Variational methods for photometric 3D-reconstruction

Yvain QUÉAU

CNRS, GREYC laboratory, University of Caen, France

Institut Henri Poincaré, Paris
October 3rd, 2019

Outline

Outline

Shape-from-shading: A Classic Ill-posed Problem

Given an image $\mathbf{I} : \Omega \subset \mathbb{R}^2 \to \mathbb{R}^m$, *Shape-from-Shading* (SfS) consists in inverting the forward photometric model (*image irradiance equation*)

$$\mathbf{I} = \mathcal{R}(z, \boldsymbol{\rho}, \boldsymbol{\ell}) \tag{1}$$

with $\mathcal{R}$ a *radiance* function depending on the unknown depth $z : \Omega \to \mathbb{R}$, surface reflectance $\boldsymbol{\rho} : \Omega \to \mathbb{R}^m$, and incident lighting $\boldsymbol{\ell} : \Omega \to \mathbb{S}^2$.

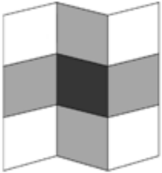

RGB image: $\mathbf{I}$

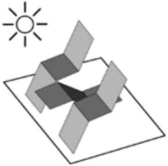

Sculptor's explanation: z

Painter's explanation: $\boldsymbol{\rho}$

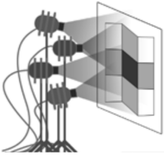

Gaffer's explanation: $\boldsymbol{\ell}$

Illustration of SfS's Ill-posedness

Even with known surface reflectance $\boldsymbol{\rho}$ and incident lighting $\boldsymbol{\ell}$, shape estimation by SfS is an ill-posed inverse problem (Horn, 1970).

Example: two solutions of $\mathbf{I} = \mathcal{R}(z, \boldsymbol{\rho}, \boldsymbol{\ell})$ with $\mathbf{I} :=$ Lena, white reflectance ($\rho \equiv 1$) and frontal lighting ($\boldsymbol{\ell} \equiv [0, 0, -1]^\top$):

Illustration of SfS's Ill-posedness

Even with known surface reflectance ρ and incident lighting $\boldsymbol{\ell}$, shape estimation by SfS is an ill-posed inverse problem (Horn, 1970).

Example: two solutions of $\mathbf{I} = \mathcal{R}(z, \boldsymbol{\rho}, \boldsymbol{\ell})$ with $\mathbf{I} :=$ Lena, white reflectance ($\rho \equiv 1$) and frontal lighting ($\boldsymbol{\ell} \equiv [0, 0, -1]^\top$):

Illustration of SfS's Ill-posedness

Even with known surface reflectance ρ and incident lighting $\boldsymbol{\ell}$, shape estimation by SfS is an ill-posed inverse problem (Horn, 1970).

Example: two solutions of $\mathbf{I} = \mathcal{R}(z, \boldsymbol{\rho}, \boldsymbol{\ell})$ with $\mathbf{I} :=$ Lena, white reflectance ($\rho \equiv 1$) and frontal lighting ($\boldsymbol{\ell} \equiv [0, 0, -1]^\top$):

Illustration of SfS's Ill-posedness

Even with known surface reflectance ρ and incident lighting ℓ, shape estimation by SfS is an ill-posed inverse problem (Horn, 1970).

Example: two solutions of $\mathbf{I} = \mathcal{R}(z, \rho, \ell)$ with $\mathbf{I} :=$ Lena, white reflectance ($\rho \equiv 1$) and frontal lighting ($\ell \equiv [0, 0, -1]^\top$):

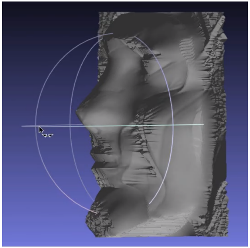

Illustration of SfS's Ill-posedness

Even with known surface reflectance ρ and incident lighting ℓ, shape estimation by SfS is an ill-posed inverse problem (Horn, 1970).

Example: two solutions of $\mathbf{I} = \mathcal{R}(z, \rho, \ell)$ with $\mathbf{I} :=$ Lena, white reflectance ($\rho \equiv 1$) and frontal lighting ($\ell \equiv [0, 0, -1]^\top$):

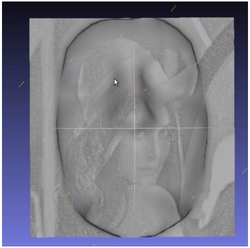

Illustration of SfS's Ill-posedness

Even with known surface reflectance ρ and incident lighting ℓ, shape estimation by SfS is an ill-posed inverse problem (Horn, 1970).

Maximal viscosity solution [Cristiani and Falcone 2007]

Variational solution [Quéau et al. 2017]

Parameterization the irradiance equation $\mathbf{I} = \mathcal{R}(z, \boldsymbol{\rho}, \boldsymbol{\ell})$

Basic Lambertian model:

RGB image

$\mathbf{I} : \Omega \to \mathbb{R}^3$

$=$

Albedo

$\boldsymbol{\rho} : \Omega \to \mathbb{R}^3$

$\odot$

Shading

$S(z, \boldsymbol{\ell}) : \Omega \to \mathbb{R}$

where albedo (Lambertian reflectance) $\equiv$ color, and shading $\equiv$ lighting-geometry interaction.

Parameterization the irradiance equation $\mathbf{I} = \mathcal{R}(z, \boldsymbol{\rho}, \boldsymbol{\ell})$

Shading $\equiv$ lighting-geometry interaction:

RGB image
$\mathbf{I} : \Omega \to \mathbb{R}^3$

$=$

Albedo
$\boldsymbol{\rho} : \Omega \to \mathbb{R}^3$

$\odot$

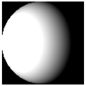

Lighting
$\boldsymbol{\ell} \in \mathbb{S}^2$

$\cdot$

Normals
$\mathbf{n}(z) : \Omega \to \mathbb{S}^2$

where the surface normal $\mathbf{n}$ relates to the depth map z in a **nonlinear** way:

$$\mathbf{n}(z) = \frac{1}{\sqrt{|f\nabla z|^2 + (-z - < \mathbf{p}, \nabla z >)^2}} \begin{bmatrix} f\nabla z \\ -z - < \mathbf{p}, \nabla z > \end{bmatrix}$$

($f > 0$: length, and $\mathbf{p} : \Omega \to \mathbb{R}^2$: centered pixel coordinates).

Parameterization the irradiance equation $\mathbf{I} = \mathcal{R}(z, \boldsymbol{\rho}, \boldsymbol{\ell})$

Extension to first-order spherical harmonics lighting $\boldsymbol{\ell} \in \mathbb{R}^4$:

 $=$  $\odot$ 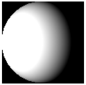 $\cdot$ 

RGB image $\mathbf{I} : \Omega \to \mathbb{R}^3$

Albedo $\boldsymbol{\rho} : \Omega \to \mathbb{R}^3$

Lighting $\boldsymbol{\ell} \in \mathbb{R}^4$

Geometry $\begin{bmatrix} \mathbf{n} \\ 1 \end{bmatrix} (z) : \Omega \to \mathbb{R}^4$

$$I = \mathcal{R}(z, \boldsymbol{\rho}, \boldsymbol{\ell}) := \boldsymbol{\rho} \left\langle \boldsymbol{\ell}, \begin{bmatrix} \mathbf{n}(z) \\ 1 \end{bmatrix} \right\rangle$$

where the surface normal $\mathbf{n}$ relates to the depth map z in a **nonlinear** way:

$$\mathbf{n}(z) = \frac{1}{\sqrt{|f\nabla z|^2 + (-z - < \mathbf{p}, \nabla z >)^2}} \begin{bmatrix} f\nabla z \\ -z - < \mathbf{p}, \nabla z > \end{bmatrix}$$

($f > 0$: length, and $\mathbf{p} : \Omega \to \mathbb{R}^2$: centered pixel coordinates).

Outline

Variational Solving of SfS Equation $I = \mathcal{R}(z, \rho, \ell)$

Assume (for now) that ρ and ℓ are known. Problem reduces to a nonlinear PDE $I := \mathcal{R}(\nabla z)$.

Horn and Brooks 1986: Regularization

Set $(p, q) := \nabla z$ over $\Omega \subset \mathbb{R}^2$

1) Estimate gradient components satisfying integrability: $\min_{p,q} \iint_\Omega (I - \mathcal{R}(p,q))^2 + \lambda(\partial_y p - \partial_x q)^2 \, \mathrm{d}x\mathrm{d}y$
2) Integrate: $\min_z \iint_\Omega \|(p,q) - \nabla z\|^2 \, \mathrm{d}x\mathrm{d}y$

Quéau et al. 2017 (EMMCVPR): Hard constraint

(p, q) is conservative *by construction* $\rightarrow$ Integrated estimation of gradient and depth:

$$\min_{p,q,z} \iint (I - \mathcal{R}(p,q))^2 \, \mathrm{d}x\mathrm{d}y$$

$$\text{s.t. } (p,q) = \nabla z$$

Regularized SfS Model

- Minimal surface regularization over Ω
- (Incomplete) depth prior over $\Omega' \subset \Omega$

$$\min_{p,q,z} \iint_{\Omega} \lambda \left(I - \mathcal{R}(p,q)\right)^2 + \nu \sqrt{1 + p^2 + q^2} \, \mathrm{d}x\mathrm{d}y$$
$$+ \iint_{\Omega'} \mu \left(z - z^0\right)^2 \mathrm{d}x\mathrm{d}y$$
$$\text{s.t. } (p,q) = \nabla z$$

By tuning λ, μ and ν, we may achieve SfS, depth denoising and inpainting, or shading-based depth refinement.

Solving the Regularized SfS Model using ADMM

$(p^{(k+1)}, q^{(k+1)})$ solution of the local, nonlinear least-squares problem: use parallel BFGS iterations

$$\min_{(p,q)} \lambda \|I - \mathcal{R}(p,q)\|^2_{\ell^2(\Omega)} + \nu \left\| \sqrt{1 + p^2 + q^2} \right\|_{\ell^1(\Omega)} + \frac{1}{2\beta} \left\| (p,q) - \nabla z^{(k)} + \theta^{(k)} \right\|^2_{\ell^2(\Omega)}$$

$z^{(k+1)}$ solution of the global, linear least-squares problem: use preconditioned conjugate gradient iterations

$$\min_z \mu \left\| z - z^0 \right\|^2_{\ell^2(\Omega')} + \frac{1}{2\beta} \left\| (p^{(k+1)}, q^{(k+1)}) - \nabla z + \theta^{(k)} \right\|^2_{\ell^2(\Omega)}$$

Auxiliary variable update

$$\theta^{(k+1)} = \theta^{(k)} + (p^{(k+1)}, q^{(k+1)}) - \nabla z^{(k+1)}$$

Application 1: Depth Refinement for MVS Techniques

Input images I^1 and I^2

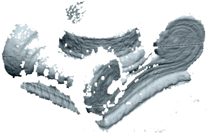

Input depth map z^2

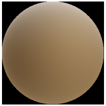

Estimated lighting

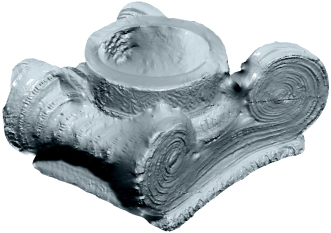

Shading-based refinement

Application 2: SfS under Natural Illumination

Input RGB image

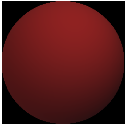

Calibrated lighting

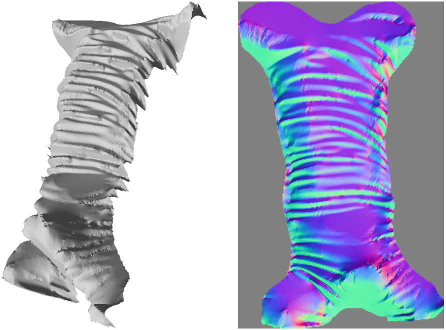

SfS 3D-reconstruction

Application 3: Depth Refinement for RGB-D Sensors

Input RGB image

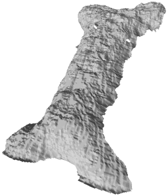

Input depth map

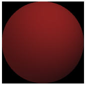

Estimated lighting

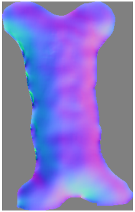

Denoised (minimal surface)

Shading-based refinement

Outline

Problem with RGB-D Sensors

Depth image

Shape

RGB image

Depth image has

- noise and quantization,
- missing areas,
- **coarse resolution.**

RGB image has

- less noise and quantization,
- no missing area,
- **high resolution.**

Goal:

Combine data to get high-resolution shape

Problem with RGB-D Sensors

Depth image

Shape

RGB image

High-resolution shape

Depth image has

- noise and quantization,
- missing areas,
- **coarse resolution.**

RGB image has

- less noise and quantization,
- no missing area,
- **high resolution.**

Goal:

Combine data to get high-resolution shape

Ill-posedness in Depth Super-Resolution

Given a low-resolution depth map $z_0 : \Omega_{LR} \to \mathbb{R}$, *Depth Super-Resolution* (SR) consists in inverting the forward downsampling model

$$z_0 = Dz$$

with $z : \Omega_{HR} \to \mathbb{R}$ the (unknown) high-resolution depth, and D a (rank-deficient) downsampling operator

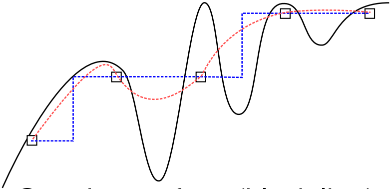

Genuine surface (black line) can be approximated in ∞ many ways (dashed lines) given sparse obervations (rectangles)

$\Rightarrow$ Use SfS to find a shape interpolation which is consistent with the high-resolution RGB image

- Peng et al. 2017 (ICCV)
- Haefner, Quéau, et al. 2018 (CVPR)
- Haefner, Peng, et al. 2019 (PAMI)

Variational Formulation

$$\min_{\substack{z:\Omega_{HR}\to\mathbb{R}\\ \rho:\Omega_{HR}\to\mathbb{R}^3\\ \ell\in\mathbb{R}^4}} \left\| \mathbf{I} - \rho < \ell, \begin{bmatrix} \mathbf{n}(z) \\ 1 \end{bmatrix} > \right\|^2_{\ell_2(\Omega_{HR})} + \mu \left\| z_0 - Dz \right\|^2_{\ell_2(\Omega_{LR})} + \nu P_1(z) + \lambda P_2(\rho)$$

P_1 is a minimal surface regularization term:

$$\begin{aligned} P_1(z) &= \|\mathrm{d}A(z)\|_{\ell_1(\Omega_{HR})} \\ &= \left\| \frac{z}{f}\sqrt{|f\nabla z|^2 + (-z- < \mathbf{p}, \nabla z >)^2} \right\|_{\ell_1(\Omega_{HR})} \end{aligned}$$

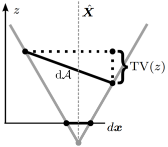

P_2 is a Potts regularization term (**nondifferentiable** and **nonconvex**),

$$P_2(\rho) = \|\nabla \rho\|_{\ell_0(\Omega_{HR})} = \sum_{\mathbf{p}\in\Omega_{HR}} \begin{cases} 0, & \text{if } |\nabla \rho(\mathbf{p})|_F = 0, \\ 1, & \text{otherwise.} \end{cases}$$

Numerical Resolution - Splitting Strategy

Normal $\mathbf{n}$ and minimal surface term $\mathrm{d}A$ depend on z **and** ∇z:

$$\mathbf{n}(z) := \mathbf{n}(z, \nabla z) = \frac{1}{\sqrt{|f\nabla z|^2 + (-z - <\mathbf{p}, \nabla z>)^2}} \begin{bmatrix} f\nabla z \\ -z - <\mathbf{p}, \nabla z> \end{bmatrix}$$

$$\mathrm{d}A(z) := \mathrm{d}A(z, \nabla z) = \frac{z}{f}\sqrt{|f\nabla z|^2 + (-z - <\mathbf{p}, \nabla z>)^2}$$

Introduce splitting $\theta := (z, \nabla z)$ to make optimization tractable:

$$\min_{\substack{z:\Omega_{HR}\to\mathbb{R} \\ \rho:\Omega_{HR}\to\mathbb{R}^3 \\ \ell\in\mathbb{R}^4 \\ \theta:\Omega_{HR}\to\mathbb{R}^3}} \left\| \mathbf{I} - \rho < \ell, \begin{bmatrix} \mathbf{n}(\theta) \\ 1 \end{bmatrix} > \right\|^2_{\ell_2(\Omega_{HR})} + \mu \left\| z_0 - Dz \right\|^2_{\ell_2(\Omega_{LR})} + \nu \left\| \mathrm{d}A(\theta) \right\|_{\ell_1(\Omega_{HR})} + \lambda \left\| \nabla\rho \right\|_{\ell_0(\Omega_{HR})}$$

$$\text{s.t.} \quad \theta = (z, \nabla z)$$

Numerical Resolution - Multi-block ADMM

Given $(\boldsymbol{\rho}^{(k)}, \boldsymbol{\ell}^{(k)}, \boldsymbol{\theta}^{(k)}, z^{(k)})$ at iteration k, we update:

$$\boldsymbol{\rho}^{(k+1)} = \underset{\boldsymbol{\rho}}{\text{argmin}} \left\| \mathbf{I} - \boldsymbol{\rho} < \boldsymbol{\ell}^{(k)}, \begin{bmatrix} \mathbf{n}(\boldsymbol{\theta}^{(k)}) \\ 1 \end{bmatrix} > \right\|^2_{\ell^2(\Omega_{HR})} + \lambda \left\| \nabla \boldsymbol{\rho} \right\|_{\ell^0(\Omega_{HR})}$$

$$\boldsymbol{\ell}^{(k+1)} = \underset{\boldsymbol{\ell}}{\text{argmin}} \left\| \mathbf{I} - \boldsymbol{\rho}^{(k+1)} < \boldsymbol{\ell}, \begin{bmatrix} \mathbf{n}(\boldsymbol{\theta}^{(k)}) \\ 1 \end{bmatrix} > \right\|^2_{\ell^2(\Omega_{HR})}$$

$$\boldsymbol{\theta}^{(k+1)} = \underset{\boldsymbol{\theta}}{\text{argmin}} \left\| \mathbf{I} - \boldsymbol{\rho}^{(k+1)} < \boldsymbol{\ell}^{(k+1)}, \begin{bmatrix} \mathbf{n}(\boldsymbol{\theta}) \\ 1 \end{bmatrix} > \right\|^2_{\ell^2(\Omega_{HR})}$$

$$+ \nu \left\| \mathrm{d} A_{\boldsymbol{\theta}} \right\|_{\ell^1(\Omega_{HR})} + \frac{\kappa}{2} \left\| \boldsymbol{\theta} - (z, \nabla z)^{(k)} + u^{(k)} \right\|^2_{\ell^2(\Omega_{HR})}$$

$$z^{(k+1)} = \underset{z}{\text{argmin}} \; \mu \left\| z_0 - D z \right\|^2_{\ell^2(\Omega_{LR})} + \frac{\kappa}{2} \left\| \boldsymbol{\theta}^{(k+1)} - (z, \nabla z) + u^{(k)} \right\|^2_{\ell^2(\Omega_{HR})}$$

$$u^{(k+1)} = u^{(k)} + \boldsymbol{\theta}^{(k+1)} - (z, \nabla z)^{(k+1)}$$

Qualitative Evaluation

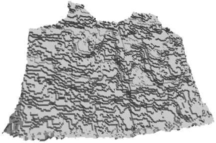

Input depth

Input RGB

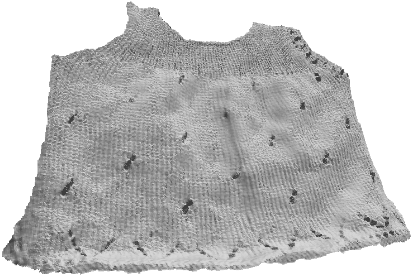

Depth estimate

Qualitative Evaluation

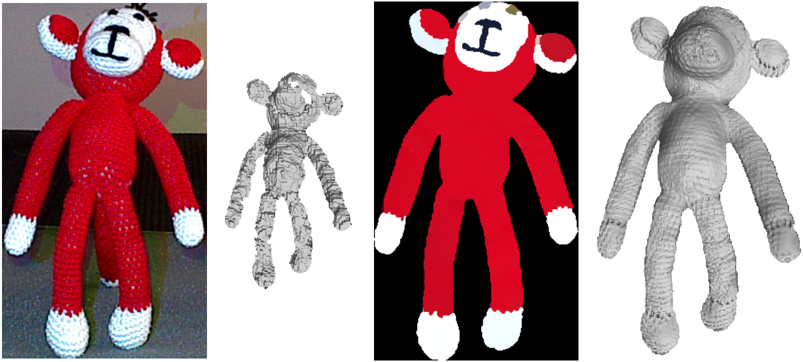

Input RGB | Input depth | Albedo estimate | Depth estimate

Qualitative Evaluation

Input RGB

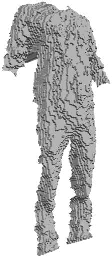

Input depth

Albedo estimate

Depth estimate

Qualitative Evaluation

Input RGB

Albedo estimate

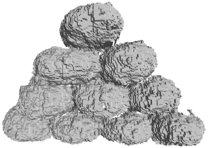

Depth image

Depth estimate

Outline

Motivation: Failure Case of the Previous Approach

When estimated albedo is highly undersegmented, albedo information propagates to geometry:

Input RGB | Input depth | Albedo estimate | Depth estimate

Beyond Regularization: Reflectance Learning

Replacing Potts regularization of reflectance by a deep learning framework circumvents the difficulties of choosing an appropriate regularizer, and simplifies numerics.

Input RGB | Input depth | CNN-albedo estimate | Depth estimate

[Haefner et al., "Photometric Depth Super-Resolution", PAMI 2019]

Reflectance Learning: Idea

Learn a black-box mapping from image to albedo to get rid of man-made Potts prior. Leave the rest (geometry and lighting estimation) to the physics-based variational approach.

$\mapsto$

RGB image (input)

Albedo (output)

Advantage: Move from “piecewise constant albedo” assumption to “same class of objects” assumption.

Reflectance Learning: Database Creation

Rendering (with Blender) of $\approx$ 5000 **faces** with known reflectance:

- 21 faces, each with 15 different expressions
- three different viewpoints
- multiple different lighting conditions

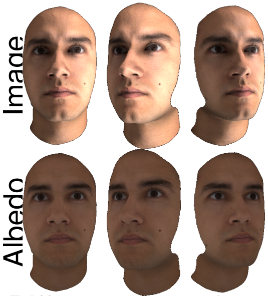

Different viewpoints

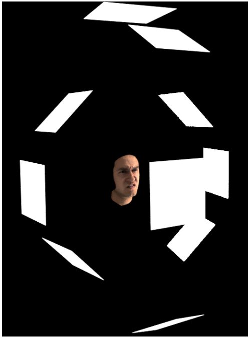
Different lighting conditions

Reflectance Learning: Results

Image-albedo mapping learnt using a U-Net CNN on the synthetic database. Testing on real-world images:

Image | Albedo estimate | Image | Albedo estimate

Reflectance Learning: Back to Depth Super-resolution

The optimization framework gets simpler, as no optimization over $\rho = \rho_{\text{CNN}}$ is needed, so the Potts term $\lambda P_2(\rho)$ disappears:

$$\min_{\substack{z:\Omega_{HR}\to\mathbb{R} \\ \cancel{\rho:\Omega_{HR}\to\mathbb{R}^3} \\ \ell\in\mathbb{R}^4}} \left\| \mathbf{I} - \rho_{\text{CNN}} < \ell, \begin{bmatrix} \mathbf{n}(z) \\ 1 \end{bmatrix} > \right\|^2_{\ell_2(\Omega_{HR})} + \mu \left\| z_0 - Dz \right\|^2_{\ell_2(\Omega_{LR})} + \nu P_1(z) + \cancel{\lambda P_2(\rho)}$$

Take Away Message

- Use variational methods if the physics-based model is simple and realistic (here, for micro-geometry estimation)
- If the phycis-based model is over-complicated or unrealistic, prefer a black-box (here, for reflectance estimation)

Qualitative Results

Input RGB | Input depth | CNN-albedo estimate | Depth estimate

Qualitative Results

Input RGB | Input depth | CNN-albedo estimate | Depth estimate

Qualitative Results

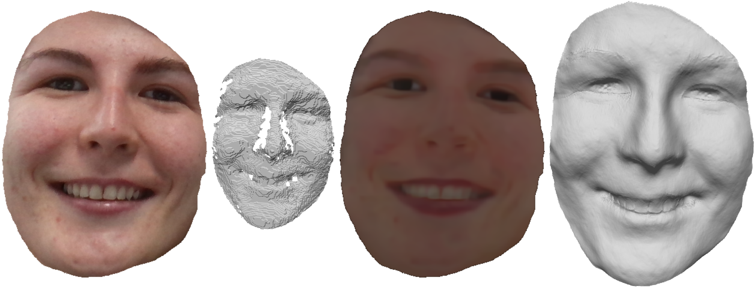

Input RGB | Input depth | CNN-albedo estimate | Depth estimate

Failure Case

RGB image is not a face ⇒ Reflectance information is propagated to geometry

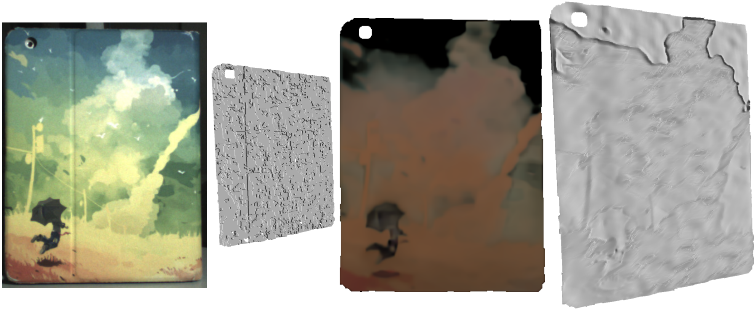

Input RGB | Input depth | CNN-albedo estimate | Depth estimate

Possible remedies: larger training set, or multi-shot approach i.e., photometric stereo.

Outline

Multi-Shot Depth Super-Resolution using Photometric Stereo

n RGB images under varying lighting

n depth images

Albedo estimate

Depth estimate

[Peng et al., "Depth super-resolution meets uncalibrated photometric stereo", ICCV 2017]

Idea

Use uncalibrated photometric stereo instead of shape-from-shading, i.e. go from a single image observation

$$\mathbf{I} = \boldsymbol{\rho} < \boldsymbol{\ell}, \begin{bmatrix} \mathbf{n}(z) \\ 1 \end{bmatrix} >$$

to multiple image observations under **varying lighting**

$$\mathbf{I}^i = \boldsymbol{\rho} < \boldsymbol{\ell}^i, \begin{bmatrix} \mathbf{n}(z) \\ 1 \end{bmatrix} >, \qquad i \in \{1, \ldots, n\}.$$

$\Longrightarrow$ Results in much more constrained $\boldsymbol{\rho}$ and z, due to their independence on $i \in \{1, \ldots, n\}$. No regularization or learning is thus needed:

$$\min_{\substack{z:\Omega_{HR}\to\mathbb{R} \\ \boldsymbol{\rho}:\Omega_{HR}\to\mathbb{R}^3 \\ \{\boldsymbol{\ell}^i\}\in\mathbb{R}^4}} \sum_{i=1}^{n} \left\| \mathbf{I}^i - \boldsymbol{\rho} < \boldsymbol{\ell}, \begin{bmatrix} \mathbf{n}(z) \\ 1 \end{bmatrix} > \right\|_{\ell_2(\Omega_{HR})}^2 + \mu \left\| z_0 - Dz \right\|_{\ell_2(\Omega_{LR})}^2 + \cancel{\nu P_1(z)} + \cancel{\lambda P_2(\boldsymbol{\rho})}$$

Qualitative Results

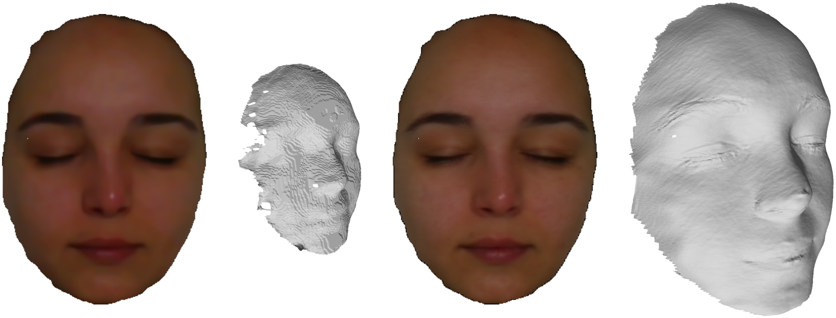

Input RGBs | Input Depth | Albedo Estimate | Depth Estimate

Qualitative Results

Input RGBs

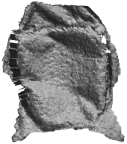

Input Depth

Albedo Estimate

Depth Estimate

Qualitative Results

Input RGBs

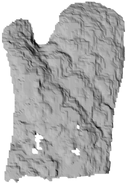

Input Depth

Albedo Estimate

Depth Estimate

Qualitative Results

Input RGBs

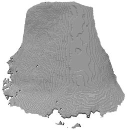

Input Depth

Albedo Estimate

Depth Estimate

Qualitative Results

Input RGBs

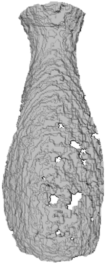

Input Depth

Albedo Estimate

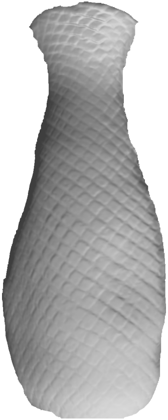

Depth Estimate

What if There is no Depth Prior ?

In theory, the depth prior is not even needed:

Theorem – Brahimi et al. 2019 (Hal 02297643)

There is a **unique** (C^2-depth,reflectance,lighting) solution of:

$$\mathbf{I}^i = \rho < \ell^i, \begin{bmatrix} \mathbf{n}(z) \\ 1 \end{bmatrix} >, \qquad i \in \{1, \ldots, n\},$$

And the solution can be found in closed-form using a spectral approach...

But, spectral approach very sensitive to noise: regularization $+$ non-convex optimization remains the best option.
State-of-the-art heuristic: ballooning initialization, then multi-block ADMM: Haefner, Ye, et al. 2019 (ICCV 2019).

Conclusion and Perspectives

Contributions:

1. A flexible splitting-based numerical framework for photometric 3D-reconstruction
2. Application to depth super-resolution for RGBD sensors

Ongoing work:

- Simultaneous 3D-reconstruction and (Chan-Vese-like) 2D-segmentation: Haefner, Quéau, and Cremers 2019 (3DV)
- Extension to multi-view stereo: Mélou et al. 2019 (SSVM)

- ... Photometric 3D-reconstruction for Arts

Ongoing Work on Cultural Heritage

Bayeux tapestry (represents the conquest of England by William):

High-resolution multi-illumination capture:

Ongoing Work on Cultural Heritage

High-resolution 3D-scanning, for 3D-copies which could be touched by visually-defficients:

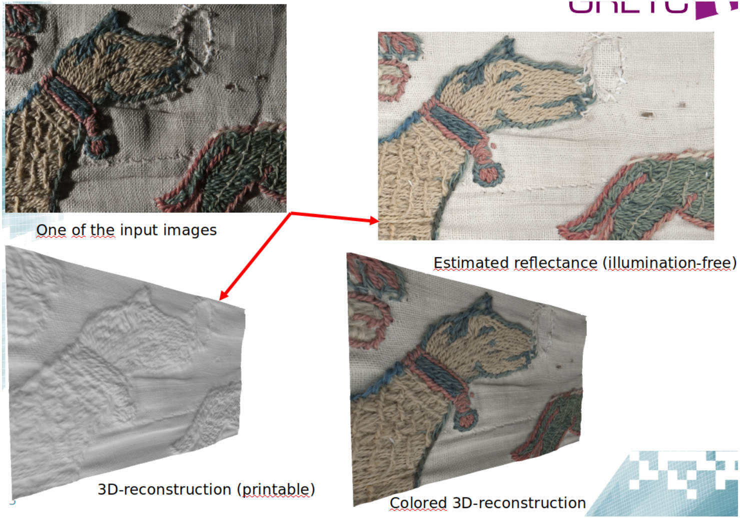

Ongoing Work on Optical Illusions

- With one image: ill-posed problem
- With many images: well-posed problem

- $\Rightarrow$ With two images: can we find a shape explaining any two images under two different lighting?

Monet

Van Gogh

Monet or Van Gogh ?

Thank you for your attention !

Brahimi, M. et al. (2019). "On the Well-Posedness of Uncalibrated Photometric Stereo under General Lighting". In: *HAL 02297643*.

Haefner, B., S. Peng, et al. (2019). "Photometric Depth Super-Resolution". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*.

Haefner, B., Y. Quéau, and D. Cremers (2019). "Photometric Segmentation: Simultaneous Photometric Stereo and Masking". In: *The IEEE International Conference on 3D Vision (3DV)*.

Haefner, B., Y. Quéau, et al. (2018). "Fight ill-posedness with ill-posedness: Single-shot variational depth super-resolution from shading". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Haefner, B., Z. Ye, et al. (2019). "On the Well-Posedness of Uncalibrated Photometric Stereo under General Lighting". In: *The IEEE International Conference on Computer Vision (ICCV)*.

Mélou, J. et al. (2019). "A splitting-based algorithm for multi-view stereopsis of textureless objects". In: *International Conference on Scale Space and Variational Methods in Computer Vision (SSVM)*.

Peng, S. et al. (2017). "Depth Super-Resolution Meets Uncalibrated Photometric Stereo". In: *The IEEE International Conference on Computer Vision (ICCV)*.

Quéau, Y. et al. (2017). "A Variational Approach to Shape-from-shading Under Natural Illumination". In: *Energy Minimization Methods for Computer Vision and Pattern Recognition (EMMCVPR)*.