Resolving Shape Ambiguities using Heat Conduction and Shading

Akihiko Oharazawa, Sriram Narayanan, Mani Ramanagopal, and Srinivasa G. Narasimhan

Abstract—Shape from shading using a single image of a Lambertian surface is inherently ambiguous. When the light source direction is known, the surface normal estimation has a cone-ambiguity, which worsens when the source is unknown. Recently, shape from heat conduction has emerged as an approach that leverages heat transport equations to estimate the Shape Laplacian operator, an intrinsic measure of shape. However, deriving surface normals from the Laplacian operator encounters a local binary convex/concave ambiguity. Our contribution introduces a novel theory to resolve these local shape ambiguities (excluding a few degeneracies) without relying on priors like smoothness, by combining the cues from shading and heat conduction. Our method ensures the mathematical constraints of both shading and the Laplacian are satisfied simultaneously, even with an unknown light source. We validate our theory through simulations of complex shapes and analyze its performance in the presence of noise.

Index Terms—Shape Reconstruction, Heat Conduction, Concave/convex Ambiguity, Thermal Video

1 Introduction

SHADING has long been a crucial cue for shape estimation in computer vision. Woodham [1] introduced photometric stereo, a technique where shape and albedo are derived from images of a Lambertian surface illuminated by at least three known distant light sources. When the directions of these sources are unknown, Belhumeur et al. [2] demonstrated that shape, albedo, and lighting can be estimated up to three degrees of freedom, known as the Generalized Bas-Relief (GBR) ambiguity. Subsequent works enhance the realism of the appearance model by modeling interreflections [3] and inverse-square fall-off [4] to resolve the GBR ambiguity.

Even with a known source direction, when only a single shading image of a Lambertian surface is available, the estimation of the surface normal is highly ambiguous — any normal along a cone around the source vector satisfies the shading/irradiance term ($n.s$). And if the source direction is unknown, the ambiguity is worse. Thus, historically, many approaches exploit additional shape (known boundary, smoothness, etc.) [5] and/or lighting priors [6] to overcome these ambiguities. Today, deep learning models recover depth maps from a single image based on supervised pre-training on large datasets of captured images and ground truth 3D models [7]. However, these models may struggle to generalize when the scene conditions differ significantly from the training data.

Recently, thermal heat conduction has been proposed as a strong cue of surface shape [8]. From a single thermal video of an object undergoing either heating or cooling, the authors estimate the Shape Laplacian-Beltrami operator. This operator is an intrinsic representation of shape (as opposed to extrinsic surface normals or depths), and the approach works under arbitrary and unknown lighting and even for non-Lambertian surfaces. But computing normals or depth from the Laplacian operator suffers from a local binary convex/concave ambiguity. The authors suggest capturing 2-4 thermal videos under different lighting directions to resolve this ambiguity [8].

• *The authors are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213.*
E-mail: {*aoharaza, snochurn, manikans, srinivas*}*@andrew.cmu.edu*
Webpage: http://imaging.cs.cmu.edu/vision-with-heat-and-light/

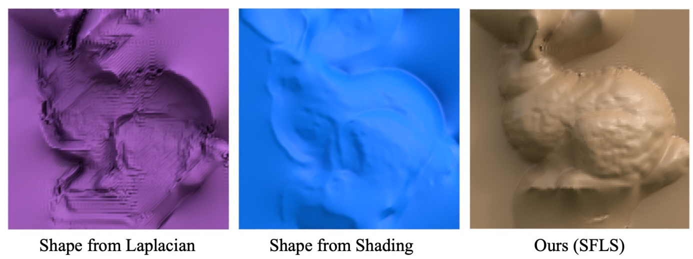

Fig. 1. Reconstructed shapes using estimation techniques based on Laplacian cues, shading cues, or both. While Laplacian cues constrain the local curvature around a pixel, it is invariant to convex/concave deformations. On the other hand, shading cues constrain the normal at each pixel to lie on a cone around the source direction, but does not constrain neighboring normals. This paper combines these complementary cues to resolve ambiguities while enabling faithful 3D reconstructions.

Our key contribution lies in developing a novel theory to resolve the ambiguities in shape recovery by combining two cues: heat conduction and shading, under an unknown but distant light source. Mathematically, we resolve the ambiguities in normal estimation by simultaneously satisfying the Shape Laplacian and the Cosine (or $n.s$) term in shading. As before, shading is obtained by capturing a visible spectrum image of the scene. Heat conduction is obtained using the thermal transient video of the scene that is illuminated (and hence heated) by the light source. For the purposes of our theory, we assume the surface albedos are either uniform or are estimated *a priori* using the visible-thermal intrinsic decomposition approach [9].

Following [8], we model the Shape Laplacian using cotangent weights on a mesh surface. Then, we develop a mathematical relationship between these weights and the local vertex surface normals. We consider the case of known source direction first. Our method estimates normals by satisfying both constraints — the equations for cotangent weights and the cosine term for irradiance — simultaneously, resolving ambiguities except in a few degenerate cases. Then, we consider the general case of an unknown source direction and develop an algorithm that uniquely estimates both the normals (see Fig. 1) and source direction.

Previous approaches use priors such as smoothness (including penalizing Laplacian or second derivatives of estimated normals, for instance) that often clash with the "data" terms and require setting a regularization hyperparameter that changes with iterations. But our shape Laplacian is actually a "data term" that provides the correct smoothness too without the need for any regularization parameters.

Our contribution is primarily theoretical and we demonstrate the theory using simulations. We conclude with a discussion of limitations and future work in making the approach practical and accurate for real-world scenes.

2 Related Work

Single-view shape recovery from shading has been studied for several decades and remains an active research topic in computer vision [10], [11], [12], [13]. Under the assumption of Lambertian reflectance, shape-from-shading is fundamentally ill-posed, especially when the light source is unknown. This is because the observed shading can be explained by a multidimensional space of possible surface normal [14]. Over the years, a variety of techniques have been proposed to address this ill-posedness, including partial differential equation (PDE) [5], optimization [15], and local analysis or linearization approaches [16]. However, these methods often suffer from practical issues such as convergence to local minima. As a result, many shape-from-shading methods either assume known illumination or introduce strong priors to constrain the solution space. For example, SIRFS [6] employs distinct lighting priors for natural versus laboratory environments and imposes surface normal priors around the occluding boundaries. Despite such advancements, fully resolving the ambiguity of shape-from-shading under the single-view and unknown-lighting scenario remains an open challenge.

Thermal imaging offers a promising avenue to alleviate some of these ambiguities. Because thermal cameras capture wavelengths outside the visible spectrum, they enable new applications that leverage material properties, temperature distributions, and radiative characteristics that are otherwise difficult to observe. Recent advances include pose estimation for humans or objects [17], navigation [18] and material classification [19]. Moreover, using the heat conduction equation, researchers have estimated diffusivity and emissivity for planar objects [20]. In [9], the authors introduce the physical relationship between a visible image and a thermal video of the scene and demonstrate that intrinsic image decomposition becomes a well-posed problem.

Several efforts have explored shape estimation via thermal imaging. Nagase et al. [21] proposed a method for 3D reconstruction using multispectral thermal radiation while accounting for air attenuation. Tanaka et al. [22] introduced a photometric stereo approach using steady-state (equilibrium) thermal images. Because the heat conduction equation strongly depends on surface curvature and boundary conditions, the transient response of thermal diffusion can, in principle, reveal shape information. Narayanan et al. [8] exploited the Laplacian operator from the heat equation to reconstruct surface shape, thereby highlighting the potential of time-resolved thermal observations. However, their approach still requires two or more light sources to resolve local concavity or convexity ambiguities. In contrast, our method assumes a single, uncalibrated light source and highlights the relationship between shading and Laplacian cues derived from heat conduction. We provide a framework for the simultaneous estimation of unknown light source direction and shape, and theoretically demonstrate that the mathematical constraints of both shading and Laplacian cues resolved local ambiguities.

3 Problem Statement

The objective of this work is to estimate the scene geometry (shape) and the light source direction ($\hat{\mathbf{s}}$) given only a single shading image (η) and the shape Laplacian (L). Similar to the original Shape-from-Shading (SfS) problem definition, we consider an orthographic camera and represent the shape using surface normals (p_i, q_i) at each pixel i. Unlike SfS, we do not know the light source direction and do not use an arbitrary smoothness term.

Let Ω denote the unknown surface in 3D. We discretize it into a mesh with V vertices, with each vertex lying along the camera ray from a corresponding pixel. We use a regular grid connectivity such that each vertex is connected to its six neighbors, as depicted in Fig. 2.

The shading η at a pixel i can be written as

$$\eta(i) = \frac{\vec{\mathbf{n}}(i)}{\|\vec{\mathbf{n}}(i)\|} \cdot \frac{\vec{\mathbf{s}}}{\|\vec{\mathbf{s}}\|} \quad \text{where} \quad \vec{\mathbf{n}}(i) = \begin{bmatrix} p_i \\ q_i \\ -1 \end{bmatrix}, \vec{\mathbf{s}} = \begin{bmatrix} p_s \\ q_s \\ -1 \end{bmatrix}. \tag{1}$$

For scenes with spatially varying albedo, one could use [9] to recover shading from the visible image and thermal video.

The discrete Laplacian on Ω, denoted by Δ_Ω, is represented by a sparse symmetric matrix $\mathrm{L} \in \mathbb{R}^{|V|\times|V|}$. Specifically, each non-zero entry represents a weight w_{ij} between vertices i and j described as [8]:

$$\mathrm{L}_{ij} = \begin{cases} w_{ij} = \dfrac{\cot \alpha_{ij} + \cot \beta_{ij}}{2} & \text{if } i \neq j \\ -\sum_{j \in \mathcal{N}(i)} w_{ij} & \text{if } i = j \end{cases} \tag{2}$$

where α_{ij} and β_{ij} are the angles of the two triangles adjacent to the edge (i, j). These edge weights can be computed directly from vertex normals, as we will show in Sec. 4.1.

4 Shape from Laplacian and Shading

We explore how the intrinsic geometry captured by Laplacian operators can be combined with photometric cues from shading to recover surface normals more accurately and unambiguously. While Laplacian-based methods offer geometric structure, they suffer from local ambiguities in normal direction. Conversely, shape-from-shading provides photometric constraints but has its own set of ambiguities tied to lighting. In this section, we introduce a novel formulation that connects Laplacian weights directly to surface normals, allowing shading information to be integrated

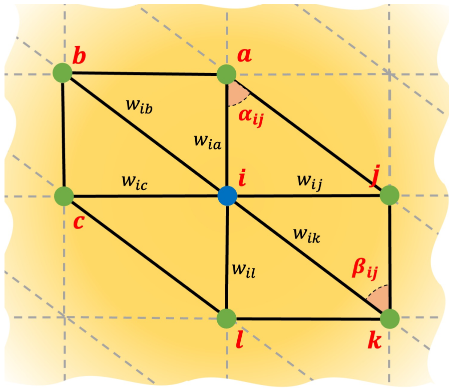

Fig. 2. A regular grid triangular mesh whose cotangent weights for any vertex i depend on its neighboring six vertices.

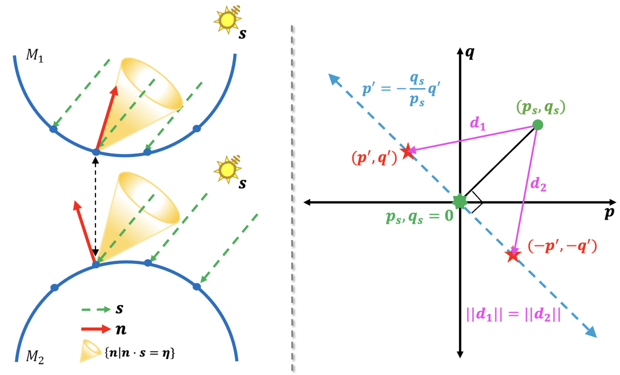

Fig. 3. **Left:** Illustrates ambiguity in shape-from-shading, where a cone of normals yields identical shading, and a local binary ambiguity in shape-from-Laplacian, where different shapes with the same cotangent weights appear as mirror reflections in normal space. The true object shape is uniquely determined at the intersection of the shading cone and Laplacian normals. **Right:** Demonstrates degenerate cases where the proposed shape-from-Laplacian-and-shading method fails.

seamlessly. We also analyze the nature of ambiguities in each modality and show how their intersection leads to a unique and consistent surface reconstruction.

4.1 Shape Laplacians using Vertex Normals

We introduce a novel parameterization of Laplacian weights in terms of vertex normals. The surface normal at a vertex can be expressed using depth gradients as $p = \frac{\partial z}{\partial x}$ and $q = \frac{\partial z}{\partial y}$. Discretizing these derivatives and using the notation illustrated in Fig. 2, we can write:

$$p_i = \frac{z_j - z_i}{\delta_p}, \quad q_i = \frac{z_l - z_i}{\delta_p} \tag{3}$$

where δ_p is the pixel pitch and z_i is the depth of vertex i.

The edge vectors connecting vertex a to its neighbors can then be written in terms of normals:

$$\vec{ai} = (0, \delta_p, z_i - z_a) \quad \vec{aj} = (\delta_p, \delta_p, z_j - z_a) \tag{4}$$

Substituting depth differences in terms of p and q:

$$\begin{aligned} z_i - z_a &= q_a \cdot \delta_p \\ z_j - z_a &= (p_i + q_a) \cdot \delta_p \end{aligned} \tag{5}$$

Using these, we compute the cotangent of the angle α_{ij} between edges $\vec{ai}$ and $\vec{aj}$:

$$\cot \alpha_{ij} = \frac{|\vec{ai} \cdot \vec{aj}|}{|\vec{ai} \times \vec{aj}|} = \frac{|1 + p_i q_a + q_a^2|}{\sqrt{p_i^2 + q_a^2 + 1}} \tag{6}$$

Similarly,

$$\cot \beta_{ij} = \frac{|\vec{ki} \cdot \vec{kj}|}{|\vec{ki} \times \vec{kj}|} = \frac{|1 + p_i q_j + q_j^2|}{\sqrt{p_i^2 + q_j^2 + 1}} \tag{7}$$

Therefore, one can write the weight w_{ij} in L from Eq. 2 as,

$$w_{ij} = \frac{1}{2}\left(\frac{|1 + p_i q_a + q_a^2|}{\sqrt{p_i^2 + q_a^2 + 1}} + \frac{|1 + p_i q_j + q_j^2|}{\sqrt{p_i^2 + q_j^2 + 1}}\right). \tag{8}$$

This establishes a direct relationship between the cotangent weights and the surface vertex normals. The corresponding equations for all six edge weights are provided in the Appendix.

4.2 Resolving Shape Ambiguities

In this section, we show that combining Laplacian-based and shading-based cues, under known distant illumination, resolves the shape ambiguity that arises when each is used independently. Specifically, we demonstrate that the intersection of their solution spaces leads to a unique normal direction at each point, except in a few degenerate cases. We begin by revisiting the types of ambiguities that arise when estimating shape from either Laplacian or shading cues alone, assuming the light source direction (p_s, q_s) and shading values η are known. While we make this assumption here for clarity, the following section describes how the source direction can also be robustly estimated through our optimization framework.

4.2.1 Ambiguities in Shape from Laplacian

As shown in [8], the Laplace operator is intrinsic to a surface and invariant to rigid transformations. This means there are infinitely many 3D embeddings that share the same Laplacian. However, when shapes are constrained to lie along the camera rays and to preserve local edge lengths, the ambiguity becomes binary. Note that flipping the sign of all vertex normals in a patch, i.e. substituting $(-p_x, -q_x)$ in place of (p_x, q_x), in Eq. 8 does not change L. This local binary ambiguity arises from isometric deformations and reflects a convex/concave uncertainty at each point.

4.2.2 Ambiguities in Shape from Shading

Shape-from-shading suffers from a different kind of ambiguity. Even when the albedo is constant and the light source direction (p_s, q_s) is known, the surface normal is not uniquely determined. Instead, all normals that make the same angle with the light source produce the same shading. Mathematically, from Eq. 1, this is written as

$$\frac{pp_s + qq_s + 1}{\sqrt{p^2 + q^2 + 1}\sqrt{p_s^2 + q_s^2 + 1}} = \eta \tag{9}$$

This defines a cone of possible normals around the light source direction, creating a continuous ambiguity distinct from the binary ambiguity of Laplacian-based methods.

4.2.3 Intersection of Shading and Laplacian Ambiguities

We seek a normal direction that satisfies both the binary Laplacian constraint and the shading equation as shown in Fig. 3 (left). If both possible Laplacian normals (p, q) and $(-p, -q)$ satisfy Equation 9, then:

$$pp_s + qq_s + 1 = -pp_s - qq_s + 1 \tag{10}$$

$$\Rightarrow \quad pp_s + qq_s = 0 \tag{11}$$

This yields the condition $pp_s = -qq_s$ corresponding to a narrow set of surface orientations in the normal space. This ambiguity persists only under special cases: (i) Fronto-parallel lighting: when $p_s = q_s = 0$, where both normal directions yield identical shading. (ii) Specific surface orientation: when $p = -\frac{q_s}{p_s}q$, representing a plane orthogonal to the bisector of (p_s, q_s) and $(-p_s, -q_s)$.

Geometrically, these cases occur when the light direction (p_s, q_s) lies equidistant between (p, q) and its mirror reflection $(-p, -q)$ in the pq space–i.e., along the line where shading distance to both normals is equal (see Fig. 3–right). In such configurations, both candidate normals fall within the same shading cone defined by Eq.9, making the ambiguity unresolvable.

However, these conditions occur only for some rare and isolated directions and typically do not persist across an entire surface. Thus, in general lighting conditions, combining Laplacian and shading cues yields a unique surface normal almost everywhere.

5 Jointly Estimating Lighting and Shape

In the previous section, we looked at how individual ambiguities in shape normals from Laplacian and shading can be solved when the source direction is given. Building on the theoretical framework, in this section, we describe our optimization process in detail that allows estimating the uncalibrated light source direction and surface normal across the entire image in a three step process. First, we identify a few locally near-planar patches for an initial estimate of the source direction. Second, we refine both the surface normals at these patches and the light source direction by minimizing photometric error under shading and Laplacian constraints, iterating until convergence. Finally, we propagate the refined surface normals across the entire image domain to ensure global consistency.

5.1 Initializing Normals for Near-Planar Patches

In our pixel-based mesh parameterization with connectivity as shown in Fig. 2 we can see that every vertex is connected to a maximum of six neighbors. Given this connectivity scheme, to estimate an initial source direction we first identify nearly flat regions based on the insight that Laplacian weights of diagonally opposite edges are approximately equal for planar patches. These near-planar regions are not known *apriori* but are automatically identified from the Laplacian. Specifically, for Fig. 2 if the shown vertices correspond to a locally flat surface then,

$$w_{ij} \approx w_{ic}, \quad w_{ik} \approx w_{ib}, \quad w_{ia} \approx w_{il}. \tag{12}$$

We choose pixels with the least sum of absolute differences between diagonally opposite edge weight at that pixel. For such a nearly planar pixel i, the Laplacian weights can be expressed in terms of the unknown patch normal (p, q) as:

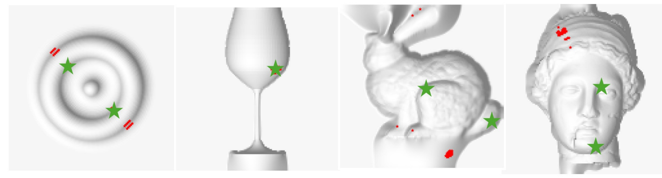

Fig. 4. The above visualization shows a subset of vertices in $\mathcal{F}$ marked with red dots and highlighted with green markers for visibility. These candidate patches were picked as near-planar regions based on the input Laplacian weights.

$$\begin{aligned} \frac{w_{ia} + w_{il}}{2} &\approx \frac{p(p+q)+1}{\sqrt{p^2+q^2+1}} = \mathrm{w}_1 \\ \frac{w_{ij} + w_{ic}}{2} &\approx \frac{q(p+q)+1}{\sqrt{p^2+q^2+1}} = \mathrm{w}_2 \\ \frac{w_{ik} + w_{ib}}{2} &\approx \frac{-pq}{\sqrt{p^2+q^2+1}} = \mathrm{w}_3 \end{aligned} \tag{13}$$

Using the symmetry in the above equations, note that

$$\frac{p^2+q^2+2}{\sqrt{p^2+q^2+1}} = w_1 + w_2 + 2w_3, \tag{14}$$

which reduces to

$$\frac{x^2+1}{x} = w_1 + w_2 + 2w_3, \tag{15}$$

where $x = \sqrt{p^2+q^2+1}$ and by definition, $x \geq 1$. We solve this quadratic equation in x and pick the solution that is greater than or equal to 1. In appendix, we show that the other solution is less than 1. To recover the values of p and q, observe that

$$\frac{p^2+1}{x} = w_1 + w_3. \tag{16}$$

From this equation, we can compute $|p| = \sqrt{(w_1+w_3)x - 1}$ and similarly, $|q| = \sqrt{(w_2+w_3)x - 1}$. Finally, to recover the sign of p, q, we use the sign of w_3 as follows. If Sign$(w_3) < 0$, then the two solutions are $(|p|, |q|)$ and $(-|p|, -|q|)$. If Sign$(w_3) > 0$, then the two solutions are $(|p|, -|q|)$ and $(-|p|, |q|)$. When $w_3 = 0$, the solutions simplify to either $p = 0$ and $q = \pm\sqrt{w_2^2 - 1}$ or $p = \pm\sqrt{w_1^2 - 1}$ and $q = 0$.

This formulation allows us to recover surface normals $\mathbf{n}_{\mathcal{F}} = \{(p_i, q_i, -1)\ \forall i \in \mathcal{F}\}$ directly from Laplacian weights in locally planar patches where $\mathcal{F}$ is the set of all vertices whose diagonal weights are approximately equal.

5.2 Estimating Source Given Normals

Once surface normals are known at a set of pixel locations, we can estimate the lighting direction using a reverse photometric stereo approach, derived from the shading equation in Eq. 9. In matrix form, this estimation can be expressed as:

$$\begin{bmatrix} \tilde{p}_0 & \tilde{q}_0 & \alpha_0 \\ \vdots & \vdots & \vdots \\ \tilde{p}_N & \tilde{p}_N & \alpha_N \end{bmatrix} \begin{bmatrix} \hat{p}_s \\ \hat{q}_s \\ \beta_s \end{bmatrix} = \begin{bmatrix} \eta_0 \\ \vdots \\ \eta_N \end{bmatrix} \tag{17}$$

Algorithm 1 Shape from Laplacian and Shading (SFLS)

procedure SFLS($\mathbf{L}, \boldsymbol{\eta}$)
    `// Normals for near-planar patches`
    $\mathcal{F}, \mathbf{n}_{\mathcal{F}} \leftarrow$ `planarPatchNormals`$(\mathbf{L})$
    `// Optimize` $\hat{\mathbf{s}}$ `and refine` $\mathbf{n}_{\mathcal{F}}$
    while $\|\Delta\hat{\mathbf{s}}\| > \epsilon_s$ or $\|\Delta\mathbf{n}_{\mathcal{F}}\| > \epsilon_n$ **do**
        $\hat{\mathbf{s}} \leftarrow$ `estimateSourceDir`$(\mathcal{F}, \mathbf{n}_{\mathcal{F}}, \boldsymbol{\eta}_{\mathcal{F}})$
        $\mathbf{n}_{\mathcal{F}} \leftarrow$ `shapeGivenSource`$(\mathcal{F}, \mathbf{n}_{\mathcal{F}}, \mathbf{L}, \hat{\mathbf{s}})$
    end while
    `// Propagate normals to full surface`
    $\mathcal{M} \leftarrow \mathcal{F}$
    while $\mathcal{M} \neq \mathcal{V}$ **do** ▷ $\mathcal{V}$: all mesh vertices
        $\mathcal{K} \leftarrow \mathcal{M} \cup \mathcal{N}(\partial\mathcal{M})$ ▷ Expand frontier
        $\mathbf{n}_{\mathcal{K}} \leftarrow$ `shapeGivenSource`$(\mathcal{K}, \mathbf{n}_{\mathcal{M}}, \mathbf{L}, \hat{\mathbf{s}})$
        $\mathcal{M} \leftarrow \mathcal{K}$
    end while
    return $\mathbf{n}_{\mathcal{V}}$
end procedure

Here, $(\hat{p}_s, \hat{q}_s, \beta_s)$ represents the estimated lighting direction, inferred from N observed pixels. The term β_s is a normalization factor. The entries $\tilde{p}$ and $\tilde{q}$ in the above matrix are normalized components of the surface normals, computed as $\tilde{p} = \alpha p$ and $\tilde{q} = \alpha q$, where $\alpha = 1/\sqrt{p^2 + q^2 + 1}$ ensures unit norm.

However, an inherent ambiguity arises because a surface normal $(\tilde{p}, \tilde{q})$ and its flipped version $(-\tilde{p}, -\tilde{q})$ yield identical shading under correspondingly flipped lighting direction $(-p_s, -q_s)$. At the same time, both versions of the surface normals satisfy the Laplacian weights. Here, the value of $(\tilde{p}, \tilde{q})$ corresponds to one of the two signed solutions from the previous subsection.

To disambiguate the correct sign for the surface normals, we apply a RANSAC [23] based approach. We randomly pick a light source direction and then determine the sign for each surface normal which produces the lower shading error. Using these signs, we compute the optimal light source direction. We repeat until both the light source direction and signs of the normals are converged. The signs of the final light source direction $[p_s, q_s, -1]^T$ and surface normals $[p_i, q_i, -1]^T$ corresponds to one of two global solutions, with the other solution being $[-p_s, -q_s, -1]^T$ and $[-p_i, -q_i, -1]^T \forall i$. In both cases, the light source and camera are on the same side of the object. Our method enumerates both these global solutions. Without loss of generality, we pick one of the light source directions and proceed to estimate its corresponding shape.

5.3 Estimating Shape Given Source Direction

In this section, we estimate the surface normals across the entire object given an estimated source direction $\hat{\mathbf{s}}$ and a few initial normals at selected pixels, while enforcing both shading and Laplacian constraints. Rather than optimizing over the full 3D space of normals, we reduce the search space by restricting normals to lie within a *shading* cone—the set of unit vectors consistent with the observed shading η and source direction $\hat{\mathbf{s}}$.

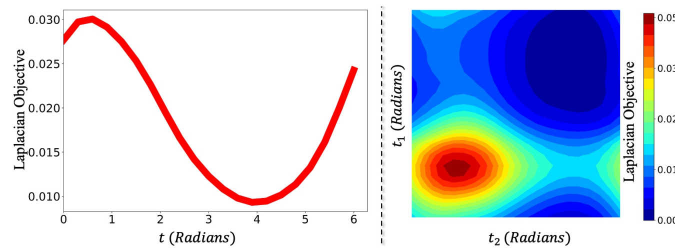

Fig. 5. **Left**: Plot of the objective function in Eq. 19 vs t_i for an unknown vertex i when the other $t_j, j \in \mathcal{N}(i)$ values are known. **Right**: Objective function shown as heatmap for two unknown t when other t values are known. These typical plots show that there exists a single global minimum that satisfies the cotangent weight constraints. Note that t represents an angle that wraps around along the boundary.

We reparametrize each normal as a point on this cone using a 1D polar angle $t \in (0, 2\pi]$:

$$\mathbf{n}(t) = \eta\,\hat{\mathbf{s}} + \sqrt{1-\eta^2}\; R(\hat{\mathbf{s}}) \begin{pmatrix} \cos t \\ \sin t \\ 0 \end{pmatrix} \tag{18}$$

In this formulation, $\eta\hat{\mathbf{s}}$ is the component of the normal aligned with the light source direction, while the perpendicular component (scaled by $\sqrt{1-\eta^2}$) lies in the plane orthogonal to $\hat{\mathbf{s}}$. The rotation matrix $R(\hat{\mathbf{s}})$ maps vectors from the standard xy-plane to this perpendicular plane, and the vector $[\cos t \;\; \sin t \;\; 0]$ sweeps out the circle of possible directions. This reparametrization ensures that shading constraints are inherently satisfied, reducing the problem to selecting the correct t values that match the Laplacian constraint.

To estimate the surface shape, we define an objective over a subset of vertices $\mathcal{S}$ for which Laplacian weights are to be matched. The optimization problem becomes:

$$\mathbf{t}^* = \arg\min_{\mathbf{t}} \sum_{i \in \mathcal{S}} \sum_{k \in \mathcal{N}(i)} \|w_{ik} - \hat{w}_{ik}(\mathbf{t})\|^2 \tag{19}$$

Here, $\mathbf{t}^* = \{t_i \mid i \in \mathcal{S}\}$ represents the optimal polar angles for the selected vertices. The set $\mathcal{N}(i)$ contains neighboring vertices of i in the mesh, w_{ik} are the known cotangent weights, and $\hat{w}_{ik}(\mathbf{t})$ are the weights derived from the current normal estimates.

To make the optimization tractable, we adopt an expanding frontier strategy—estimating normals for a subset of regions at a time. We begin with a coarse multi-dimensional grid search over $\mathbf{t} \in \prod_{i \in \mathcal{S}} (0, 2\pi]$ to find a promising initialization, then refine the estimate via gradient descent. The resulting $\mathbf{t}^*$ provides accurate normal estimates that satisfy both shading and Laplacian constraints. The next subsection details how this localized estimation integrates into our global shape reconstruction pipeline.

5.4 Overall Framework

Algorithm 1 outlines the pseudo-code of our complete Shape from Laplacian and Shading (SFLS) method. We follow a three-stage approach that systematically reconstructs surface normals and light source direction while ensuring consistency with both Laplacian weights and shading observations.

In the first stage, we identify near-planar patches $\mathcal{F}$ on the surface as shown in Fig. 4 and compute their initial normal estimates $\mathbf{n}_{\mathcal{F}}$ using the method described in Section 5.1. These planar regions provide reliable starting points for our reconstruction since their geometric properties are well-captured by the distinctive patterns in their Laplacian weights.

The second stage involves an alternating optimization between lighting estimation and normal refinement. Using the initial normals from planar patches, we estimate the light source direction $\hat{\mathbf{s}}$ through the reverse photometric stereo approach detailed in Section 5.2. With this lighting estimate, we then refine the normals $\mathbf{n}_{\mathcal{F}}$ to better satisfy both shading constraints (consistency with observed intensities $\boldsymbol{\eta}_{\mathcal{F}}$) and Laplacian constraints (consistency with mesh geometry). This alternating process continues until change in both the lighting direction and normal estimates converge within thresholds ϵ_s and ϵ_n, respectively.

In the final stage, we propagate normal estimates from the confidently reconstructed regions to the entire surface through an expanding frontier strategy. Starting with the set of planar patches $\mathcal{M} = \mathcal{F}$ whose normals have been reliably estimated, we iteratively expand to include neighboring vertices $\mathcal{N}(\partial\mathcal{M})$ along the boundary $\partial\mathcal{M}$ of the current region. For each expansion step, we compute normals for the expanded set $\mathcal{K}$ using the method described in Section 5.3, which ensures the new normals respect both the observed shading and the geometric constraints imposed by the Laplacian.

To geometrically interpret our normal propagation method, consider the mesh connectivity illustrated in Fig. 2. When estimating normals for vertices adjacent to those with known normals, we solve for the optimal polar angles $\mathbf{t}^*$ in our cone parametrization. This typically involves a one- or two-dimensional search over the parameter space. Figure 5 shows empirical plots of the objective function described in Eq. 19 for vertices requiring either one- or two-dimensional search of $\mathbf{t}^*$ values, indicating a single local minimum as theoretically described in Section 4.2.

Through this progressive expansion, normals are propagated across the entire mesh while maintaining consistency with both local geometry (encoded in the Laplacian) and the global lighting model. The algorithm terminates when all vertices in the mesh $\mathcal{V}$ have been assigned normals, resulting in a complete surface reconstruction that accurately captures both the fine geometric details and the overall shape of the object.

5.5 Discussion on Limitations of Theory

In this section, we discuss the limitations of our theory.

Case 1: Our light source estimation using reverse photometric stereo requires at least three non-coplanar near-planar patches. In the extreme case when the shape is simply a plane, we cannot estimate the light source direction and consequently cannot unambiguously estimate shape. However, if the light source direction is provided, even a simple plane can be unambiguously reconstructed.

Case 2: In the special case of fronto-parallel lighting, the surface normals for both concave and convex shapes lie within the shading cone at all points on the surface and the ambiguities cannot be resolved.

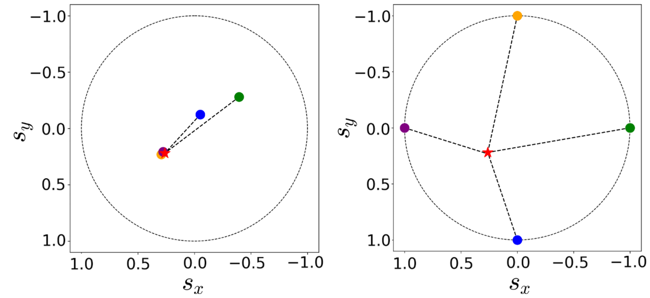

Fig. 6. **Left:** Estimated light source directions for complex-shaped objects, initialized using their near-planar surface patches and progressively refined to converge toward the true source (indicated by a star). **Right:** Convergence behavior of our refinement strategy starting from extreme initial guesses. Here, colors blue, orange, purple and green indicate droplet wave, wineglass, bunny and Igea objects respectively.

Case 3: For directional lighting, surface normals along the line $p = -\frac{q_s}{p_s}q$ and their flipped counterparts produce the same shading and Laplacian weights. In the extreme case when the shape only contains normals along this line, we cannot resolve the ambiguity. Also note that, if all normals in a shape are collinear in the pq space, the light source can also not be estimated as in Case 1. However, for general shapes, a nearby pixel would have a normal outside the above line. In this case, our propagation strategy allows picking the correct sign for the normal based on the nearby pixel.

Outside these special cases, the combination of shading and Laplacians provides a powerful way to resolve inherent shape ambiguities in each modality. Recall that, for unknown light source direction, a global binary ambiguity still exists. However, our approach estimates both possible light source direction and its corresponding shape.

We note that the regular six-neighbor mesh was chosen for symmetry and computational convenience, and the derived equations in Sec. 4.1 are not restricted to this configuration – it can be extended to other mesh connectivities.

6 Experimental Results

We evaluate the theory and effectiveness of our proposed method through simulations on a variety of objects and under varying levels of input noise. In Sec. 6.1, we analyze our light source estimation algorithm. Next, in Sec. 6.2, we compare our shape estimation algorithm against three baselines: (A) Shape from Laplacian, which optimizes (p_i, q_i) to match the given Laplacian weights, (B) Shape from Shading, which optimizes (p_i, q_i) to the match the given shading assuming known source direction, and (C) Analysis-by-synthesis, which optimizes for both (p_s, q_s) and (p_i, q_i) using the sum of Laplacian error and Shading error as loss. In Sec. 6.3, we analyze the robustness of our algorithm to noise. Finally, in Sec. 6.4, we analyze the performance of our algorithm in the presence of inter-reflections and complex BRDFs. Our current Python implementation is not optimized for computational time and the approach takes several tens of minutes per scene.

6.1 Source Estimation Results

Table 1 presents the angular error between the estimated and ground truth light source directions across various simulated objects. Initial estimates based on near-planar patch

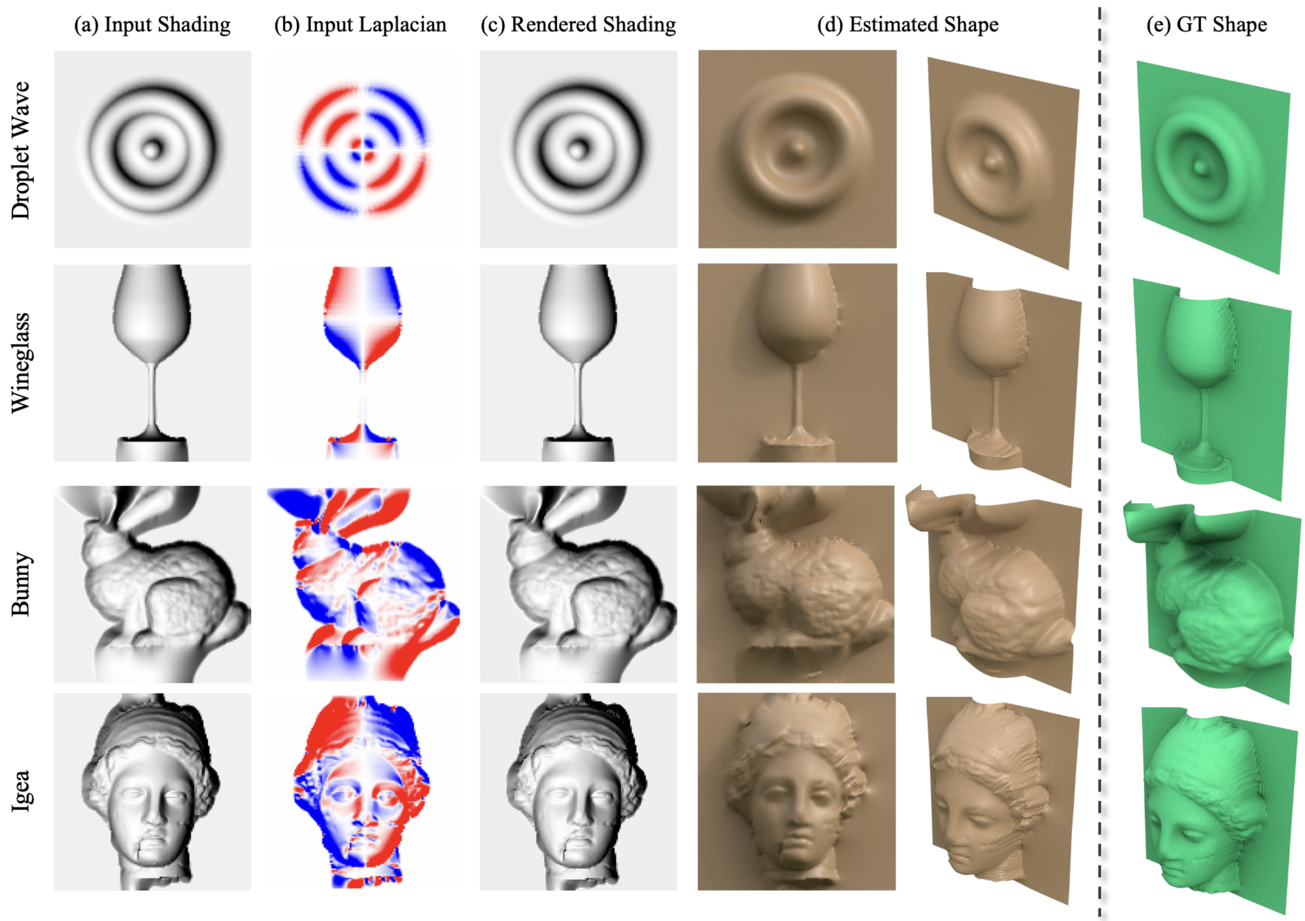

Fig. 7. Shape estimation results for a simulated object using shading and Laplacian as inputs. Figures **(a)** and **(b)** show the inputs cues: shading and Lapalcian, respectively. The Laplacian is visualized as the magnitude of Gaussian curvature, which is an intrinsic geometric property of shape. Figure **(c)** presents the rendering of the estimated lighting and reconstructed surface according to the algorithm described in Sec 5. Note that this rendering is derived entirely from the unknown lighting and unknown shape. Figure **(d)** shows the depth map recovered from the reconstructed surface, and it can be seen that it aligns with the ground truth in **(e)**. The estimated shape and ground truth are both scaled, and the mean distance between them was at most 1.58%. Detailed quantitative results are summarized in Table 2.

TABLE 1
Angular error (in degrees) between estimated and true lighting direction, evaluated after a single iteration of our refinement strategy and after convergence.

	Droplet Wave	Wineglass	Bunny	Igea
Initial Estimate	27.02	2.026	1.154	48.94
After Refinement	0.088	0.029	0.038	0.038

TABLE 2
Average percentage distance error with the ground truth for shapes estimated using shape from Laplacian [8], shape from shading [11] (Fig. 8), and proposed shape from Laplacian and shading (Fig. 7), computed after scaling each shape to fit within a unit bounding box.

	Droplet Wave	Wineglass	Bunny	Igea
Shape from Laplacian	3.05%	3.16%	8.51%	6.60%
Shape from Shading	4.17%	2.75%	6.20%	8.60%
Analysis by Synthesis	3.01%	3.82%	5.94%	11.34%
SFLS	0.05%	0.23%	0.26%	1.58%

TABLE 3
Quantitative results for a droplet wave object with increasing Gaussian noise applied to the inputs of both the shading and Laplacian weights.

Added Gaussian Noise	$\sigma = 0.01$	$\sigma = 0.1$	$\sigma = 0.2$
Avg. distance error	$5e^{-3}$	$1e^{-2}$	$5e^{-2}$
% Distance error	0.45%	1.29%	5.05%

TABLE 4
Quantitative results for a droplet wave object reconstructed from simulated noisy thermal videos with Gaussian noise and Gaussian + fixed pattern noise that is commonly seen in thermal cameras.

	Gaussian	Gaussian + Fixed Pattern
Avg. distance error	$8.1e^{-3}$	$1.2e^{-2}$
% Distance error	0.81%	1.20%

normals result in errors as high as $49°$ for some cases. However, by iteratively refining the estimates through alternating reverse photometric stereo and Laplacian-constrained normal refinement, the angular error is reduced dramatically to below $0.09°$ as visualized in Fig. 6. Additionally, Fig. 6 (right) demonstrates that our method remains effective even when starting from extreme initial guesses, highlighting the effectiveness of our iterative strategy.

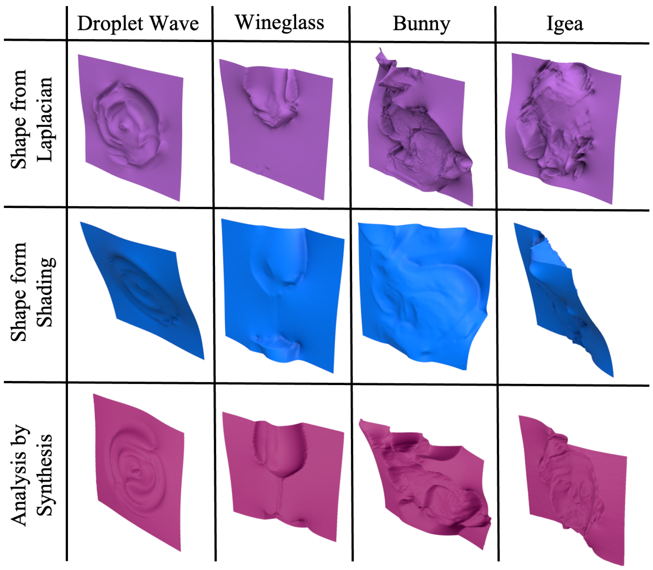

Fig. 8. Shape estimation results using only Laplacian or shading cues, and joint optimization over lighting and shape using an analysis-by-synthesis approach. Due to the non-convex nature of the underlying optimization, the results exhibit ambiguities—specifically, locally binary ambiguities from Laplacian constraints and conic ambiguities from shading cues—reflecting the presence of multiple global minima. The light source estimation errors with the analysis-by-synthesis approach were $96.12°$, $25.68°$, $17.70°$, and $15.20°$ from left to right object, compared to our approach shown in Table 1.

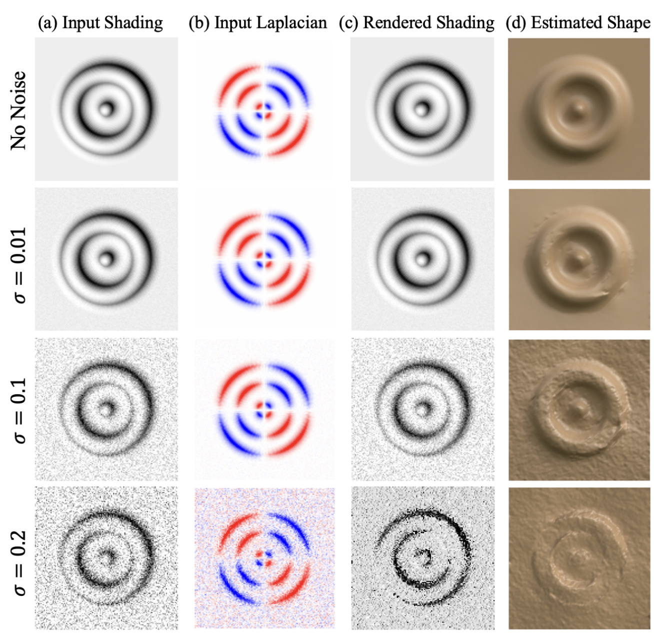

Fig. 9. Shape estimation with varying levels of additive gaussian noise. **(a)** and **(b)** are the input shading and Laplacian, **(c)** the rendered shading image estimated from lighting and shape, and **(d)** the estimated depth-map. From the top to bottom, results are shown for no noise, and for Gaussian noise with standard derivations $\sigma = 0.01$, $\sigma = 0.1$, and $\sigma = 0.2$, respectively. Although greater noise levels introduce more local distortion in the reconstructed shape, our proposed method still yields a reasonable estimation overall. Table 3 presents the corresponding quantitative numbers.

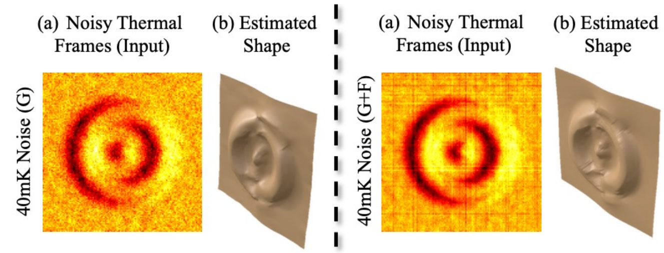

Fig. 10. We simulated and reconstructed shape from the noisy thermal videos with Gaussian noise **(a,b)** and the Gaussian + fixed pattern noises **(c,d)**. We used the algorithm from [8] to estimate Shape Laplacian and we apply median filtering and total variation denoising before using them as input to our shape estimation.

TABLE 5
Quantitative results for a droplet wave for two cases (1) interreflections on a diffuse surface and (2) interreflections on a surface with complex BRDF, with roughness parameter $\alpha = 0.4$.

	$\theta = 30$	$\theta = 50$	$\theta = 70$
Interreflections	0.06%	0.31%	1.23%
BRDF with interreflections	0.84%	0.74%	0.81%

6.2 Shape Estimation Results

Figure 8 shows the qualitative results of shape estimation using the three baselines. Shape from Laplacian estimates the correct local curvature while the overall shape is incorrect due to local concave/convex ambiguities. On the other hand, Shape from Shading estimates shapes that often have lower curvature than the true shape. The analysis-by-synthesis approach, i.e. joint optimization over lighting and shape, is simple in principle but gets stuck in local minima. Note that for every source direction, there exist many corresponding shapes that exactly satisfies the shading equation. Thus, fixing the source direction is crucial for converging to the correct solution. Unlike the shading equation, the Shape Laplacians at near-planar pixels can be solved in closed form, which we then use to estimate source. Also, joint optimization requires a hyperparameter to balance photometric error and Laplacian error, while our proposed t formation satisfies the shading equation by construction. The corresponding results from our proposed method is shown in Fig. 7 along with the ground truth shape. By combining the correct curvature information from Shape Laplacians and the correct direction information from Shading, our algorithm faithfully recovers the 3D shapes even for complex geometries.

To quantitatively evaluate the performance, we compare the estimated meshes against ground truth meshes using MeshLab after global registration. Table 2 reports the average percentage distance errors. Our proposed Shape-from-Laplacian-and-Shading (SFLS) method achieves high recon-

TABLE 6
Quantitative results for a droplet wave object reconstructed from simulated BRDFs with interreflections. The lighting was set at $\theta = 30$.

	$\alpha = 0.4$	$\alpha = 0.3$	$\alpha = 0.2$
BRDFs with interreflections	0.84%	1.42%	1.77%

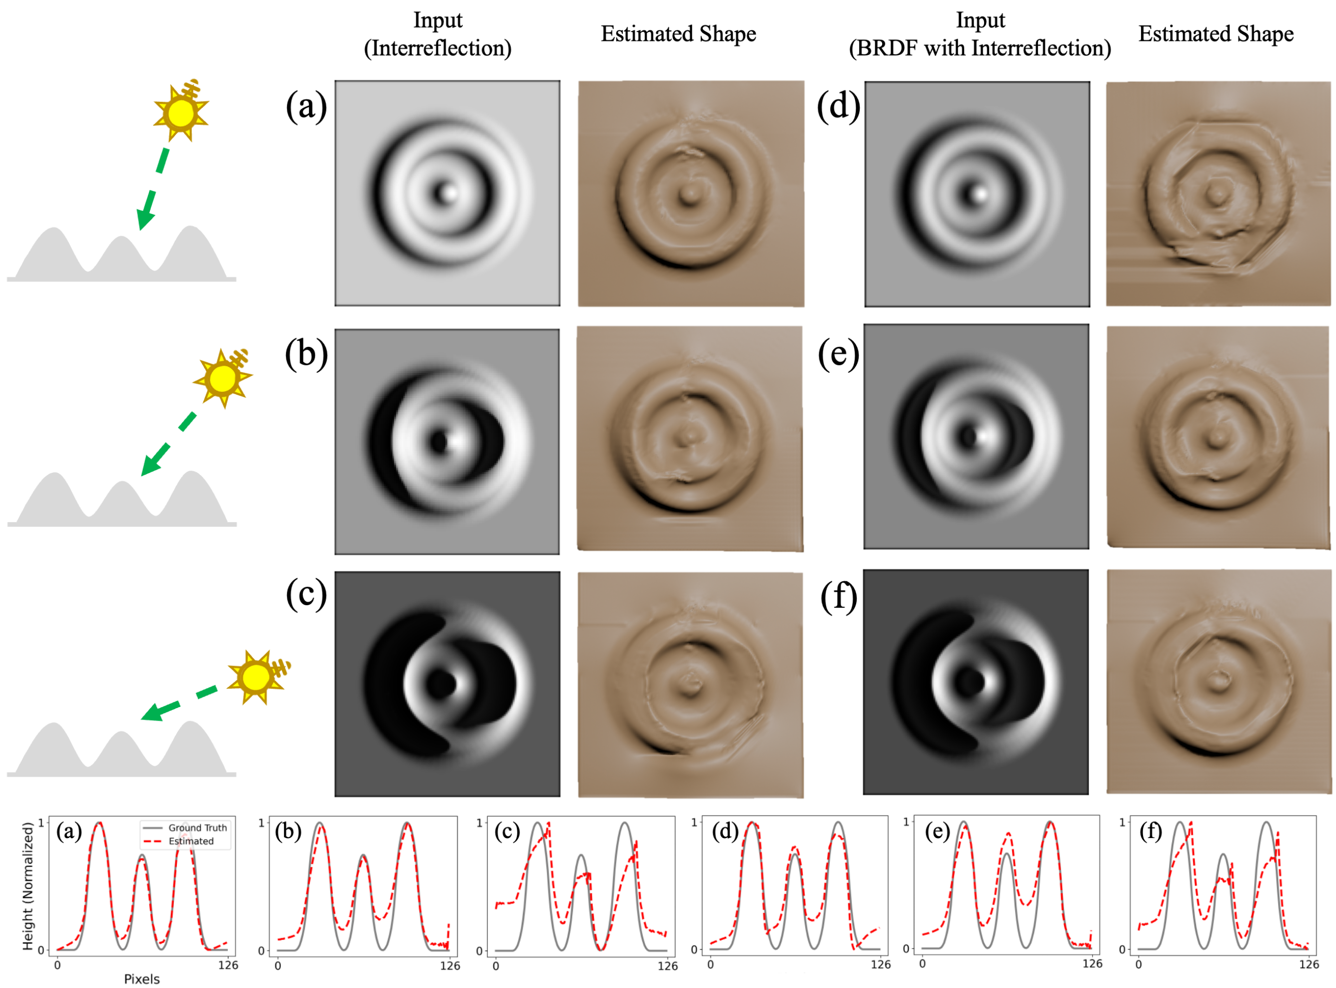

Fig. 11. Shape estimation results for a simulated object under known lighting conditions, using Laplacian and shading as inputs for two cases: (1) shading includes interreflections with the Lambertian BRDF and (2) shading includes interreflections with a non-lambertian BRDF with roughness parameter $\alpha = 0.4$. Figures **(a)**, **(b)** and **(c)** show pairs of the input shading image rendered using Mitsuba and the estimated shape. Similarly, Figures **(d)**, **(e)** and **(f)** show the corresponding results for the non-lambertian BRDF case. Below the visual results, we also show the corresponding depth profiles comparing the estimated shapes with the ground truth along the vertical center line of the image. The mean distance between them was at most 1.23%. Detailed quantitative results are in Table 5.

struction accuracy, with errors as low as 1.6% for complex geometries such as the Igea object.

6.3 Robustness to Noise

The inputs to our algorithm are the Shading image and Shape Laplacian weights. To evaluate robustness, we added Gaussian noise to both these inputs. Table 3 shows the percentage distance errors for the reconstructed droplet shape under increasing noise levels. Despite significant input degradation, our method maintains errors below 1.3%, demonstrating good resilience. Visual results in Fig. 9 confirm the robustness of the proposed pipeline.

While additive gaussian noise is a reasonable noise model for Shading images, the noise model for Shape Laplacians could be more complex as they are estimated from thermal videos which suffer from both gaussian and fixed pattern noises. Therefore, we performed additional experiments by simulating thermal videos and estimating Shape Laplacians using the method in [8]. We additionally use median filtering and total variation denoising on the estimated Shape Laplacians as pre-processing steps. Fig. 10 shows the reconstructed shapes from thermal video with added Gaussian noise, and Gaussian and fixed pattern noise. Table 4 shows the distance errors for each case. The errors are within 1.2% even under noise conditions commonly observed in real-world thermal imaging.

6.4 Interreflections, shadows and complex BRDFs

In this section, we analyze our algorithm in the presence of complex light transport effects such as shadows, interreflections and non-Lambertian BRDFs. Our light source estimation algorithm based on reverse photometric stereo fails in most of these cases with errors above $40°$. Therefore, we restrict our analysis to the case of known lighting.

The behavior of our algorithm varies depending on the amount of interreflection present in the scene. In regions where the direct component (for Lambertian surfaces, this is the cosine shading) dominates interreflections, our approach performs reasonably even if it produces slightly shallower shapes (as noted in classical works [24]). But when interreflections dominate, using just the cosine shading model can result in larger errors. In particular, in shadow regions where there are only interreflections, the surface patches may be estimated to be facing away from the source when

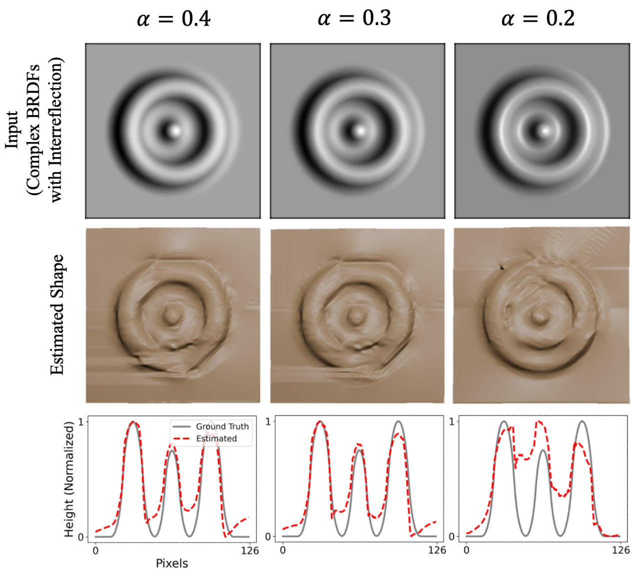

Fig. 12. Shape estimation results using shading images with interreflections under different roughness parameters as input. As the surface becomes smoother (i.e., as α decreases), the specular highlights in the shading introduces rapid changes in surface normals. The depth profiles comparing the estimated shapes with the ground truth along the vertical center line of the image. Quantitative results are summarized in Table 6.

they are actually oriented towards the source. Thus, in such cases, the binary ambiguities in Shape Laplacian will not be resolved correctly using our simple shading model.

To experimentally confirm these observations, we simulate the Droplet Wave shape under three different lighting conditions with interreflections rendered using Mitsuba. We further analyze the same scenes with a non-Lambertian BRDF by varying the roughness (α) parameter. Fig. 11 shows the corresponding shape estimation results. When the lighting is close to fronto-parallel, most regions have a dominant direct component and our shape estimation resolves the concave / convex ambiguity. As lighting goes toward lower polar angles, more regions are shadowed and the shape estimation progressively degrades for both Lambertian and non-Lambertian BRDFs. Table 5 reports the corresponding quantitative errors for these experiments. Fig. 12 compares the shape estimation results with varying roughness parameters. As shown in Table 6, the errors increase with lower roughness (more specular). The profile plots in the bottom row show that the specular regions in the shading introduces rapid changes in surface normals near specular highlights.

7 Conclusion

Heat conduction has recently emerged as a novel way to determine the Shape Laplace-Beltrami operator [8], which is an intrinsic representation of shape, under unknown lighting and even for non-Lambertian surfaces. However, recovering the explicit shape suffers from a local binary convex/concave ambiguity. In contrast, estimating the shape from a single shading image is a well-known ill-posed problem even for a Lambertian scene with uniform albedo and known source direction. At each pixel, there is a cone of possible normals around the source direction that produce the same shading image. In this paper, we have shown that, outside special cases, only one of the convex/concave solutions to the Shape Laplacians lies on the Shading cone.

Our key insight is that the Laplacian weights can be directly computed from the vertex normals that is used in the shading equation, thus enabling a seamless integration of two complimentary cues. Unlike Shape-from-Shading algorithms that use an arbitrary smoothness term, the shape Laplacian encodes the true curvature of the shape. As a result, combining shading and Laplacian constraints naturally produces a smooth optimization landscape with only one global minimum. Interestingly, for homogeneous albedo, shading is directly proportional to the absorbed light [9]. This enables consistent integration of shading with the Laplacian cues only using heat conduction without requiring a separate shading image. Extending our theory to more complex light transport (near-source lighting, interreflections, complex BRDFs) and perspective cameras is an interesting avenue for future work.

Acknowledgments

The research was supported by NSF Award IIS-2107236 and NSF-NIFA AI Institute for Resilient Agriculture (AIIRA). Kao Corporation supported Akihiko Oharazawa's effort.

References

[1] R. J. Woodham, "Photometric stereo: A reflectance map technique for determining surface orientation from image intensity," in *Image understanding systems and industrial applications I*, vol. 155. SPIE, 1979, pp. 136–143.

[2] P. N. Belhumeur, D. J. Kriegman, and A. L. Yuille, "The bas-relief ambiguity," *International journal of computer vision*, vol. 35, no. 1, pp. 33–44, 1999.

[3] M. K. Chandraker, F. Kahl, and D. J. Kriegman, "Reflections on the generalized bas-relief ambiguity," in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 1. IEEE, 2005, pp. 788–795.

[4] E. Prados and O. Faugeras, "Shape from shading: a well-posed problem?" in *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, vol. 2. IEEE, 2005, pp. 870–877.

[5] H. K. Berthold, *Shape from Shading: A Method for Obtaining the Shape of a Smooth Opaque Object from One View*, 1970.

[6] J. T. Barron and J. Malik, "Shape, illumination, and reflectance from shading," *IEEE transactions on pattern analysis and machine intelligence*, vol. 37, no. 8, pp. 1670–1687, 2014.

[7] X. Han, T. Zickler, and K. Nishino, "Multistable shape from shading emerges from patch diffusion," *Advances in Neural Information Processing Systems*, vol. 37, pp. 34 686–34 711, 2024.

[8] S. Narayanan, M. Ramanagopal, M. Sheinin, A. C. Sankaranarayanan, and S. G. Narasimhan, "Shape from heat conduction," in *Computer Vision – ECCV 2024*. Springer Nature Switzerland, 2025, pp. 426–444.

[9] M. Ramanagopal, S. Narayanan, A. C. Sankaranarayanan, and S. G. Narasimhan, "A theory of joint light and heat transport for lambertian scenes," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2024, pp. 11 924–11 933.

[10] K. Ikeuchi and B. K. Horn, "Numerical shape from shading and occluding boundaries," *Artificial intelligence*, vol. 17, no. 1-3, pp. 141–184, 1981.

[11] R. T. Frankot and R. Chellappa, "A method for enforcing integrability in shape from shading algorithms," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 10, no. 4, pp. 439–451, 1988.

[12] K. M. Lee and C.-C. J. Kuo, "Shape from shading with a generalized reflectance map model," *Computer vision and image understanding*, vol. 67, no. 2, pp. 143–160, 1997.

[13] R. Zhang, P.-S. Tsai, J. E. Cryer, and M. Shah, "Shape-from-shading: a survey," *IEEE transactions on pattern analysis and machine intelligence*, vol. 21, no. 8, pp. 690–706, 1999.

[14] K. Heal, J. Wang, S. J. Gortler, and T. Zickler, "A lighting-invariant point processor for shading," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 94–102.

[15] S. Richard, "Fast shape from shading," *CVGIP: Image Understanding*, vol. 53, no. 2, pp. 129–153, 1991.

[16] T. Ping-Sing and S. Mubarak, "Shape from shading using linear approximation," *Image and Vision Computing*, vol. 12, no. 8, pp. 487–498, 1994.

[17] T. Maeda, Y. Wang, R. Raskar, and A. Kadambi, "Thermal non-line-of-sight imaging," in *2019 IEEE International Conference on Computational Photography (ICCP)*, 2019, pp. 1–11.

[18] M. Sheinin, A. C. Sankaranarayanan, and S. G. Narasimhan, "Projecting trackable thermal patterns for dynamic computer vision," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 25 223–25 232.

[19] E. Kerr, T. McGinnity, and S. Coleman, "Material classification based on thermal properties — a robot and human evaluation," in *2013 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, 2013, pp. 1048–1053.

[20] A. Dashpute, V. Saragadam, E. Alexander, F. Willomitzer, A. Katsaggelos, A. Veeraraghavan, and O. Cossairt, "Thermal spread functions (tsf): Physics-guided material classification," *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun 2023.

[21] Y. Nagase, T. Kushida, K. Tanaka, T. Funatomi, and Y. Mukaigawa, "Shape from thermal radiation: Passive ranging using multi-spectral lwir measurements," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 12 651–12 661.

[22] K. Tanaka, N. Ikeya, T. Takatani, H. Kubo, T. Funatomi, V. Ravi, A. Kadambi, and Y. Mukaigawa, "Time-resolved far infrared light transport decomposition for thermal photometric stereo," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 43, no. 6, pp. 2075–2085, 2021.

[23] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, p. 381–395, 1981.

[24] S. K. Nayar, K. Ikeuchi, and T. Kanade, "Shape from interreflections," *Int. J. Comput. Vis.*, vol. 6, no. 3, pp. 173–195, 1991. [Online]. Available: https://doi.org/10.1007/BF00115695

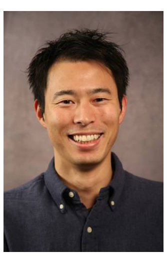

Akihiko Oharazawa is a visiting researcher at the Robotics Institute, Carnegie Mellon University, where he works with Prof. Srinivas Narasimhan. He is also a Research Scientist at the Skincare Laboratory of Kao Corporation. He received his M.S. in Comprehensive Human Science from the University of Tsukuba. His current research interests lie in physics-based vision, computer vision, and computational imaging.

Sriram Narayanan is a Ph.D. student in the Robotics Institute at Carnegie Mellon University, advised by Prof. Srinivasa Narasimhan. His research lies at the intersection of computer vision, computational imaging, and robotics, with a focus on physics-based methods for heat and light transport. Prior to CMU, he was a Research Scholar at NEC Laboratories America. He is a recipient of the Uber Presidential Fellowship.

Mani Ramanagopal is a Project Scientist at the Robotics Institute at Carnegie Mellon University. He obtained his Ph.D. in Robotics from the University of Michigan in 2022 and a B.Tech & M.Tech in Electrical Engineering from Indian Institute of Technology Madras in 2016. He is the recipient of the Richard F. and Eleanor A. Towner Prize for Distinguished Academic Achievement for Outstanding Graduate Student. His research interests are in robotics, vision, and imaging, with a focus on physics-based approaches.

Srinivasa G. Narasimhan is the Interim Director and U. A and Helen Whitaker Professor of the Robotics Institute at Carnegie Mellon University. He obtained his Ph.D. from Columbia University in Dec 2003. His group focuses on novel techniques for imaging and illumination to enable applications in vision, graphics, robotics, agriculture, intelligent transportation, and medical imaging. His works have received more than a dozen Best Paper or Best Demo or Honorable mention awards at major conferences. In addition, he has received the Ford URP Award (2013), Okawa Research Grant (2009) and the NSF CAREER Award (2007). He is the co-inventor of programmable headlights, Aqualux 3D display, Assorted pixels, Motion-aware cameras, Episcan360, Episcan3D, EpiToF3D, and programmable triangulation light curtains. He served on the editorial board of the International Journal of Computer Vision and serves as Senior or Lead Area Chair of top computer vision conferences.

APPENDIX

Laplacian weights from vertex normals: Here, we provide the detailed equations for the Laplacian weights in Fig. 2 for all six edges of a vertex.

$$
\begin{aligned}
w_{ia} &= \frac{1}{2}\left(\frac{p_b^2 + p_b q_a + 1}{\sqrt{1+q_a^2+p_b^2}} + \frac{p_i^2 + p_i q_a + 1}{\sqrt{1+q_a^2+p_i^2}}\right) \\
w_{ij} &= \frac{1}{2}\left(\frac{q_a p_i + q_a^2 + 1}{\sqrt{1+q_a^2+p_i^2}} + \frac{q_j p_i + q_j^2 + 1}{\sqrt{1+q_j^2+p_i^2}}\right) \\
w_{ik} &= \frac{1}{2}\left(\frac{-q_j p_i}{\sqrt{1+q_j^2+p_i^2}} - \frac{p_l q_i}{\sqrt{1+p_l^2+q_i^2}}\right) \\
w_{il} &= \frac{1}{2}\left(\frac{p_l q_i + p_l^2 + 1}{\sqrt{1+p_l^2+q_i^2}} + \frac{p_c q_i + p_c^2 + 1}{\sqrt{1+p_c^2+q_i^2}}\right) \\
w_{ic} &= \frac{1}{2}\left(\frac{q_i^2 + q_i p_c + 1}{\sqrt{1+p_c^2+q_i^2}} + \frac{q_b p_c + q_b^2 + 1}{\sqrt{1+p_c^2+q_b^2}}\right) \\
w_{ib} &= \frac{1}{2}\left(\frac{-p_c q_b}{\sqrt{1+p_c^2+q_b^2}} - \frac{p_b q_a}{\sqrt{1+q_a^2+p_b^2}}\right)
\end{aligned} \tag{20}
$$

Near-planar patch solutions: Let $D = w_1 + w_2 + 2w_3$. The two solutions to Eq. 15 are

$$
x = \frac{D \pm \sqrt{D^2 - 4}}{2}. \tag{21}
$$

As the solutions need to be real, $D^2 - 4 \geq 0$. From Eq. 14, $D > 0$ as all the terms on the left are positive. The above two inequalities imply $D \geq 2$. We can then write

$$
\begin{aligned}
(D-2)^2 &< (D-2)(D+2) \\
D - 2 &< \sqrt{D^2-4} \\
D - \sqrt{D^2-4} &< 2 \\
\frac{D - \sqrt{D^2-4}}{2} &< 1
\end{aligned} \tag{22}
$$

Since $D > 2$, it is trivial to verify that

$$
x = \frac{D}{2} + \frac{\sqrt{D^2-4}}{2} \geq 1. \tag{23}
$$

When $D = 2$, both solutions are identical ($x = 1$).