High-Quality 3D Reconstruction from RGB-D Sensors

Computer Vision I: Variational Methods

Current Research

Robert Maier

robert.maier@in.tum.de https://vision.in.tum.de/members/maierr

January 25, 2018







- RGB-D Sensors & 3D Reconstruction
- Signed Distance Functions
- Real-Time Camera Tracking and 3D Reconstruction Using Signed Distance Functions (Bylow et al, RSS 2013)
- De-noising, Stabilizing and Completing 3D Reconstructions On-the-go using Plane Priors (Dzitsiuk et al, ICRA 2017)
- Intrinsic3D: High-Quality 3D Reconstruction by Joint Appearance and Geometry Optimization with Spatially-Varying Lighting (Maier et al, ICCV 2017)





- RGB-D Sensors & 3D Reconstruction
- Signed Distance Functions
- Real-Time Camera Tracking and 3D Reconstruction Using Signed Distance Functions (Bylow et al, RSS 2013)
- De-noising, Stabilizing and Completing 3D Reconstructions On-the-go using Plane Priors (Dzitsiuk et al, ICRA 2017)
- Intrinsic3D: High-Quality 3D Reconstruction by Joint Appearance and Geometry Optimization with Spatially-Varying Lighting (Maier et al, ICCV 2017)



RGB-D: color (RGB) + depth (metric!)
 @ 30 fps





- RGB-D: color (RGB) + depth (metric!)
 @ 30 fps
- Structured Light / Time-of-flight
- Low-cost!







Asus Xtion



- RGB-D: color (RGB) + depth (metric!)
 @ 30 fps
- Structured Light / Time-of-flight
- Low-cost!





Microsoft Kinect v1



Asus Xtion



Intel RealSense R200



Occipital Structure Sensor



- RGB-D: color (RGB) + depth (metric!) @ 30 fps
- Structured Light / Time-of-flight
- Low-cost!



Google Tango





Asus Xtion









Lenovo Phab 2 Pro

Asus ZenFone AR

RGB-D based 3D Reconstruction



 Task: given a stream of RGB-D frames of a real-world scene, compute its 3D shape that maximizes the geometric consistency

RGB-D based 3D Reconstruction



- Task: given a stream of RGB-D frames of a real-world scene, compute its 3D shape that maximizes the geometric consistency
- SLAM: Simultaneous Localization and Mapping (RGB-D-SLAM)

RGB-D based 3D Reconstruction



- Task: given a stream of RGB-D frames of a real-world scene, compute its 3D shape that maximizes the geometric consistency
- **SLAM**: Simultaneous Localization and Mapping (RGB-D-SLAM)
- Fusion of RGB-D frames in dense volumetric 3D representation (focus in this talk: Signed Distance Fields)



KinectFusion [Newcombe et al., ISMAR 2011]

• **Robotics**: real-time 3D reconstruction of largescale environments (e.g. autonomous drones)



Parrot AR drone

- **Robotics**: real-time 3D reconstruction of largescale environments (e.g. autonomous drones)
- Requirement of high-quality 3D content for Augmented Reality, Virtual Reality, ...



HTC Vive

Microsoft HoloLens



- **Robotics**: real-time 3D reconstruction of largescale environments (e.g. autonomous drones)
- Requirement of high-quality 3D content for Augmented Reality, Virtual Reality, ...
 - •Usually: manual modelling (e.g. Maya)



HTC Vive

Microsoft HoloLens



NVIDIA VR Funhouse

- **Robotics**: real-time 3D reconstruction of largescale environments (e.g. autonomous drones)
- Requirement of high-quality 3D content for Augmented Reality, Virtual Reality, ...
 - •Usually: manual modelling (e.g. Maya)
 - Wide availability of **commodity RGB-D sensors**: efficient methods for 3D reconstruction of real-word scenes



HTC Vive

Microsoft HoloLens



- **Robotics**: real-time 3D reconstruction of largescale environments (e.g. autonomous drones)
- Requirement of high-quality 3D content for Augmented Reality, Virtual Reality, ...
 - •Usually: manual modelling (e.g. Maya)
 - Wide availability of commodity RGB-D sensors: efficient methods for 3D reconstruction of real-word scenes
- Challenge: how to reconstruct high-quality 3D
 models from low-cost depth sensors?



HTC Vive

Microsoft HoloLens



NVIDIA VR Funhouse





- RGB-D Sensors & 3D Reconstruction
- Signed Distance Functions
- Real-Time Camera Tracking and 3D Reconstruction Using Signed Distance Functions (Bylow et al, RSS 2013)
- De-noising, Stabilizing and Completing 3D Reconstructions On-the-go using Plane Priors (Dzitsiuk et al, ICRA 2017)
- Intrinsic3D: High-Quality 3D Reconstruction by Joint Appearance and Geometry Optimization with Spatially-Varying Lighting (Maier et al, ICCV 2017)

Signed Distance Functions¹



Volumetric 3D model representation

• Voxel grid: dense (e.g. KinectFusion) or sparse (e.g. Voxel Hashing)



¹ "A volumetric method for building complex models from range images", Curless and Levoy, SIGGRAPH 1996

Signed Distance Functions¹



Volumetric 3D model representation

- Voxel grid: dense (e.g. KinectFusion) or sparse (e.g. Voxel Hashing)
- Each voxel stores:
 - Signed Distance Function (SDF): signed distance to closest surface
 - Color values
 - Weights

 $D(\mathbf{x}) < 0$ $D(\mathbf{x}) = 0$ $D(\mathbf{x}) > 0$





Fusion of depth maps

• Integrate depth maps into SDF with their estimated camera poses

- Integrate depth maps into SDF with their estimated camera poses
- Voxel updates using weighted average

$$\begin{array}{l} D \leftarrow \frac{WD+wd}{W+w} \\ C \leftarrow \frac{WC+wc}{W+w} \\ W \leftarrow W+w \end{array}$$





- Integrate depth maps into SDF with their estimated camera poses
- Voxel updates using weighted average

$$\begin{array}{l} D \leftarrow \frac{WD+wd}{W+w} \\ C \leftarrow \frac{WC+wc}{W+w} \\ W \leftarrow W+w \end{array}$$





- Integrate depth maps into SDF with their estimated camera poses
- Voxel updates using weighted average

$$\begin{array}{l} D \leftarrow \frac{WD+wd}{W+w} \\ C \leftarrow \frac{WC+wc}{W+w} \\ W \leftarrow W+w \end{array}$$





- Integrate depth maps into SDF with their estimated camera poses
- Voxel updates using weighted average

$$\begin{array}{l} D \leftarrow \frac{WD+wd}{W+w} \\ C \leftarrow \frac{WC+wc}{W+w} \\ W \leftarrow W+w \end{array}$$





Fusion of depth maps

- Integrate depth maps into SDF with their estimated camera poses
- Voxel updates using weighted average

$$\begin{array}{l} D \leftarrow \frac{WD+wd}{W+w} \\ C \leftarrow \frac{WC+wc}{W+w} \\ W \leftarrow W+w \end{array}$$

 Extract ISO-surface with Marching Cubes² (triangle mesh)

² "Marching cubes: A high resolution 3D surface construction algorithm", Lorensen and Cline, SIGGRAPH 1987





Real-Time Camera Tracking and 3D Reconstruction Using Signed Distance Functions

E. Bylow¹, J. Sturm², C. Kerl², F. Kahl¹, D. Cremers²

¹ Lund University



² Technical University of Munich

Robotics: Science and Systems (RSS) 2013, Berlin, Germany















29

System Pipeline













Camera tracking



- Estimate current camera pose from input RGB-D frame
- KinectFusion: synthetic depth map from SDF (raycasting) + ICP alignment

Camera tracking



- Estimate current camera pose from input RGB-D frame
- KinectFusion: synthetic depth map from SDF (raycasting) + ICP alignment
- Novel direct camera tracking against SDF: **direct minimization of error** between input depth map and SDF:



Camera tracking



- Estimate current camera pose from input RGB-D frame
- KinectFusion: synthetic depth map from SDF (raycasting) + ICP alignment
- Novel direct camera tracking against SDF: direct minimization of error between input depth map and SDF:





Extension: CopyMe3D

CopyMe3D: Scanning and Printing Persons in 3D (Sturm et al, GCPR 2013)

CopyMe3D: Scanning and Printing Persons in 3D

Jürgen Sturm, Erik Bylow, Fredrik Kahl, Daniel Cremers

German Conference on Pattern Recognition (GCPR) September 2013



Computer Vision Group Department of Computer Science Technical University of Munich



CopyMe3D Printed 3D figures












CopyMe3D2 Full-body scanning



De-noising, Stabilizing and Completing 3D Reconstructions On-the-go using Plane Priors

M. Dzitsiuk^{1,2}, J. Sturm², R. Maier¹, L. Ma¹, D. Cremers¹





² Technical University of Munich

International Conference on Robotics and Automation (ICRA), May 2017, Singapore



Real-time 3D reconstruction on a mobile device



Chisel: Real Time Large Scale 3D Reconstruction Onboard a Mobile Device using Spatially Hashed Signed Distance Fields [Klingensmith et al., 2015]





Motivation



Problems with real-time 3D reconstruction: noisy, incomplete, no segmentation



Motivation



Solution idea: detect and use planes

Approach

• Input: Signed Distance Field divided into voxel blocks





Approach

- Input: Signed Distance Field divided into voxel blocks
- 1. Find local plane candidates





- Approach
- Input: Signed Distance Field divided into voxel blocks
- 1. Find local plane candidates

• 2. Merge planes









Approach

- Input: Signed Distance Field divided into voxel blocks
- 1. Find local plane candidates

• 2. Merge planes

• 3. De-noising









Results: de-noising



Before





After



Results

Hole filling





Results

Hole filling



Segmentation



Classify reconstruction geometry:

- Floor or wall (area, angle with gravity)
- Object (mesh connected components)



Results: real-time 3D reconstruction on mobile device



Intrinsic3D: High-Quality 3D Reconstruction by Joint Appearance and Geometry Optimization with Spatially-Varying Lighting

R. Maier^{1,2}, K. Kim¹, D. Cremers², J. Kautz¹, M. Nießner^{2,3}



International Conference on Computer Vision (ICCV) October 2017, Venice, Italy









RGB-D based 3D Reconstruction

 Challenge: how to reconstruct high-quality 3D models with best-possible geometry and color from low-cost depth sensors?



RGB-D based 3D Reconstruction

- Challenge: how to reconstruct high-quality 3D models with best-possible geometry and color from low-cost depth sensors?
- Real-time, robust, fairly accurate geometric reconstructions



KinectFusion, 2011

"KinectFusion: Real-time Dense Surface Mapping and Tracking", Newcombe et al., ISMAR 2011.



RGB-D based 3D Reconstruction

- Challenge: how to reconstruct high-quality 3D models with best-possible geometry and color from low-cost depth sensors?
- Real-time, robust, fairly accurate geometric reconstructions



KinectFusion, 2011

"KinectFusion: Real-time Dense Surface Mapping and Tracking", Newcombe et al., ISMAR 2011.



DynamicFusion, 2015

"DynamicFusion: Reconstruction and Tracking of Non-rigid Scenes in Real-time", Newcombe et al., CVPR 2015.

RGB-D based 3D Reconstruction

- Challenge: how to reconstruct high-quality 3D models with best-possible geometry and color from low-cost depth sensors?
- Real-time, robust, fairly accurate geometric reconstructions



KinectFusion, 2011

"KinectFusion: Real-time Dense Surface Mapping and Tracking", Newcombe et al., ISMAR 2011.



DynamicFusion, 2015

"DynamicFusion: Reconstruction and Tracking of Non-rigid Scenes in Real-time", Newcombe et al., CVPR 2015.



BundleFusion, 2017

"BundleFusion: Real-time Globally Consistent 3D Reconstruction using On-the-fly Surface Re-integration", Dai et al., ToG 2017. 54



- Baseline RGB-D based 3D reconstruction framework
 - initial camera poses
 - sparse SDF reconstruction





- Baseline RGB-D based 3D reconstruction framework
 - initial camera poses
 - sparse SDF reconstruction
- Challenges:
 - (Slightly) inaccurate and over-smoothed geometry





- Baseline RGB-D based 3D reconstruction framework
 - initial camera poses
 - sparse SDF reconstruction
- Challenges:
 - (Slightly) inaccurate and over-smoothed geometry
 - Bad colors





- Baseline RGB-D based 3D reconstruction framework
 - initial camera poses
 - sparse SDF reconstruction
- Challenges:
 - (Slightly) inaccurate and over-smoothed geometry
 - Bad colors
 - Inaccurate camera pose estimation





- Baseline RGB-D based 3D reconstruction framework
 - initial camera poses
 - sparse SDF reconstruction
- Challenges:
 - (Slightly) inaccurate and over-smoothed geometry
 - Bad colors
 - Inaccurate camera pose estimation
 - Input data quality (e.g. motion blur, sensor noise)





- Baseline RGB-D based 3D reconstruction framework
 - initial camera poses
 - sparse SDF reconstruction
- Challenges:
 - (Slightly) inaccurate and over-smoothed geometry
 - Bad colors
 - Inaccurate camera pose estimation
 - Input data quality (e.g. motion blur, sensor noise)





- Baseline RGB-D based 3D reconstruction framework
 - initial camera poses
 - sparse SDF reconstruction
- Challenges:
 - (Slightly) inaccurate and over-smoothed geometry
 - Bad colors
 - Inaccurate camera pose estimation
 - Input data quality (e.g. motion blur, sensor noise)
- Goal: High-Quality Reconstruction of Geometry and Color









High-Quality Colors [Zhou2014]





Optimize camera poses and image deformations to optimally fit initial (maybe wrong) reconstruction

But: HQ images required, no geometry refinement involved

"Color Map Optimization for 3D Reconstruction with Consumer Depth Cameras", Zhou and Koltun, ToG 2014



High-Quality Colors [Zhou2014]





Optimize camera poses and image deformations to optimally fit initial (maybe wrong) reconstruction

But: HQ images required, no geometry refinement involved

"Color Map Optimization for 3D Reconstruction with Consumer Depth Cameras", Zhou and Koltun, ToG 2014

High-Quality Geometry [Zollhoefer2015]



Adjust camera poses in advance (bundle adjustment) to improve color

Use shading cues (RGB) to refine geometry (shading based refinement of surface & albedo)

But: RGB is fixed (no color refinement based on refined geometry)

"Shading-based Refinement on Volumetric Signed Distance Functions", Zollhoefer et al., ToG 2015



High-Quality Colors [Zhou2014]





Optimize camera poses and image deformations to optimally fit initial (maybe wrong) reconstruction

But: HQ images required, no geometry refinement involved

"Color Map Optimization for 3D Reconstruction with Consumer Depth Cameras", Zhou and Koltun, ToG 2014

High-Quality Geometry [Zollhoefer2015]



Adjust camera poses in advance (bundle adjustment) to improve color

Use shading cues (RGB) to refine geometry (shading based refinement of surface & albedo)

But: RGB is fixed (no color refinement based on refined geometry)

"Shading-based Refinement on Volumetric Signed Distance Functions", Zollhoefer et al., ToG 2015

Idea: **jointly optimize for geometry, albedo and image formation model** to simultaneously obtain high-quality geometry and appearance!







• Temporal view **sampling & filtering** techniques (input frames)





- Temporal view **sampling & filtering** techniques (input frames)
- Joint optimization of
 - **surface & albedo** (Signed Distance Field)
 - image formation model





- Temporal view sampling & filtering techniques (input frames)
- Joint optimization of
 - **surface & albedo** (Signed Distance Field)
 - image formation model





- Temporal view **sampling & filtering** techniques (input frames)
- Joint optimization of
 - **surface & albedo** (Signed Distance Field)
 - image formation model
- Lighting estimation using Spatially-Varying Spherical Harmonics (SVSH)





- Temporal view sampling & filtering techniques (input frames)
- Joint optimization of
 - **surface & albedo** (Signed Distance Field)
 - image formation model
- Lighting estimation using Spatially-Varying Spherical Harmonics (SVSH)
- **Optimized colors** (by-product)



Approach

Overview










Overview





Overview





Shading-based Refinement (Shape-from-Shading)

Overview



Shading-based Refinement (Shape-from-Shading)

Overview



Shading-based Refinement (Shape-from-Shading)

> Spatially-Varying Lighting Estimation

Overview





Shading-based Refinement (Shape-from-Shading)

> Spatially-Varying Lighting Estimation

Joint Appearance and Geometry Optimization

- surface
- albedo
- image formation model

High-Quality 3D

Reconstruction

Approach

Overview





Shading-based Refinement (Shape-from-Shading)

> Spatially-Varying Lighting Estimation

Joint Appearance and Geometry Optimization

- surface
- albedo
- image formation model





Shading-based Refinement (Shape-from-Shading)

Spatially-Varying Lighting Estimation

Joint Appearance and Geometry Optimization

- surface
- albedo
- image formation model

High-Quality 3D Reconstruction







Temporal view sampling / filtering

80

Approach Overview



Shading-based Refinement (Shape-from-Shading)

> Spatially-Varying Lighting Estimation

Joint Appearance and Geometry Optimization

- surface
- albedo
- image formation model

High-Quality 3D Reconstruction







Keyframe Selection



• Compute per-frame blur score (for color image)³





Frame 81

Frame 84

• Select frame with best score within a fixed size window as keyframe

³ "The blur effect: perception and estimation with a new no-reference perceptual blur metric", Crete et al., SPIE 2007.

Sampling / Filtering

Sampling of voxel observations

- Sample from selected **keyframes only**
- Collect observations for voxel in input views:

$$c_i^v = \mathcal{C}_i(\pi(\mathcal{T}_i^{-1}\boldsymbol{v}_{iso})).$$





Input keyframes

Sampling / Filtering

Sampling of voxel observations

- Sample from selected keyframes only
- Collect observations for voxel in input views:

$$c_i^v = \mathcal{C}_i(\pi(\mathcal{T}_i^{-1}\boldsymbol{v}_{\mathrm{iso}})).$$

Voxel center transformed and projected into input view





Sampling / Filtering

Sampling of voxel observations

- Sample from selected **keyframes only**
- Collect observations for voxel in input views:

$$c_i^v = \mathcal{C}_i(\pi(\mathcal{T}_i^{-1}\boldsymbol{v}_{\mathrm{iso}})).$$

Voxel center transformed and projected into input view

• Observation weights: view-dependent on normal and depth $\cos(\theta)$

$$w_i^{\boldsymbol{v}} = \frac{\cos(\theta)}{d^2}$$

 Filter observations: keep only best 5 observations by weight





Reconstruction



Overview







Double-hierarchical (coarse-to-fine: SDF Volume / RGB-D)

Shading-based Refinement (Shape-from-Shading)

> Spatially-Varying Lighting Estimation

Joint Appearance and Geometry Optimization

- surface
- albedo
- image formation model

High-Quality 3D Reconstruction







• Shading equation:

$$\mathbf{B}(\boldsymbol{v}) = \mathbf{a}(\boldsymbol{v}) \sum_{m=1}^{b^2} l_m H_m(\mathbf{n}(\boldsymbol{v})),$$



• Shading equation:





• Shading equation:

$$\mathbf{B}(\boldsymbol{v}) = \mathbf{a}(\boldsymbol{v}) \sum_{m=1}^{b^2} l_m H_m[\mathbf{n}(\boldsymbol{v})],$$





















- Shading-based refinement:
 - Intuition: high-frequency changes in surface geometry \rightarrow shading cues in input images





- Shading-based refinement:
 - Intuition: high-frequency changes in surface geometry \rightarrow shading cues in input images
 - Estimate **lighting** given **surface** and **albedo** (intrinsic material properties)





- Shading-based refinement:
 - Intuition: high-frequency changes in surface geometry \rightarrow shading cues in input images
 - Estimate **lighting** given **surface** and **albedo** (intrinsic material properties)
 - Estimate surface and albedo given the lighting: minimize difference between estimated shading and input luminance

Overview





Shading-based Refinement (Shape-from-Shading)

> Spatially-Varying Lighting Estimation

Joint Appearance and Geometry Optimization

- surface
- albedo
- image formation model









Spherical Harmonics (SH)

- Encode incident lighting for a given surface point
- Smooth for Lambertian surfaces



Spherical Harmonics (SH)

- Encode incident lighting for a given surface point
- Smooth for Lambertian surfaces
- SH Basis functions *H_m* parametrized by **unit normal** *n*

$$\mathbf{B}(oldsymbol{v}) = \mathbf{a}(oldsymbol{v}) \sum_{m=1}^{b^2} l_m H_m(\mathbf{n}(oldsymbol{v}))$$



Spherical Harmonics (SH)

- Encode incident lighting for a given surface point
- Smooth for Lambertian surfaces
- SH Basis functions H_m parametrized by unit normal n $\mathbf{B}(\boldsymbol{v}) = \mathbf{a}(\boldsymbol{v}) \sum_{m=1}^{b^2} l_m H_m(\mathbf{n}(\boldsymbol{v}))$
- Good approx. using only 9 SH basis functions (2nd order)



0

1

Spherical Harmonics (SH)

- Encode incident lighting for a given surface point
- Smooth for Lambertian surfaces
- SH Basis functions H_m parametrized by unit normal n $\mathbf{B}(\boldsymbol{v}) = \mathbf{a}(\boldsymbol{v}) \sum_{m=1}^{b^2} l_m H_m(\mathbf{n}(\boldsymbol{v}))$
- Good approx. using only 9 SH basis functions (2nd order)
- Estimate SH coefficients: $E_{\text{light}}(\mathbf{l}) = \sum_{\mathbf{v} \in \mathbf{D}_0} (B(\mathbf{v}) \mathbf{I}(\mathbf{v}))^2$





0

1



Spherical Harmonics (SH)

- Encode incident lighting for a given surface point
- Smooth for Lambertian surfaces
- SH Basis functions H_m parametrized by unit normal n $\mathbf{B}(\boldsymbol{v}) = \mathbf{a}(\boldsymbol{v}) \sum_{m=1}^{b^2} l_m H_m(\mathbf{n}(\boldsymbol{v}))$
- Good approx. using only 9 SH basis functions (2nd order)
- Estimate SH coefficients: $E_{\text{light}}(\mathbf{l}) = \sum_{\mathbf{v} \in \mathbf{D}_0} (B(\mathbf{v}) \mathbf{I}(\mathbf{v}))^2$
- Shortcoming: purely directional → cannot represent scene lighting for all surface points simultaneously





0

1





Subvolume Partitioning





Subvolume Partitioning

 Partition SDF volume into subvolumes



Subvolume Partitioning

- Partition SDF volume into subvolumes
- Estimate independent SH coefficients for each subvolume



Subvolume Partitioning

- Partition SDF volume into subvolumes
- Estimate independent SH coefficients for each subvolume
- Obtain **per-voxel SH coefficients** through tri-linear interpolation



Joint Optimization



Joint Optimization

• Estimate SVSH coefficients for all subvolumes jointly:

$$E_{\text{lighting}}(\boldsymbol{l}_1,\ldots,\boldsymbol{l}_K) = E_{\text{appearance}} + \lambda_{\text{diffuse}} E_{\text{diffuse}}.$$



Joint Optimization

• Estimate SVSH coefficients for all subvolumes jointly:

$$E_{\text{lighting}}(\boldsymbol{l}_1,\ldots,\boldsymbol{l}_K) = E_{\text{appearance}} + \lambda_{\text{diffuse}} E_{\text{diffuse}}.$$

Data term:

$$E_{\text{appearance}} = \sum_{\boldsymbol{v} \in \mathbf{D}_0} (\mathbf{B}(\boldsymbol{v}) - \mathbf{I}(\boldsymbol{v}))^2.$$

Similarity between estimated shading and input luminance



Joint Optimization

• Estimate SVSH coefficients for all subvolumes jointly:

$$E_{\text{lighting}}(\boldsymbol{l}_1,\ldots,\boldsymbol{l}_K) = E_{\text{appearance}} + \lambda_{\text{diffuse}} E_{\text{diffuse}}.$$

Data term:

$$E_{\text{appearance}} = \sum_{\boldsymbol{v} \in \mathbf{D}_0} (\mathbf{B}(\boldsymbol{v}) - \mathbf{I}(\boldsymbol{v}))^2.$$

Similarity between estimated shading and input luminance

Laplacian regularizer:

$$E_{\text{diffuse}} = \sum_{s \in \mathcal{S}} \sum_{r \in \mathcal{N}_s} (\boldsymbol{l}_s - \boldsymbol{l}_r)^2.$$

Smooth illumination changes

Overview





Spatially-Varying Lighting Estimation

Joint Appearance and Geometry Optimization

- surface
- albedo
- image formation model










Shading-based SDF optimization

• Joint optimization of geometry, albedo and image formation model (camera poses and camera intrinsics):

$$E_{\text{scene}}(\mathcal{X}) = \sum_{\boldsymbol{v} \in \tilde{\mathbf{D}}_0} \lambda_g E_g + \lambda_v E_v + \lambda_s E_s + \lambda_a E_a$$

with $\mathcal{X} = (\mathcal{T}, \tilde{\mathbf{D}}, \mathbf{a}, f_x, f_y, c_x, c_y, \kappa_1, \kappa_2, \rho_1)$



Shading-based SDF optimization

• Joint optimization of geometry, albedo and image formation model (camera poses and camera intrinsics):

$$E_{\text{scene}}(\mathcal{X}) = \sum_{\boldsymbol{v} \in \tilde{\mathbf{D}}_0} \lambda_g E_g + \lambda_v E_v + \lambda_s E_s + \lambda_a E_a$$

with $\mathcal{X} = (\mathcal{T}, \tilde{\mathbf{D}}, \mathbf{a}, f_x, f_y, c_x, c_y, \kappa_1, \kappa_2, \rho_1)$

Gradient-based shading constraint (data term)



Shading-based SDF optimization

• Joint optimization of geometry, albedo and image formation model (camera poses and camera intrinsics):

$$\begin{split} E_{\text{scene}}(\mathcal{X}) &= \sum_{\boldsymbol{v} \in \tilde{\mathbf{D}}_0} \lambda_g E_g + \lambda_v E_v + \lambda_s E_s + \lambda_a E_a \\ & \text{with } \mathcal{X} = (\mathcal{T}, \tilde{\mathbf{D}}, \mathbf{a}, f_x, f_y, c_x, c_y, \kappa_1, \kappa_2, \rho_1) \end{split}$$

Gradient-based shading constraint (data term)

Volumetric regularizer: smoothness in distance values (Laplacian)

$$E_v(\boldsymbol{v}) = (\Delta \tilde{\mathbf{D}}(\boldsymbol{v}))^2$$



Shading-based SDF optimization

• Joint optimization of geometry, albedo and image formation model (camera poses and camera intrinsics):

$$E_{\text{scene}}(\mathcal{X}) = \sum_{\boldsymbol{v}\in\tilde{\mathbf{D}}_0} \lambda_g E_g + \lambda_v E_v + \lambda_s E_s + \lambda_a E_a$$

with $\mathcal{X} = (\mathcal{T}, \tilde{\mathbf{D}}, \mathbf{a}, f_x, f_y, c_x, c_y, \kappa_1, \kappa_2, \rho_1)$

Gradient-based shading constraint (data term)

Volumetric regularizer: smoothness in distance values (Laplacian) Surface Stabilization constraint: stay close to initial distance values

$$E_s(\boldsymbol{v}) = (\tilde{\mathbf{D}}(\boldsymbol{v}) - \mathbf{D}(\boldsymbol{v}))^2$$



Shading-based SDF optimization

• Joint optimization of geometry, albedo and image formation model (camera poses and camera intrinsics):

$$E_{\text{scene}}(\mathcal{X}) = \sum_{\boldsymbol{v}\in\tilde{\mathbf{D}}_0} \lambda_g E_g + \lambda_v E_v + \lambda_s E_s + \lambda_a E_a$$

with $\mathcal{X} = (\mathcal{T}, \tilde{\mathbf{D}}, \mathbf{a}, f_x, f_y, c_x, c_y, \kappa_1, \kappa_2, \rho_1)$

Gradient-based shading constraint (data term)

Volumetric regularizer: smoothness in distance values (Laplacian) Surface Stabilization constraint: stay close to initial distance values Albedo regularizer: constrain albedo changes based on chromaticity (Laplacian)

$$E_a(\boldsymbol{v}) = \sum_{\boldsymbol{u} \in \mathcal{N}_{\boldsymbol{v}}} \phi(\boldsymbol{\Gamma}(\boldsymbol{v}) - \boldsymbol{\Gamma}(\boldsymbol{u})) \cdot (\mathbf{a}(\boldsymbol{v}) - \mathbf{a}(\boldsymbol{u}))^2$$



Shading Constraint (data term)

• Idea: maximize consistency between estimated voxel shading and sampled intensities in input luminance images (gradient for robustness)

$$E_g(\boldsymbol{v}) = \sum_{\mathcal{I}_i \in \mathcal{V}_{\text{best}}} w_i^{\boldsymbol{v}} \| \nabla \mathbf{B}(\boldsymbol{v}) - \nabla \mathcal{I}_i(\pi(v_i)) \|_2^2$$



Shading Constraint (data term)

 Idea: maximize consistency between estimated voxel shading and sampled intensities in input luminance images (gradient for robustness)

$$E_g(\boldsymbol{v}) = \sum_{\mathcal{I}_i \in \mathcal{V}_{\text{best}}} w_i^{\boldsymbol{v}} \| \nabla \mathbf{B}(\boldsymbol{v}) - \nabla \mathcal{I}_i(\pi(v_i)) \|_2^2$$

Best views for voxel and respective view-dependent weights



Shading Constraint (data term)

 Idea: maximize consistency between estimated voxel shading and sampled intensities in input luminance images (gradient for robustness)

$$E_g(\boldsymbol{v}) = \sum_{\mathcal{I}_i \in \mathcal{V}_{\text{best}}} w_i^{\boldsymbol{v}} \| \nabla \mathbf{B}(\boldsymbol{v}) - \nabla \mathcal{I}_i(\pi(v_i)) \|_2^2$$

Best views for voxel and respective view-dependent weights

Shading: allows for optimization of surface (through normal) and albedo



Shading Constraint (data term)

 Idea: maximize consistency between estimated voxel shading and sampled intensities in input luminance images (gradient for robustness)

$$E_g(\boldsymbol{v}) = \sum_{\mathcal{I}_i \in \mathcal{V}_{\text{best}}} w_i^{\boldsymbol{v}} \| \nabla \mathbf{B}(\boldsymbol{v}) - \nabla \mathcal{I}_i(\pi(v_i)) \|_2^2$$

Best views for voxel and respective view-dependent weights

- Shading: allows for optimization of surface (through normal) and albedo
- Voxel center transformed and projected into input view



Shading Constraint (data term)

• Idea: maximize consistency between estimated voxel shading and sampled intensities in input luminance images (gradient for robustness)

$$E_g(\boldsymbol{v}) = \sum_{\mathcal{I}_i \in \mathcal{V}_{\text{best}}} w_i^{\boldsymbol{v}} \| \nabla \mathbf{B}(\boldsymbol{v}) - \nabla \mathcal{I}_i(\pi(v_i)) \|_2^2$$

Best views for voxel and respective view-dependent weights

- Shading: allows for optimization of surface (through normal) and albedo
- Voxel center transformed and projected into input view
- Sampling: allows for optimization of camera poses and camera intrinsics

Recolorization

Optimal colors





Recolorization



Optimal colors

- Recompute voxel colors after optimization at each level
- Sampling
 - Sample from **keyframes only**
 - Collect, weight and filter observations

Recolorization

Optimal colors

- Recompute voxel colors after optimization at each level
- Sampling
 - Sample from **keyframes only**
 - Collect, weight and filter observations
- Weighted average $c_{\boldsymbol{v}}^* = \operatorname*{arg\,min}_{c_{\boldsymbol{v}}} \sum_{(c_i^{\boldsymbol{v}}, w_i^{\boldsymbol{v}}) \in \mathcal{O}_v} w_i^{\boldsymbol{v}} (c_{\boldsymbol{v}} c_i^{\boldsymbol{v}})^2.$





Frog (synthetic)





Frog (synthetic)





Frog (synthetic)





Frog (synthetic)

- Quantitative surface accuracy evaluation
- Color coding: absolute distances (ground truth)

Zollhöfer et al. 15





Frog (synthetic)

- Quantitative surface accuracy evaluation
- Color coding: absolute distances (ground truth)

Zollhöfer et al. 15



Ours





Frog (synthetic)

- Quantitative surface accuracy evaluation
- Color coding: absolute distances (ground truth)

Zollhöfer et al. 15



Ours



Mean absolute deviation:

- Ours: 0.222mm (std.dev. 0.269mm)
- Zollhöfer et al: 0.278mm (std.dev. 0.299mm)
 - \rightarrow 20.14% more accurate

Relief (geometry)

Input Color











Fusion



Zollhöfer et al. 15





Ours







Fountain (appearance)

Input Color









Fusion







Zollhöfer et al. 15







Ours







129



Lion

Input Color















Geometry (ours)



Fusion

Ours

Fusion

Ours

130

Tomb Statuary

Input Color









Geometry (ours)





Ours



Appearance (ours)



Ours







Luminance





Luminance



133





Luminance



Shading





Global SH



134





Luminance





Shading

 $\mathbf{B}_{ ext{diff}} = |\mathbf{B}(\boldsymbol{v}) - \mathbf{I}(\boldsymbol{v})|$



Global SH

SVSH

Albedo

Conclusion



High-Quality 3D Reconstruction of Geometry and Appearance

- Temporal view **sampling & filtering** techniques
- Spatially-Varying Lighting estimation
- Joint optimization of surface & albedo (SDF) and image formation model
- **Optimized texture** as by-product

Conclusion

High-Quality 3D Reconstruction of Geometry and Appearance

- Temporal view sampling & filtering techniques
- Spatially-Varying Lighting estimation
- Joint optimization of surface & albedo (SDF) and image formation model
- Optimized texture as by-product

Robert Maier

Technical University of Munich Computer Vision Group

robert.maier@in.tum.de https://vision.in.tum.de/members/maierr

Thank you!

Questions?

