alexander Oebel, Hannah Droege, Patrick Stotko, and Reinhard Klein. 2025. Preconditioned Deformation Grids. In *Comput. Graph. Forum*, Vol. 44. Wiley Online Library.
- J Karatson and L Loczi. 2005. Sobolev gradient preconditioning for the electrostatic potential equation. *Computers & Mathematics with Applications* 50, 7 (2005).
- Tero Karras, Samuli Laine, and Timo Aila. 2019. A style-based generator architecture for generative adversarial networks. In *Proc. CVPR*.
- Shahar Z Kovalsky, Meirav Galun, and Yaron Lipman. 2016. Accelerated quadratic proxy for geometric optimization. *ACM Trans. Graph.* 35, 4 (2016).
- Samuli Laine, Janne Hellsten, Tero Karras, Yeongho Seol, Jaakko Lehtinen, and Timo Aila. 2020. Modular Primitives for High-Performance Differentiable Rendering. *ACM Trans. Graph.* 39, 6 (2020).
- Tianye Li, Timo Bolkart, Michael J Black, Hao Li, and Javier Romero. 2017. Learning a model of facial shape and expression from 4D scans. *ACM Trans. Graph.* 36, 6 (2017), 194–1.
- Selena Ling, Merlin Nimier-David, Alec Jacobson, and Nicholas Sharp. 2025. Stochastic Preconditioning for Neural Field Optimization. *arXiv preprint arXiv:2505.20473* (2025).
- Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. 2015. SMPL: a skinned multi-person linear model. *ACM Trans. Graph.* 34, 6 (2015).
- Li Luo and Xiao-Chuan Cai. 2023. Preconditioned Inexact Newton with Learning Capability for Nonlinear System of Equations. *SIAM Journal on Scientific Computing* 45, 2 (2023).
- Tobias Martin, Pushkar Joshi, Miklós Bergou, and Nathan Carr. 2013. Efficient Non-linear Optimization via Multi-scale Gradient Filtering. In *Comput. Graph. Forum*, Vol. 32. Wiley Online Library.
- Baback Moghaddam and Alex Pentland. 1997. Probabilistic visual learning for object representation. *IEEE TPAMI* (1997).
- Reinhard Nabben and Cornelis Vuik. 2006. A comparison of deflation and the balancing preconditioner. *SIAM Journal on Scientific Computing* 27, 5 (2006).
- JW Neuberger. 1985. Steepest descent and differential equations. *Journal of the Mathematical Society of Japan* 37, 2 (1985).
- Baptiste Nicolet, Alec Jacobson, and Wenzel Jakob. 2021. Large steps in inverse rendering of geometry. *ACM Trans. Graph.* 40, 6 (2021).
- Stanley Osher, Bao Wang, Penghang Yin, Xiyang Luo, Farzin Barekat, Minh Pham, and Alex Lin. 2022. Laplacian smoothing gradient descent. *Research in the Mathematical Sciences* 9, 3 (2022).
- Jea-Hyun Park, Abner J Salgado, and Steven M Wise. 2021. Preconditioned accelerated gradient descent methods for locally Lipschitz smooth objectives with applications to the solution of nonlinear PDEs. *Journal of Scientific Computing* 89, 1 (2021).
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. PyTorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32, Article 721 (2019).
- Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. 2019. Expressive body capture: 3d hands, face, and body from a single image. In *Proc. CVPR*.
- Pascal Paysan, Reinhard Knef, Brian Amberg, Sami Romdhani, and Thomas Vetter. 2009. A 3D face model for pose and illumination invariant face recognition. In *IEEE International Conference on Advanced Video and Signal based Surveillance*.
- Alex Pentland, Baback Moghaddam, and Thad Starner. 1994. View-based and modular eigenspaces for face recognition. In *Proc. CVPR*.
- Ulrich Pinkall and Konrad Polthier. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental mathematics* 2, 1 (1993).
- Boris T Polyak. 2007. Newton's method and its use in optimization. *European Journal of Operational Research* 181, 3 (2007).
- Josua Sassen, Henrik Schumacher, Martin Rumpf, and Keenan Crane. 2024. Repulsive shells. *ACM Trans. Graph.* 43, 4 (2024).
- Liam Schoneveld, Zhe Chen, Davide Davoli, Jiapeng Tang, Saimon Terazawa, Ko Nishino, and Matthias Nießner. 2025. SHeP: Self-Supervised Head Geometry Predictor Learned via 2D Gaussians. *arXiv preprint arXiv:2504.12292* (2025).
- Xing Shen, Runyuan Cai, Mengxiao Bi, and Tangjie Lv. 2024. Preconditioned nonlinear conjugate gradient method for real-time interior-point hyperelasticity. In *Proc. ACM SIGGRAPH*.
- Miroslava Slavcheva, Maximilian Baust, and Slobodan Ilic. 2018. Sobolevfusion: 3d reconstruction of scenes undergoing free non-rigid motion. In *Proc. CVPR*.
- Gabriel Taubin. 1995. A signal processing approach to fair surface design. In *Proc. ACM SIGGRAPH*.
- Chris Yu, Caleb Brakensiek, Henrik Schumacher, and Keenan Crane. 2021. Repulsive surfaces. *arXiv preprint arXiv:2107.01664* (2021).
- Yi Zhang and Jeff G Schneider. 2010. Projection penalties: dimension reduction without loss. In *Proc. ICML*.
- Yufeng Zhu, Rook Bridson, and Danny M Kaufman. 2018. Blended cured quasi-newton for distortion optimization. *ACM Trans. Graph.* 37, 4 (2018).
- Wojciech Zielonka, Timo Bolkart, and Justus Thies. 2022. Towards Metrical Reconstruction of Human Faces. In *Proc. ECCV*.




Fig. 6. Additional results on inverse rendering from a *single-view* surface normal map. Our method enables more detailed deformations that accurately capture the target geometry. target: ©3DScanStore

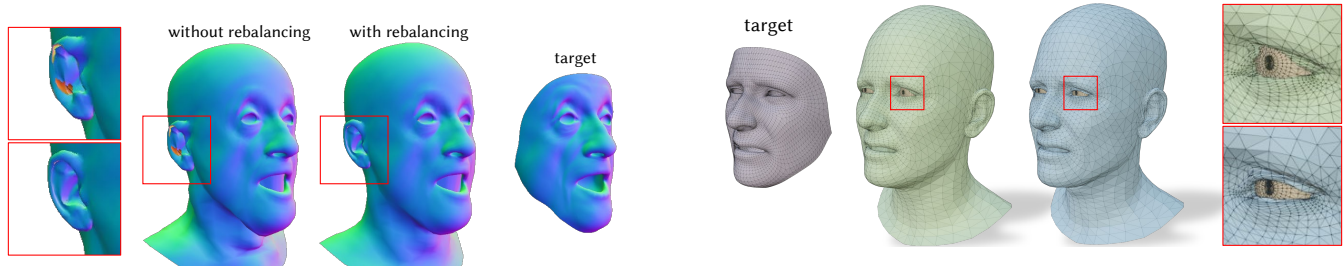


Fig. 7. Effect of eigenmode rebalancing. Without rebalancing, unobserved regions exhibit spurious deformations rarely seen in the training data, driven by low-variance eigenmodes. Rebalancing mitigates this artifact by emphasizing dominant eigenmodes. target: ©3DScanStore

Fig. 9. Automatic eyeball adjustment. When optimization is only applied to the skin region, Laplacian-based methods (green) cannot propagate updates to disconnected components such as the eyeballs, resulting in collisions. In contrast, our method (blue) automatically adjusts the eyeballs as well through learned statistical correlations and avoids such artifacts. target: ©3DScanStore




Fig. 8. Effect of hyperparameter  $\beta$ . Once  $\beta$  is set sufficiently large, our method consistently produces stable and natural results. target: ©3DScanStore




Fig. 10. Limitations. Our method cannot reliably guide deformations toward high-frequency details, such as wrinkles, which are rarely captured by deformation modes of a PCA model. Each pair shows the target surface normal map (left) and our result (right).



Fig. 11. Additional 3D face reconstruction results on in-the-wild images. Input photos are from the FFHQ dataset [Karras et al. 2019]