

MIT Open Access Articles

Decomposing global light transport using time of flight imaging

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Di Wu, M. O Toole, A. Velten, A. Agrawal, and R. Raskar. "Decomposing global light transport using time of flight imaging." In 2012 IEEE Conference on Computer Vision and Pattern Recognition, 366-373. Institute of Electrical and Electronics Engineers, 2012.

As Published: <http://dx.doi.org/10.1109/CVPR.2012.6247697>

Publisher: Institute of Electrical and Electronics Engineers

Persistent URL: <http://hdl.handle.net/1721.1/80724>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



Decomposing Global Light Transport using Time of Flight Imaging

Di Wu^{1,2} Matthew O’Toole^{1,3} Andreas Velten^{1,4} Amit Agrawal⁵ Ramesh Raskar¹

¹MIT Media Lab

²Tsinghua University

³University of Toronto

⁴Morgridge Institute, University of Wisconsin-Madison

⁵Mitsubishi Electric Research Labs

Abstract

Global light transport is composed of direct and indirect components. In this paper, we take the first steps toward analyzing light transport using high temporal resolution information via time of flight (ToF) images. The time profile at each pixel encodes complex interactions between the incident light and the scene geometry with spatially-varying material properties. We exploit the time profile to decompose light transport into its constituent direct, subsurface scattering, and interreflection components.

We show that the time profile is well modelled using a Gaussian function for the direct and interreflection components, and a decaying exponential function for the subsurface scattering component. We use our direct, subsurface scattering, and interreflection separation algorithm for four computer vision applications: recovering projective depth maps, identifying subsurface scattering objects, measuring parameters of analytical subsurface scattering models, and performing edge detection using ToF images.

1. Introduction

Light transport analysis is a lasting and challenging area in computer vision and graphics. Existing methods for light transport analysis usually rely on structured light techniques that use active spatially and/or angularly varying illumination patterns. They enable a broad range of applications, including material acquisition [8, 11], light transport separation [13, 18], and robust geometry acquisition [5, 6, 7, 23]. On the other hand, ToF imaging is traditionally limited to geometry acquisition. Emerging femto-second photography techniques for multiple bounce acquisition enable new applications in computer vision, such as single view BRDF capture [12].

In this paper, we propose a separation method based on ToF imaging. Our approach is new in two aspects. First, we utilize ToF imaging to measure 3D light transport; second, we use the time profile of each pixel to decompose different

light transport effects in a very simple way. The key insight of this research is to exploit ultra-high time resolution in light transport by differentiating light transport effects with different time profiles. The time models of different light effects are analyzed and verified in this paper.

Contributions We show how to decompose ToF images into direct, subsurface scattering, and interreflection components for the following applications:

- (i) recovering projective depth from the direct component in the presence of global scattering;
- (ii) identifying and label different types of global illumination effects;
- (iii) measuring parameters of subsurface scattering materials from a single point of view;
- (iv) performing edge detection using spatial-temporal ToF information.

2. Related work

Global illumination effects, including interreflections, caustics, and subsurface scattering, are important visual features of real-world scenes. Structured light methods utilize global light transport when decomposing multi-bounce light transport into individual bounces [2, 18], extracting geometric information from second-bounce light transport [10], and compensating for global illumination [14, 20]. A method for separating high-frequency direct transport from low-frequency global transport requires as little as a single photo under structured lighting [13]. These direct/global separation techniques find uses in range acquisition applications for robustly capturing scene depth information in the presence of global illumination [5, 7, 23]. Specifically designed structured illumination patterns also allow for robust geometry acquisition that accounts for both interreflections and subsurface scattering light transport effects within a scene [6].

ToF imaging methods provide another means for analyzing global transport effects. ToF data used in LiDAR and

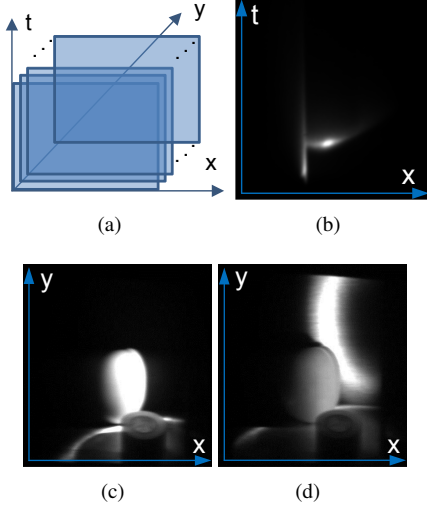


Figure 1. The ToF image format. (a) The ToF acquisition method consists of capturing multiple $x-t$ (streak) images, and combining the images into a single $x-y-t$ volume. (b) Example of one $x-t$ image. (c) $x-y$ image of the tomato tape scene (iv) (see Section 5 for details) when the direct component appears on the tomato at time t_1 . (d) $x-y$ image when the direct component appears on the back wall at time t_2 , with $t_2 > t_1$.

gated viewing systems exploit the direct component of the returned light to obtain a depth measurement [3, 4]. Coherent LiDAR relies on the wave interference of the returned light to recover depth [1, 9, 21, 22], so the technique can detect only the direct component, as global illumination poorly preserves coherence.

Recently, new cameras capture ToF images containing direct and global components for reconstructing geometry of occluded objects “around a corner” through the arrival time of multi-bounce light transport [15, 19]. In addition to geometry acquisition, ToF cameras allow for in-the-wild BRDF acquisition without any encircling equipment [12].

3. Light transport analysis using ToF

3.1. 3D ToF images

We store our ToF images as $x-y-t$ volumes, as shown in Figure 1a, where x and y correspond to pixel locations in a conventional image and t corresponds to time. An $x-y$ slice of the volume produces an image at a given time t , as shown in Figure 1c and 1d. The time profile $P(t)$ for a pixel (x, y) returns the pixel’s intensity as a function of time. We detail the capture process of these ToF images in Section 5. We use datasets from [16].

3.2. Composition of time profile

A photo is a composition of multiple light transport effects that occur within a scene. Similarly, the observed time profile $P(t)$ within our ToF image is a sum of time profiles,

including the direct time profile $D(t)$, subsurface scattering profile $S(t)$, and interreflection profile $I(t)$. Each of these individual time profiles have unique features.

Among all light paths that arrive at a particular pixel (x, y) at different times t , the direct light path, the path consisting of a single bounce, travels the shortest distance among all possible paths. It follows that the first impulse observed in the time profile $P(t)$ is the direct component.

Interreflections also produce individual impulses, though they occur later in time and with reduced intensity when compared to the direct impulse. For higher-order interreflections, individual peaks overlap, producing a smoothly decaying interreflection term.

Subsurface scattering materials exponentially attenuate the intensity of a ray of light with respect to its path length. By assuming path lengths directly correspond to the time of arrival, we expect intensity to decay exponentially as a function of time. In general, subsurface scattering is very complex and depends on several factors such as scene geometry, lighting conditions, and material properties. We empirically find that an exponentially decaying function closely approximates these subsurface scattering effects, as demonstrated in Figure 2b.

Decomposing a measured time profile into its individual global effects is a non-trivial problem in practice. Ideally, the measured time profile of a direct reflection or interreflections is an impulse of infinite short duration. Due to the limitation of ToF imaging, a spatio-temporal point spread function blurs our measured time profiles $P(t)$, as illustrated in Figure 2. Moreover, the high-dynamic range of the ToF images produces large amounts of saturation. A decomposition algorithm must robustly estimate the individual time profiles $D(t)$, $S(t)$, and $I(t)$ even when subjected to large blur kernels and amounts of saturation. A mixture model approach of decomposing the time profile into a mixture of Gaussians (as demonstrated in Figure 2a) and exponentially decaying functions is sensitive to blur and noise in the ToF images, and produces poor results; we take a more relaxed approach.

3.3. Direct and global separation

Our algorithm separates direct and global time profiles by localizing the direct component within the time profile $P(t)$, and extracting the direct component $D(t)$ by imposing smoothness constraints on the global component $G(t)$. Note that some pixels within a ToF image may receive no direct illumination. We thus reject time profiles without a direct component by thresholding the first peak intensity.

In Algorithm 1, steps 1 and 2 localize the arrival time t_{start} of the direct component. Step 3 finds the apex of the direct component at time t_{middle} . Step 4 returns the departure time t_{end} of the direct component, assuming the apex occurs at the center of the direct impulse. The pro-




Figure 2. (a) *Tomato tape scene*. A Gaussian closely approximates the time profile of a point containing only the direct component. (b) *Backlit tomato scene*. The time profile of a point lit by subsurface scattered light decays exponentially over time. A line closely approximates the log of the time profile. (c) *Mounted tomato scene*. A time profile containing both direct and subsurface scattering light transport is a sum of a Gaussian and an exponentially decaying function.

Algorithm 1 Direct/global separation

Require: the time profile $P(t)$

1. Compute the derivative of the time profile $P'(t)$
 2. Find the start of the direct component by solving $t_{start} = \arg \min_t P'(t) > \alpha$, for some tolerance α (typical value of α : $0.05 \max_t P'(t)$)
 3. Find the apex of the direct component by solving $t_{middle} = \arg \min_t P'(t) < \beta$ subject to $t > t_{start}$, for some tolerance β (typical value of β : 0.0001)
 4. Compute the end of the direct component $t_{end} = t_{start} + 2(t_{middle} - t_{start})$
 5. Smoothly interpolate time profile values between t_{start} and t_{end} using C^0 and C^1 continuity to generate the global component $G(t)$
 6. Extract direct component $D(t) = P(t) - G(t)$
-

file values $P(t)$ for $t \notin [t_{start}, t_{end}]$ are global component values; step 5 predicts the unknown global values $G(t)$ for $t \in [t_{start}, t_{end}]$ through interpolation. Lastly, step 6 subtracts the global component from the time profile to produce the direct component, as in Figure 3 and 4a.

3.4. Subsurface scattering and interreflection

The ToF images provide more information than simply the direct and global components; they are a composition of multiple global light transport effects, including interreflections and subsurface scattering. Our interreflection/subsurface scattering separation algorithm takes as input the global time profile $G(t)$ and the departure time of the direct component t_{direct} . Our separation algorithm localizes and extracts the interreflection component from the global time profile.

Algorithm 2 Interreflection/subsurface scattering separation

Require: global time profile $G(t)$ and end of direct component t_{direct}

1. Compute the derivative of the time profile $G'(t)$
 2. Find the start of the interreflection component by solving $t_{start} = \arg \min_t G'(t) > \alpha_G$ subject to $t > t_{direct}$, for some tolerance α_G (typical value of α_G : $0.05 \max_t P'(t)$)
 3. Find the apex of the interreflection component by solving $t_{middle} = \arg \min_t G'(t) < \beta_G$ subject to $t > t_{start}$, for some tolerance β_G (typical value of β_G : 0.0001)
 4. Compute the end of the interreflection component $t_{end} = t_{start} + 2(t_{middle} - t_{start})$
 5. Smoothly interpolate time profile values between t_{start} and t_{end} using C^0 and C^1 continuity to generate the subsurface scattering component $S(t)$
 6. Extract interreflection component $I(t) = G(t) - S(t)$
-

Algorithm 2 finds the start of a new impulse, identifies the apex of the impulse, computes the end of the impulse, and interpolates time profile values between the start and end points. Figure 4b decomposes the global time profile from Figure 4a into its interreflection and subsurface scattering components.

Both algorithms rely on the ability to detect individual impulses. High-order bounces have diminished intensity, and lose temporal locality (appears less impulse-like). As a result, our approach only reliably separates second-bounce interreflection from the global time profile.

4. Exploiting light transport components

Decomposing a single ToF image into its direct, subsurface scattering, and interreflection components has many applications in computer vision. Our paper focuses on four such applications: recovering projective depth from direct component, labelling image regions to identify subsurface scattering objects, measuring scattering coefficients, and edge detection.

Projective depth from ToF Monostatic LiDAR systems use a colocated light source and sensor to recover geometry using ToF. These devices capture depth by measuring the time elapsed between emitting a pulse of light onto a point in a scene, and the backscattered light returning to the sensor. LiDAR systems typically sweep the light across an area




Figure 3. Our algorithm takes a pixel’s time profile (black), computes its gradient (green), identifies the direct component region (pink), and decomposes the original time profile (black) into direct (red) and global (blue) components. The illustrated time profiles represent: (a) the direct component, (b) direct and global, (c) direct and subsurface scattering (with saturation), and (d) direct and subsurface scattering (without saturation). Time profiles (a-d) correspond respectively to points 1-4 in Figure 8.




Figure 4. The time profile corresponds to point 5 in Figure 8c. (a) Algorithm 1 decomposes the profile into its direct and global components. (b) Algorithm 2 further decomposes the global time profile into the interreflection and subsurface scattering components.

in the scene to recover a full depth map. Our ToF imaging device for capturing depth in a single image is similar to a bistatic LiDAR system [17], in that it requires no colocation between sensor and light source. The laser illuminates an unknown point in the scene and produces a virtual light source (see Section 5). This virtual light source illuminates the scene, and the resulting direct component provides sufficient information to recover a projective depth map of the scene.

The laser intersects a (unknown) point $L = (L_x, L_y, L_z)$ in world space at a (unknown) time t_0 in the scene, with the camera centered at point $C = (0, 0, 0)$. If the camera receives the first bounce of light at time t at camera pixel (x, y) , the direct reflection point $P = (P_x, P_y, P_z)$ must satisfy the following equation:

$$\|P - L\|_2 + \|P - C\|_2 = k(t - t_0) \quad (1)$$

Note k is a constant that maps units in time to units in space. The intersection of the ellipsoid in Equation 1 with a ray emitted from pixel (x, y) produces a unique depth map value, representing the depth value for camera pixel (x, y) .

In practice, we calibrate the system to recover L , t_0 , and k . Our calibration approach fits these parameters to multiple points on the surface, each containing the same depth value. Once these constants are known, our depth recovery algorithm solves for the unknown depth value given a time

t for each pixel, as Figure 5¹.

Labelling subsurface scattering regions The time profile for subsurface and volumetric scattering objects after receiving an initial light pulse has a unique appearance. When light hits an object in a scene, either transmission, reflection, or absorption occurs. Certain objects with subsurface scattering properties will scatter the transmitted light internally and slowly release the light over time.

To distinguish subsurface scattering objects from other objects within a scene, our algorithm labels subsurface scattering regions by evaluating the following equation for each pixel (x, y) :

$$T \geq Ratio(x, y) = \frac{\int_{t=0}^T P(t)_{x,y}}{\max_{t=0}^T P(t)_{x,y}} \geq 1 \quad (2)$$

When the function returns a value near 1, the energy concentrates to a single instant in time within the time profile, indicating the corresponding point has no subsurface scattering properties. A larger $Ratio(x, y)$ value corresponds to distributed energy within a time profile, a characteristic of subsurface scattering points.

Our algorithm for identifying subsurface scattering materials relies on thresholding the $Ratio(x, y)$ function at each pixel. When the function returns a value larger than a particular threshold T (0.35 for Figure 6), the algorithm identifies the pixel as a point with subsurface scattering.

¹ Figure 5d shows the projective depth of simulation ToF data for “corner” scene (We modify PBRT to simulate ToF of scenes with interreflections). The number of unknown variables is 5: (L_x, L_y, L_z, t_0, k) . We choose pixels lie on a horizontal line on the bottom plane, solving the above parameters based on Equation 1, using local search method *fminsearch* in Matlab. However, giving selected pixel constraints, it is hard to optimize variables up to 5. As in Figure 5d, the projective depth result has some error. In real world data, it is ever harder to pick pixels with same depth even for planar object, as the surface might be slightly tilted. Therefore, we reduce the unknown variables by setting $L = (1, 0, 0)$ and recover t_0 and k by choosing a set of pixels with the same depth value. (In Figure 5c, we choose the pixels along a vertical line on the back wall.) The projective depth corresponds to another coordinate system transformed based on constrain $L = (L_x, L_y, L_z)$ in the original coordinate system to be $L = (1, 0, 0)$ in the new one. In Figure 5c, the depth value preserve the plane structure, observed from another point of view.




Figure 5. Projective depth computation from direct component. (a) The points on a plane in the scene with the same ToF are distributed on an ellipse. (b) Direct peak time map of the back and bottom plane in the corner scene. (c) The corresponding projective depth map. (d) The projective depth of simulated ToF data. From left to right: scene illustration; direct component intensity map, the right wall has no direct component for serving as virtual light source; the tof of time profile peaks, we do not consider right wall which does not has direct component; projective depth from tof of direct component. The color corresponds to values in color bar in Figure 6a.




Figure 6. Labelling subsurface scattering for the mounted tomato scene (ii). (a) The image is given by computing the ratio per pixel using Equation 2. Note that higher values correspond to more subsurface scattering. (b) The image after thresholding. Cyan identifies subsurface scattering regions, magenta indicates no subsurface scattering, and black corresponds to background regions.

Measuring scattering model parameters An important application of ToF images is material acquisition, and in particular, modelling volumetric scattering. When a short pulse of light illuminates a volumetric scattering medium, the time profile shows an exponential decay in the measured irradiance that lasts from a few picoseconds to several nanoseconds. The time profile provides sufficient information to accurately model light scattering within a homogeneous medium.

Our model for volumetric scattering mediums is an exponential $\alpha e^{\gamma t} = e^{\log(\alpha) + \gamma t}$, where γ is the scattering coefficient and α is the incident illumination. Our algorithm

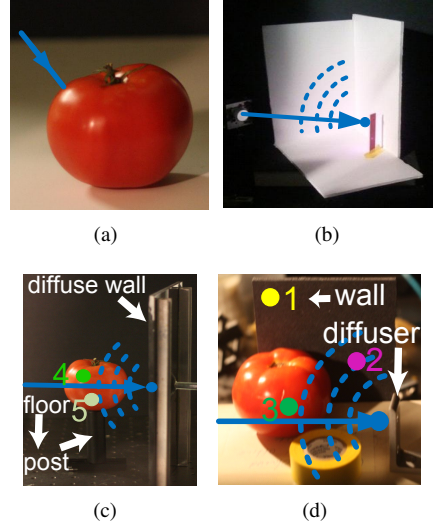


Figure 8. Images of setup for (a) backlit tomato scene, (b) the corner scene, (c) mounted tomato scene, (d) tomato tape scene. The blue arrows simulate the incident laser beam, each producing a virtual point light source within the scene. The separation plots for point 1-4 is in Figure 3, and point 5 is in Figure 4.

for capturing the scattering coefficient requires two steps. First, the algorithm computes $\log(P(t))$. Second, a line of the form $\log(\alpha) + \gamma t$ is fit to the logarithmic time profile $\log(P(t))$ to produce the scattering coefficient. Given the scattering coefficient, the time profile of the volumetric scattering medium is a function of the incident illumination.

Edge detection using ToF The local spatial properties of the time profile can improve edge detection methods. It is difficult to distinguish between a depth edge and a shadow edge. The time profile of a depth edge has two distinct direct time values, whereas shadow edges have at most one distinct peak. By analyzing the time profile of a set of pixels, our algorithms distinguish between the two edges.

To identify a depth edge, we analyze the time profiles for points within a 7×7 sliding window, as illustrated in Figure 7. For each time profile within the window, our algorithm computes the position and intensity of the maximum, and evaluates the variance of these quantities within the window, denoted as V_{apex} and V_{time} respectively. The higher value of $V_{time} \times V_{apex}$, the higher possibility that it belongs to a depth edge.

5. Experiment data

The acquisition system for capturing the data is described in separate work [16]. Please refer to [16] for more original moives.

We illuminate our scene with a *virtual point light source*, a diffuse point within the scene directly lit by the laser. We




Figure 7. The time profiles correspond to points in the following patches: (a) patch on tomato, (b) patch on perimeter of tomato, (c) patch on edge of shadow, (d) patch crossing another shadow edge.

intentionally occlude the virtual point light source within the scene; otherwise, the brightness of the virtual light source produces severe lens flare artifacts in the ToF image. The time resolution for the following dataset is 4, 4, 3, 4 picosecond respectively.

(i) Backlit Tomato: pure subsurface scattering The laser illuminates the backside of a tomato, as shown in Figure 8a. The dominant transport effect within the tomato scene is subsurface scattering.

(ii) Mounted Tomato: subsurface and simple multi-bounce scattering Figure 8c contains a diffuse wall and a tomato mounted on a post. The laser produces a virtual point light source on the diffuse wall, which then illuminates the tomato. This scene contains primarily direct, subsurface scattering and interreflection light transport effects.

(iii) Corner: complex multibounce scattering Figure 8b consists of three diffuse planes arranged to form a concave corner. The laser illuminates an occluded point on the wall. Because of the large number of near discrete bounces occurring in this setup, decomposing the time profile into its individual components is challenging.

(iv) Tomato and Tape: complex subsurface and multi-bounce scattering Figure 8d combines several light transport effects: direct, interreflection, and subsurface scattering. The scene consists of a roll of tape, a tomato and a wall. The laser creates a virtual light source on a diffuser. Complex interreflections occur in the roll of tape, and subsurface scattering effects appear on the tomato.

6. Results and performance validation

Separation results We perform direct and global separation for the mounted tomato scene (ii), corner scene (iii), and tomato and tape scene (iv). We show several $x-y$ slices of the separated scenes in Figures 9, 10a, and 10b respectively. In general, we observe the direct component in the earlier frames, whereas subsurface and multibounce scattering dominate the later frames.

We capture a ToF image for the mounted tomato scene (ii), and separate the image into its direct, subsurface scat-


tering, and interreflection components. Figure 9a shows direct/global separation results on an $x-y$ image at a time when the direct component intersects the tomato. The tomato produces strong subsurface scattering effects as illustrated by the rightmost image. Figure 9b shows subsurface scattering/interreflection separation on an $x-y$ image at a later time when no direct component exists. The interreflection originates from light transport between the floor and the tomato.

The corner scene (iii) in Figure 10a contains interreflections between the three walls forming a concave corner. The time profile for pixels near the corner of the wall behaves similarly to a subsurface scattering profile that decays slowly over time. Our separation algorithm identifies the direct and global components, although cannot distinguish between interreflection and subsurface scattering in this example.


We also analyze the tomato and tape scene (iv) in Figure 10b. The figure shows the direct and global components at two different time frames. The tomato produces subsurface scattering effects, and we observe interreflections in the inner part of the tape roll and between the wall and tomato. The light from the subsurface scattering glow of the tomato illuminates the rest of the scene over a prolonged period of time. The algorithm once again separates the image into its direct and global components. Please refer to the supplementary video for additional results.

Characterization of subsurface scattering The subsurface scattering properties associated with different concentrations of milk in water can be characterized as a decaying exponential. We analyze subsurface scattering by first filling a tank with 3.6 liters of water, and gradually adding milk (2% reduced fat) to the water to increase the scattering properties of the medium. Our camera captures an $x-t$ image for 0, 5, 10, 20, 40, 80, and 160 ml concentrations of milk within the tank. The milk contains emulgated fat bubbles that serve as scattering centers. Observing a point in the scatterer, the time profile of the returned light allows us to distinguish between the different milk emulsions.

A straight line closely approximates the logarithm of the



(a) Direct/global decomposition




(b) Subsurface scattering/interreflection decomposition

Figure 9. Global light transport decomposition of the mounted tomato scene (ii) at two different times. (a) Direct/global separation. Algorithm 1 takes a ToF image (left) and decomposes the image into its direct (middle) and global (right) components. (b) Algorithm 2 takes a global image (left) and returns two images: the subsurface scattering component (middle) and the interreflection component (right). Light bouncing between the floor and the tomato produces the interreflection component. Note this global image occurs at a time when the direct component has passed.




(a) Direct/global decomposition: corner scene (iii)



(b) Direct/global decomposition: tomato tape scene (iv)


Figure 10. Direct/global separation results at two different times for each of (a) corner scene and (b) tomato tape scene. Our algorithm separates the original x - y images (left) into its direct component (middle) and global component (right).

observed time profile $\log(P(t))$ for different concentrations of milk in Figure 11b. Figure 11c illustrates the exponen-



(a)

(b)



(c)

(d)




Figure 11. Subsurface scattering in milky water. (a) The time profiles of the same point in 7 different concentrations of milky water. (b) Lines are fit to the log plot of the time profiles. (c) Log-log plot for slope of lines versus concentration. (d) 4 raw streak images captured for milky water scene.




Figure 12. Differentiation of depth edges and shadow edges. (a) Edge detection using Canny edge detector for the tomato tape scene. (b) V_{apex} map, (c) V_{time} map, (d) $V_{time} \times V_{apex}$ map.

tial relationship between the decay constant and the milk concentration.

Validation of edge detection Edge detection algorithms aim to find edges through variations in depth, texture, and lighting. Different vision applications require different types of edges. Typical image edge detection algorithms extract edges due to brightness variations, which are coupled for all depth, texture, and light edges. In our ToF images, the depth edge points and shadow edge points have different local properties. We compare our edge detection based on 3D light transport with the Canny algorithm acting in the time integrated ToF image in Figure 12 for the tomato tape scene (iv). The results show that our ToF approach to edge detection generates cleaner depth edges and shadow edges compared to the Canny algorithm.

Limitations of ToF imaging The fast attenuation of the direct component in the scene produces scenes with a high dynamic range. Because our 12-bit camera sensor cannot capture the full range of intensity values, we allow some

saturation in the ToF images. The time resolution dictates the ability to recover direct and multibounce components, especially in the presence of many interreflections. A one hour capture time for a 600 slices ToF image limits our measurements to static scenes.

Limitations of the separation method When considering separation of discrete bounces from subsurface scattering, it is important to note that these are not two completely distinct phenomena. Most surfaces contain at least some subsurface scattering. A second or third bounce striking a subsurface scatterer causes subsurface scattering just like the direct component does. Light emerging from a subsurface scatterer can afterwards undergo discrete bounces in the scene. This light is not distinguishable from immediate subsurface scattering light using our methods.

7. Conclusion

This work represents the first steps toward using ToF information in light transport analysis. We decompose and label scene elements according to the type of global light transport, measure coefficients of volumetric scattering and produce depth maps from direct light transport.

In the future, combining ToF imaging with structured illumination could more accurately identify transport properties. Compact and low cost ultra-fast systems are becoming available due to massive commercial interest in optical communication and optical computing. We believe the addition of the time dimension in computer vision, computational photography, and active illumination techniques will enable a new range of applications.

8. Acknowledgements

The work of the MIT affiliated coauthors was funded by the Media Lab Consortium Members, DARPA through the DARPA YFA grant, and the Institute for Soldier Nanotechnologies and U.S. Army Research Office under contract W911NF-07-D-0004.

References

- [1] N. Abramson. Light-in-flight recording by holography. *Opt. Lett.*, 3(4):121–123, 1978.
- [2] J. Bai, M. Chandraker, T. Ng, and R. Ramamoorthi. A dual theory of inverse and forward light transport. *Proc. ECCV*, pages 294–307, 2010.
- [3] G. Brooker. Introduction to sensors for ranging and imaging. page 717, 2009.
- [4] J. Busck and H. Heiselberg. Gated viewing and high-accuracy three-dimensional laser radar. *Applied optics*, 43(24):4705–4710, 2004.
- [5] T. Chen, H. P. A. Lensch, C. Fuchs, H. peter Seidel, and M. informatik. Polarization and phaseshifting for 3d scanning of translucent objects. In *Proc. CVPR*, pages 1–8, 2007.
- [6] M. Gupta, A. Agrawal, A. Veeraraghavan, and S. Narasimhan. Structured light 3d scanning in the presence of global illumination. In *Proc. CVPR*, pages 713–720, 2011.
- [7] M. Gupta, Y. Tian, S. Narasimhan, and L. Zhang. (de) focusing on global light transport for active scene recovery. *Proc. CVPR*, pages 2969–2976, 2009.
- [8] M. Holroyd, J. Lawrence, and T. Zickler. A coaxial optical scanner for synchronous acquisition of 3d geometry and surface reflectance. *ACM Trans. Graph.*, 2010.
- [9] D. Huang, E. Swanson, C. Lin, J. Schuman, W. Stinson, W. Chang, M. Hee, T. Flotte, K. Gregory, and C. Puliafito. Optical coherence tomography. *Science*, 254(5035):1178–1181, 1991.
- [10] S. Liu, T. Ng, and Y. Matsushita. Shape from second-bounce of light transport. *Proc. ECCV*, pages 280–293, 2010.
- [11] Y. Mukaigawa, Y. Yagi, and R. Raskar. Analysis of light transport in scattering media. In *Proc. CVPR*, pages 153 – 160, 2010.
- [12] N. Naik, S. Zhao, A. Velten, R. Raskar, and K. Bala. Single view reflectance capture using multiplexed scattering and time-of-flight imaging. In *ACM Trans. Graph.*, 2011.
- [13] S. Nayar, G. Krishnan, M. Grossberg, and R. Raskar. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Trans. Graph.*, 25(3):935–944, 2006.
- [14] T.-T. Ng, R. S. Pahwa, J. Bai, K.-H. Tan, and R. amamoorthi. From the rendering equation to stratified light transport inversion. *IJCV*, 2011.
- [15] R. Pandharkar, A. Velten, A. Bardagjy, E. Lawson, M. Bawendi, and R. Raskar. Estimating motion and size of moving non-line-of-sight objects in cluttered environments. In *Proc. CVPR*, pages 265–272, 2011.
- [16] R. Raskar. Visualizing light in motion at a trillion frames per second, 2011. raskar.info/trillionfps/.
- [17] E. Repasi, P. Lutzmann, O. S. M. Elmqvist, B. Ghler, and G. Anstett. Advanced short-wavelength infrared range-gated imaging for ground applications in monostatic and bistatic configurations. *Applied Optics*, 48:5956–5969, 2009.
- [18] S. Seitz, Y. Matsushita, and K. Kutulakos. A theory of inverse light transport. In *Proc. ICCV*, pages 1440–1447, 2005.
- [19] A. Velten, T. Willwacher, O. Gupta, A. Veeraraghavan, M. G. Bawendi, and R. Raskar. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Communication*, 3:745–758, 2012.
- [20] G. Wetzstein and O. Bimber. Radiometric compensation through inverse light transport. In *Pacific Graphics*, pages 391–399, 2007.
- [21] J. C. Wyant. White light interferometry. In *SPIE*, volume 4737, pages 98–107, 2002.
- [22] H. Xia and C. Zhang. Ultrafast ranging lidar based on real-time Fourier transformation. *Optics Letters*, 34:2108–2110, 2009.
- [23] L. Zhang and S. Nayar. Projection defocus analysis for scene capture and image display. In *ACM Trans. Graph.*, volume 25, pages 907–915, 2006.