

# Face Reflectance and Geometry Modeling via Differentiable Ray Tracing

Abdallah Dib<sup>1</sup>, Gaurav Bharaj<sup>2</sup>, Junghyun Ahn<sup>2</sup>,  
Cédric Thébaud<sup>1</sup>, Philippe-Henri Gosselin<sup>1</sup>, Louis Chevallier<sup>1</sup>

<sup>1</sup> InterDigital, Rennes, France

<sup>2</sup> Technicolor, New York, NY

**Introduction.** A classic problem in computer vision is to reconstruct 3D faces from monocular image and video data. This needs disentangling scene and face parameters such as light, facial geometry (pose, identity, and expression), and reflectance, while accounting for camera depth ambiguity. Image pixel saturation due to self-shadows and specularities further add to the problem complexity.

To model the problem, optimization based [2] and more recently deep learning based [6] methods have been proposed. A photo-consistency loss difference between a parametric 3D face render and input facial image helps estimate the scene and face parameters. Often, the optimization cost is regularized with statistical priors and mathematical models of the parameters. [7] train a deep neural net to reconstruct faces and rely on high-quality training data obtained via a photometric camera rig. Both [7] and [3] present reflectance models with diffuse and specular albedo. They do not handle self-shadows explicitly.

We present a novel method to model 3D faces from monocular images, with parameterized ray traced scene light(s), and explicit disentangling of facial geometry (pose, identity and expression), reflectance (diffuse and specular albedo), and self-shadows (see Figure 2). We take inspiration from [1] and model the scene light as a *virtual light stage* with pre-oriented area lights. This setup is used in conjunction with a differentiable Monte-Carlo ray tracer [4], to optimize the scene and face parameters. Due to the disentangling of the parameters, we can not only obtain robust results, but also gain explicit control over them, with several practical applications. For example, we can change facial expression with accurate resultant self-shadows or relight the scene and obtain accurate specular reflection.




Figure 1: A *virtual light stage*

lights and resultant self-shadows.

In order to render physically accurate images via ray tracing, scene geometry and light must be known. To this end, we create a parameterized virtual light stage with  $N$  fixed oriented (pointing towards origin) area lights placed on vertex positions of an isotropically sampled hemisphere (see Figure 1). Parameterized face geometry is placed at the light stage’s origin, facing the camera. We then employ [4]’s differentiable ray tracing method to render an image. Finally, during face reconstruction optimization, backpropagation automatically calculates the gradient for the various parameters w.r.t photo-consistency cost loss between the input and rendered images.

The scene light parameters are the *color* intensities  $\{\mathbf{l}_i\} \in N(=20)$  and  $\mathbf{l}_i \in \mathbb{R}^3$ . A *light-off* (Fig 1) implies intensity  $\mathbf{0}^T$ . The face geometry parameters are  $\alpha \in \mathbb{R}^{80}$  which control a 3D model morphable based facial identity,  $\beta \in \mathbb{R}^{90}$  controls facial blend shape expressions, and  $[\mathbf{R} \ \mathbf{t}] \in \mathbb{R}^6$  the geometry transform. The skin reflectance model is given by Cook-Torrance reflectance  $f_r = k_d f_{\text{lambertian}}(\mathbf{l}_j, \delta, \mathbf{n}_i) + k_s f_{\text{cooktorrance}}(\mathbf{l}_j, \mathbf{c}_{\text{spec}}, \mathbf{n}_i)$ . Here,  $\delta \in \mathbb{R}^{80}$  are the morphable model’s vertex color albedo,  $\mathbf{n}_i(\alpha, \beta, \mathbf{R}, \mathbf{t})$  the vertex normal and  $\mathbf{c}_{\text{spec}} \in \mathbb{R}^{32 \times 32}$  are the pixel values of a vectorized gray-scale specular texture map.

**Optimization.** The photo-consistency optimization is given by:

$$\operatorname{argmin} E_{\text{data}} + w_1 E_{\text{prior}} + w_2 E_{\text{light}}, \quad (1)$$

where  $E_{\text{data}} = \sum_{i \in \mathcal{I}} \|\mathbf{p}_i^T - \mathbf{p}_i^R\|_2^2$  and  $E_{\text{light}} = \sum_{j=0}^N \|\mathbf{l}_j - \mathbf{m}_j\|_2^2$ . The vectors  $\mathbf{p}_i^T, \mathbf{p}_i^R$  are ray traced and real image ( $\mathcal{I}$ ) pixel color values, respectively

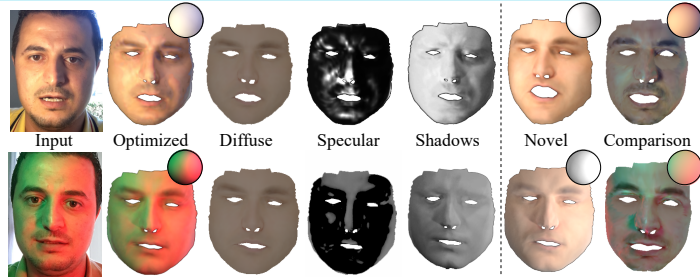


Figure 2: (From left to right) Given an input image, we optimize the scene light (shown via spherical environmental map), reflectance (diffuse and specular albedo) parameters with self-shadows. Top: We optimize for uniform color light(s). The 6th column shows updated shadows, as we transform the geometry and change the facial expression. Bottom: We optimize for multi color lights. The 6th last column shows a relit render with accurate reflectance and shadows, while the last column shows the reconstructed albedo and illumination via Garrido et al. [2]’s method.

and  $\mathbf{m}_j$  is the mean intensity of the  $j^{\text{th}}$  light. In order to make the optimization tractable, we use landmark and statistical prior  $E_{\text{prior}}$  defined in [2]. Our method relies on geometry, reflectance, and light parameters to calculate the final render pixel color values  $\mathbf{p}_i^T = \mathcal{F}(f_r, \mathbf{l}_j, \alpha, \beta, \mathbf{R}, \mathbf{t})$ , where  $\mathcal{F}$  is the Monte-Carlo estimator of the lighting equation. Generally, it is desirable to minimize light flux, to this end, we add a light flux conservation regularizer  $E_{\text{light}}$  to Equation 1.

We optimize the parameter in a sequential strategy as described in [2]’s Algorithm 1 - Multi-Step Optimization Strategy. In addition, we add the following steps to the optimization sequence: (1) Cook-Torrance reflectance based specularities, and (2) light ( $\mathbf{l}_j$ ) optimization. The ‘shadowed’ regions of the 3D face geometry are calculated by the ray tracer automatically.

**Results and Conclusion.** In Figure 2 (Top), we show face reconstruction with disentangled reflectance (diffuse & specular albedo), light (visualized as a spherical environment map), geometry (identity & expression), and self-shadow, especially near the nose region. While, Figure 2 (Bottom) shows accurate self-shadow based on optimized multi-color lights. Notice how the change in scene light (6th column) results in correct updated reflectance & self-shadows. Disentangled light, reflectance & shadows result in accurate reconstruction of face geometry regions that correspond to image patches with saturated pixel colors (due to shadows or specularities) and does not deform geometry unnaturally.

We present a novel strategy to automatically reconstruct faces with ray traced light, specularities, and shadows. The method robustly handles scenes with strong directional lights, partial occlusion due to shadows and specularities. By using a differentiable ray tracer, we accurately optimize the photo-consistency loss.

However, our method is limited by the ray tracer’s sampling rate. A low sampling rate results in a noisy render, while higher sampling rate results in slow optimization convergence. Another limitation is that our reflectance model does not model skin’s subsurface scattering properties.

- [1] P Debevec, T Hawkins, et al. *Acquiring the Reflectance Field of a Human Face*. Computer Graphics Interactive Techniques, 2000.
- [2] P Garrido, M Zollhöfer, D Casas, L Valgaerts, et al. *Reconstruction of Personalized 3D Face Rigs from Monocular Video*. ACM TOG, 2016.
- [3] P Gotardo, J Riviere, D Bradley, et al. *Practical Dynamic Facial Appearance Modeling and Acquisition*. ACM SIGGRAPH Asia, 2018.
- [4] TM Li, M Aittala, et al. *Differentiable Monte Carlo Ray Tracing through Edge Sampling*. ACM SIGGRAPH Asia, 2018.
- [5] M Pharr, W Jakob, and G Humphreys. *Physically based Rendering: From Theory to Implementation*. Morgan Kaufmann, 2016.
- [6] A Tewari, F Bernard, P Garrido, G Bharaj, et al. *FML: Face Model Learning from Videos*. CVPR, 2019.
- [7] S Yamaguchi, S Saito, et al. *High-fidelity Facial Reflectance and Geometry Inference from an Unconstrained Image*. ACM TOG, 2018.