5d time-light transport matrix: What can we reason about scene properties?

Ramesh Raskar, James Davis

Abstract

Can time-of-flight cameras be used for performing multi-path analysis of global light transport? Traditional cameras estimate intensity per pixel, I(x,y,). In this position paper, we argue that emerging technologies will soon support cameras with a temporal-profile at each pixel at picosecond resolution, allowing us an ultra-high speed capture of the time-image, I(x,y,t). This time-image contains the time profile of irradiance at a pixel. Importantly, the speed of light is relevant, and the transient properties of light transport come into play. We propose that measurements of the transient properties of light will enable new methods for reasoning about scene content, and we show a few example scenarios in which transient reasoning exposes scene properties that are beyond the reach of traditional machine vision.

1. Introduction

Machine vision research has long used cameras to observe and interpret scenes. The digital models recovered have included such properties as 3D geometry, scene lighting, and surface reflectance, and found applications in robotics, industrial sensing, security, and user interfaces. New sensors and methods for interpreting scenes will clearly be of benefit in many application areas. This paper introduces a framework for reasoning about *transient light transport* and shows that it can allow new properties of scenes to be observed and interpreted.

Traditional cameras sense a limited 2D projection of the complete light transport in a scene. Many distinct scenes result in identical projections, and thus identical recorded intensity values on the imager. It is very difficult to use traditional imaging to estimate properties such as the depth of a mirror, the position of a sheet of glass, or the overall scale of a scene because these properties are not directly observable in the RGB values reported by a traditional camera.

Specialized sensors such as laser scanners and LIDAR systems have been built to directly and accurately estimate scene depth. These have been very successful in some circumstances, but they work well only for certain types of surface reflectances, and do little to help estimate other global properties such as relationship between scene patches. In addition they are used in restrictive configurations by carefully placing emitters near the receivers. What is needed is a generalized sensor which fundamentally records a greater portion of the light transport in a scene. This sensor could then be used to design new algorithms and specialized sensing methods.

Impulse versus steady-state response: Critical to our work is the distinction between steady-state and transient light transport. Steady-state transport corresponds to the familiar case in computer vision or graphics, in which the speed of light is conventionally assumed to be infinite, taking no time to cross any distance. We interpret the value of a pixel as the amount of light received at that pixel over the exposure duration but the irradiant flux (rate of incident photons) is constant and not a function of time. Videos may be interpreted as a sequence of images of different but static worlds because the exposure time of each frame is sufficiently long. Fundamentally, steady-state light transport describes an amount of energy, a number of photons, or the irradiance at a pixel.

In transient light transport, we assume that the speed of light is some finite value. As light scatters around a scene, it takes different paths, and longer paths take a longer time to traverse. Even a single pulse of light can evolve into a complex pattern in time, as shown in Figure 1. We call the pattern of light resulting from an impulse of light the *transient photometric response function* (TPRF). Fundamentally, transient transport describes power, a rate of incoming photons, or irradiant flux at a pixel, which, importantly, is measured as a function of time.

Traditional video cameras sample light very slowly compared to the time scale at which the transient properties of light come into play. Even a microsecond exposure time is long enough for a light impulse to fully traverse all possible paths and be fully integrated by the imager for common scenes. This work proposes an imaging model which samples light on the picosecond scale, which is equivalent to light travel on the order of millimeters. Light travels 0.3 millimeter in one picosecond. At this scale, it is possible to reason about the individual paths light takes within a scene. This allows direct observation of properties such as distance, and the difference between primary and secondary reflections.

Cameras beyond 2D: Current sensors sample discrete dimensions of the light transport function. If we consider all the light which arrives in the neighborhood of our "camera", it is a function of spatial location, incoming angle, and time. A typical photo-detector simply integrates across all of these dimensions to arrive at a single intensity reading. Traditional cameras use multiple pixels to explicitly sample incoming angles, from a single spatial location. Importantly, traditional cameras integrate light across time into a single value at each pixel. That is, cameras sample in some dimensions, but integrate light in others. A Lightfield camera samples both the incoming angle of light and spatial location, resulting in increased information with which to reason about the scene. In this work, we seek to sample time, so that we can measure and reason about transient properties of the light transport in a scene that depend on light taking a measurable amount of time to cross extended distances.

Contribution: The primary contribution of this work is conceptual rather than experimental. Our goal is to influence the direction of future research both in terms of sensor design and algorithms for scene understanding. We propose a theoretical foundation for sensing and reasoning using transient light transport, as well as example scenarios in which transient reasoning exposes scene properties that are beyond the reach of traditional machine vision.

2. Related Work

SONAR: SONAR (SOund Navigation And Ranging), is a technique that uses sound propagation in a medium such as air or water to detect and locate remote objects. The speed of sound is six orders of magnitude slower than the speed of light, and therefore easier to detect. Nevertheless work in SONAR has produced intricate models of the effects of many surfaces with complicated scattering properties [Russell96].

LIDAR: LIDAR (LIght Detection And Ranging), is roughly the light analog of SONAR. Short pulses of laser light from an emitter can be used to trigger time-delayed reflections from remote objects in a scene [Kamermann 1993]. Compared with SONAR, most LIDAR models are extremely simple. For example, flash-LIDAR systems intended for the low-cost consumer market frequently assume a single surface and measure the phase of a reflected pulse or sinusoidal signal [Canesta, MESA, 3DV, PMD, Yahav00, Bamji01, Iddan01, Lange01, Gvili03]. More sophisticated reasoning is often limited to improving phase estimation, for instance by reducing the effects of ambient light [Miyagawa97, Schroeder99, Kawakita00, Davis03, Davis04], or by simulating the expected shape of the reflected signal [Jutzi06].

Some LIDAR systems do measure the transient photometric response function explicitly. For example, the depth of both forest canopy foliage and the ground can be determined independently by separately detecting multiple peaks in the sensor response [Blair99, Hofton00] and surface discontinuities can be detected through waveform analysis [Vandapel04]. However, all of these methods reason locally about the sensor response in a single direction, rather than about the global scene structure.

This paper proposes that complex global reasoning about scene contents is possible given a measured TPRF.

Time gated imaging: Time gated imaging allows a reflected pulse of light to be integrated over extremely short windows, effectively capturing I(x, y, t_{delta}). Multiple captures while incrementing the time window, t_{delta} , allow

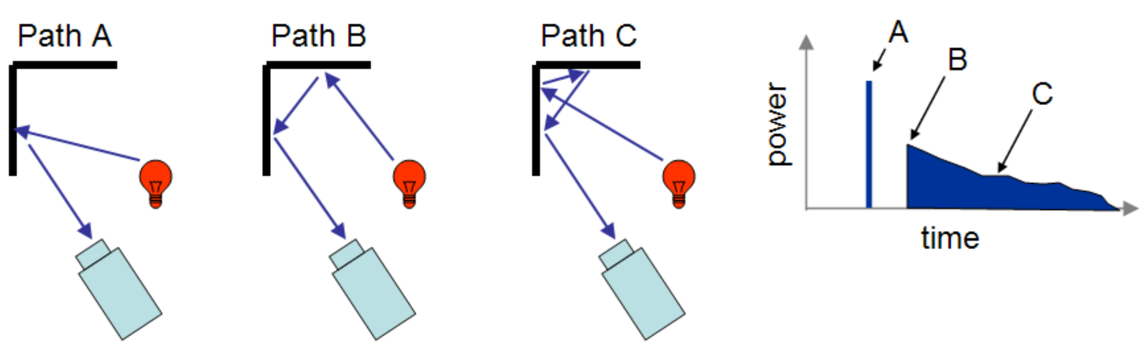

Figure 1. Light transport that proceeds along different scene paths takes different amounts of time to reach the camera. Traditional cameras integrate all of this light into a single pixel value. The transient photometric response function captured by a time camera (shown here for a single pixel), allows these paths to be sampled separately, leading to additional scene understanding.

I(x,y,t) to be captured.

While gated imaging is related to LIDAR, it has uses beyond 3D imaging. Nanosecond windows are used for imaging tanks at the range of kilometers [Andersson06]. Picosecond gating allows imaging in turbid water [McLean95]. Femtosecond windows allow ballistic photons to be separated from scattered photons while imaging in biological tissue [Farsiu07, Das93].

Most applications make limited use of global reasoning about scene characteristics, instead using a single time-gated window to improve signal to noise ratio while imaging.

It is possible to construct the transient photometric response function using gated imagers, e.g. Busck et. al show a TPRF measured to 100 picosecond accuracy [Busck04].

Streak cameras: Streak cameras are ultrafast photonic recorders which deposit photons across a spatial dimension, rather than integrating them in a single pixel. Using a 2D array, I(x, y_{delta} ,t) can be measured. Sweeping the fixed direction, y_{delta} , allows I(x, y, t) to be captured.

Picosecond streak cameras have been available for decades [Campillo83]. Modern research systems can function in the attosecond range [Itatani02]. Commercially available products image in the femtosecond regime [Hamamatsu].

Global light transport: Light often follows a complex path between the emitter and sensor. Computer vision researchers have developed complex models for reasoning about this path. This has led to work on reconstructing specular scenes [Kutulakos05], transparent scenes [Morris07], Lambertian scenes [Nayar90], reflectance properties [Yu99], and scattering properties [Narasimhan06]. All of this work has made use of traditional cameras which provide measurements *only* of the steady-state light transport phenomena. In this work, we propose that transient light transport can both be observed and meaningfully used to improve estimates of scene properties.

In computer graphics, a description of steady-state light transport in a scene is refered to as the "rendering equation" [Kajiya 1986]. Extensions have been described to include transient light transport [Arvo93], but no rendering work has yet built on this foundation.

Capturing light transport: Recent work in image-based modeling and computational photography has shown several methods for capturing steady-state light transport [Sen05, Garg06, Masselus03, Debevec00]. The incident illumination is represented as a 4D illumination field and the resultant radiance is represented as a 4D view field. Taken together, the 8D reflectance field represents all time-invariant interaction of a scene.

More relevant to this paper, the light transport has been decomposed into direct and indirect components [Nayar06] as well as into multi-bounce components [Seitz05] under strong assumptions. Nayar et. al. assume that the scene has no high-frequency reflectance object (like mirrors) so that the entire indirect component is low frequency. Seitz et. al. assume that the scene is made up of purely diffuse components so that the appearance is view independent.

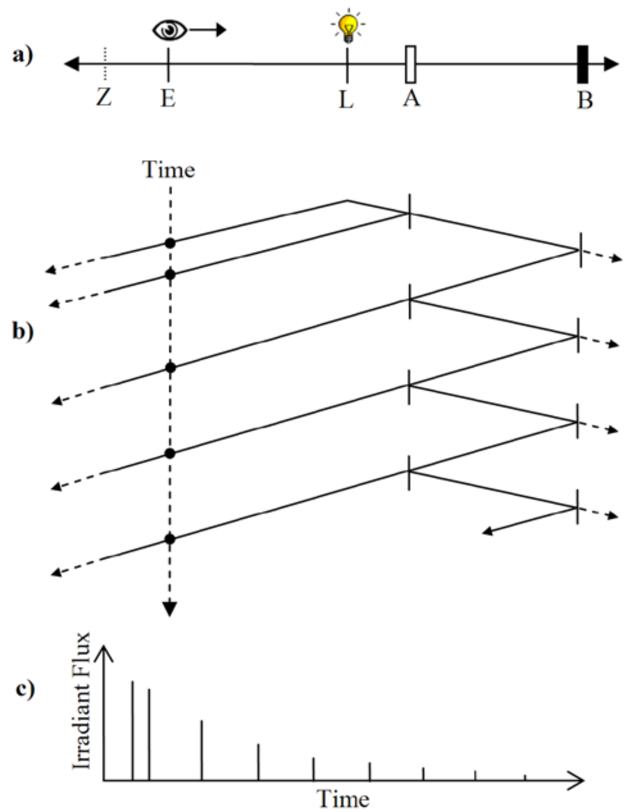

Figure 2. Transient light transport in a 1D scene. (a) 1-D world with real interfaces A and B, eye point E, light source L, and boundary Z (b) Result of transient transport (c) Time profile at E looking to the right

3. Theoretical Model

We first introduce the concept of transient photometric response function and then introduce a model of the sensor to model captured data.

3.1. Transient Photometric Response Function

The 4D light field in a scene is described using a simplified version of the rendering equation [Kajiya86] as follows. If we assume that the scene is discretized into small patches, then the lightfield at patch i can be decomposed into direct and global components. We can write them in a matrix form. The discussion below assumes a single wavelength but each wavelength can be treated separately. The direct component is the emissive lightfield at patch *i*, *E[i]*. The global component, *G[i]*, is the contribution of light from all other patches to patch *i*. The proportion of irradiance directed from patch *j* that contributes to the light field at *i* is denoted by $\rho(i,j)$. It includes bidirectional reflectance distribution function (BRDF) at *i* and visibility between the two patches (Figure 3).

$$L[i] = E[i] + \sum_j \rho(i,j) L[j]$$

In case of transient rendering, we update the rendering equation with the time-varying lightfield equation.

$$L[i,t] = E[i,t] + \sum_j \rho(i,j) L[j, t - d_{ij}],$$

where, d_{ij} , is the time for propagation from patch *j* to patch *i*.

The transient transport in our case is based on impulse illumination, represented as $\delta(0)$ at point *P*. Hence each patch, the direct illumination is $E[i,t] = \rho(p,i)\delta(0 - d_{pi})$.

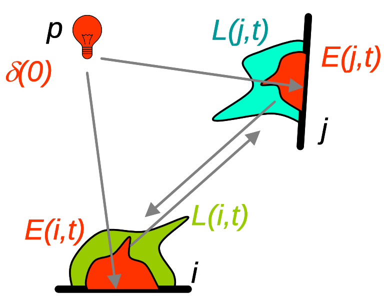

Figure 3. The global transport with time varying light field.

It is important to note the difference with respect to the steady-state transport. In traditional models, each bounce is considered for each iteration of the recursive matrix multiplication. But for transient transport, we must consider the path length and the propagation time d_{ij}. For traditional models, the equation for all patches can be written together in a compact matrix form as $L = E + \rho L$. For transient transport, the global term $\rho(i,j)L[j, t - d_{ij}]$ involves a change in L coordinates for each patch making it a non-linear equation. Nevertheless, the equation provides significantly more constraints on scene geometry and photometry. In later sections, we describe how we can infer global scene parameters in our initial explorations.

3.2. Sensor Response

The standard camera samples the 5D (4D plenoptic function plus time) light field via a 2D projection. We can instead capture a 3D projection. Figure 1 shows an example of temporal profile at a pixel aimed at near the corner of two diffuse surfaces.

Let us look at a simple 1D example shown in Figure 2 to understand the process better. We can plot the temporal profile due to reflection of an impulse from a semi-transparent glass at A and a diffuser at B. Here the $\rho(.,.)$ functions are simpler.

$$L[i,t] = \rho(p,i)\delta(0 - d_{pi}) + \sum_j \rho(i,j) L[j, t - d_{ij}],$$

Hence, we can expand the three recursive terms.

$$L[i,t] = \delta(0 - d_{pi}) + \rho(i,a)L[a, t - d_{ia}] + \rho(i,b)L[b, t - d_{ib}]$$

$$L[a,t] = \delta(0 - d_{pa}) + \rho(a,b)L[b, t - d_{ab}]$$

$$L[b,t] = \delta(0 - d_{pb}) + \rho(b,a)L[a, t - d_{ba}]$$

The recursion is clearly seen in the time-profile in Figure 2. After the first peaks, we see a same pattern due to A and B repeated after regular intervals.

4. Emerging Sensor Solutions

Significant advances in emitters and detectors have made ultra-fast imaging a growing and profitable subfield of imaging. As discussed earlier in Related Works section, based on existing technology it is nearly impossible to create a camera that will create a single shot time profile for all the pixels. However, by scanning in space and time at a high rate, one can generate time-image data. Figure 4 shows a proposal for a design that will make pico-second sensing possible based on incremental improvements in existing technology. A short pulse laser illuminating the scene is mixed with a delayed reference in a non-linear optical crystal. Non-linear optical crystal is a common solution for interferometric measurements so that the output is a convolution of the emitted (delayed) pulse and the reflected signal. By changing the delay via computer controlled translation stage, one can scan upto the maximum required delay values where the transient response eventually behaves like a steady-state response.

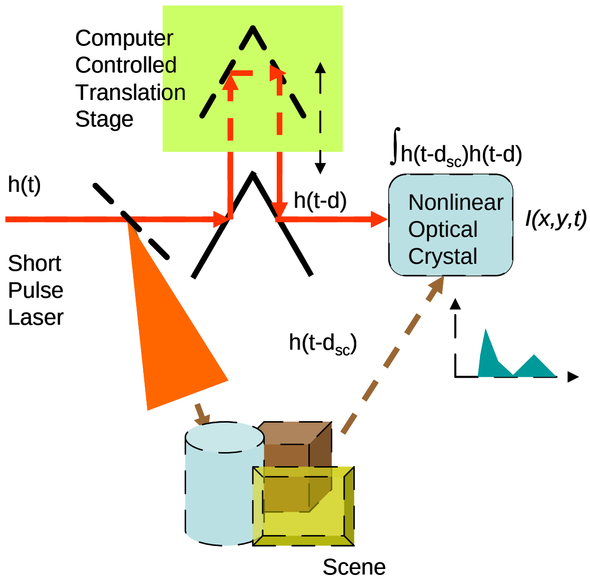

Figure 4. A design for a future time-camera based on incremental improvements to available components. The short duration pulse is used to estimate the impulse response along with a high speed controlled delay via a translation stage. The non-linear optics allows gating of reflected light so that a high speed time profile can be captured. Optional streak-apparatus can record the I(x,y,t) data.

5. Applications to sensing

We can exploit the transient values captured in multiple ways. The model of imaging is sufficiently general and can be used with single pixel cameras, multi-pixel cameras, single emitters and projector-like directional emitters. We can also generate the light impulse in a more controlled way. In the examples below we use a directional laser for generating the impulse.

5.1. Mirror and Shiny Surfaces

Consider the example shown in Figure 5(top) with an arbitrary reflectance (marked as a 'shiny surface' which can have a range of reflectances, from a mirror to a diffuse reflectance). We will discuss these cases in flat-land but with minor modification they are application to 3D configurations. In traditional cameras, it is difficult to estimate distance to a mirror or highly shiny surface because there is no direct reflection received at the camera. But using transient response, we can estimate the distance to shiny surfaces by observing indirect radiance.

If we aim the laser towards the shiny surface in a known direction, it will strike an unknown point *X*. The reflected light will in turn illuminate one (if it is a mirror) or several points. Let us consider one such observable point on a diffuser, *W*. When the laser illuminates the point *W*, we can estimate the 3D location of *W* via stereo triangulation using known angle of the laser beam and the camera pixel observing *W*. From that we can also estimate the time of flight to and from *W*, and hence the path length, i.e., *c* and *d*, as shown. When the laser illuminates *X*, the total path length sensed at pixel observing *W* is ($a+b+d$). Since, *d* is known, the point *X* lies on the locus such that ($PX + XW$) = ($a+b$), which is an ellipse with *P* and *W* as foci. Given the known ray direction <*PX*>, we intersect the ray with the ellipse to recover *X*.

We also get the bidirectional reflectance at *X* with respect to *XP* and *XW* based on the two intensities measured at the pixel observing *W*: when the laser illuminates *X* versus *W*. By illuminating points in the neighborhood of *X*, we can estimate their position independent of the BRDF and in turn we can estimate surface curvature. All these tasks will be difficult in standard cameras without tremendous instrumentation of the scene via large known reflectors.

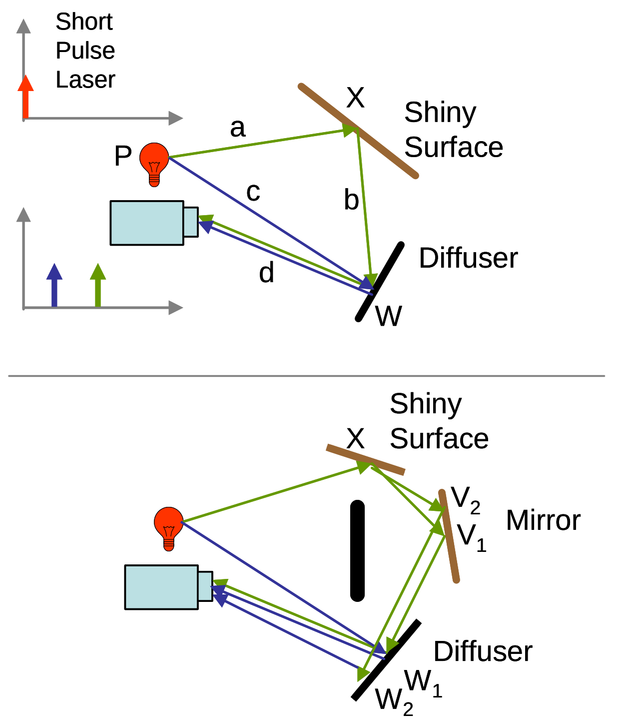

Figure 5. Estimating properties of shiny unknown reflectance objects and hidden objects. (Top) Shiny surfaces are indirectly observed via path lengths to observable diffuse surfaces. (Bottom) Hidden points, V_i , are estimated from multiple observable points.

5.2. Hidden Surfaces

Consider the example shown in Figure 5(bottom) which contains points V_i that are hidden behind an occluder. If we assume those points lie on a planar mirror, can we estimate them from global measurements? Note that this case is different from the hidden surface case discussed in [Sen05] where they used a light source that can directly illuminate the hidden points.

We again make no assumptions about the BRDF of the other scene surfaces but assume that we can indirectly illuminate neighboring points, V_i, say via the point X, which in turn illuminate observable points W_i on a diffuser. (The points V_i can be indirectly illuminate by a distinct set of points X_i for the discussion below, but for simplicity let us assume a single common point X). Mirrors induce shortest paths for light propagation. We can find shortest path, ($XV_i + V_iW_k$) , between X and a specific diffuser point W_k using first onsets in time profile at the pixel observing W_k. The corresponding point, V_k ,with shortest path, lies on an ellipse with foci X and W_k. Note that a normal to the ellipse is found by bisecting the angle between XV_i and W_iV_i, and a tangent perpendicular to the normal behaves like a mirror. If we plot the ellipse for each X-W_i pair with the appropriate shortest path length, the line corresponding to V_i's will be the common tangent line for the ellipses. The tangent gives the location of the hidden mirror plane. In addition, assuming diffuse reflectance at $W_{i,}$ the ratio of pixel intensities recovers the bidirectional reflectance at V_i points.

Thus, a combination of (i) known path lengths based on onsets in time-profile, (ii) constraints on directions induced by camera pixel or laser beam and (iii) ratios of intensities, allows us to estimate a part of the global representation of the scene.

6. Conclusion

We have presented a conceptual framework for exploring new opportunities in multi-path analysis using time-of-flight sensors. A time-image camera described here is not available today but the pico-second resolution impulse response can be captured by scanning in time or space. Emerging trends in femto-second accurate emitters, detectors and non-linear optics may support single-shot time-image cameras.

The goal of this paper is to explore the opportunities in multi-path analysis of the transient response. We developed the theoretical basis for analysis and demonstrated potential methods for recovering scene properties using simple examples. But a complete procedure for estimating scene parameters remains future work. The contribution of this work is conceptual rather than experimental. We hope to influence the direction of future research in time-of-flight systems both in terms of sensor design and algorithms for scene understanding.

Acknowledgements We thank Gavin Miller, Adam Smith and James Skorupski for discussions. This update to the first draft of July 2007 report now includes contribution from Adam and James's tech report to Figure 2 and some text in Section 1 and 2 of this report.

References

[3DV] http://www.3dvsystems.com/

[Andersson06] Pierre Andersson, Long-range three-dimensional imaging using range-gated laser radar images, Opt. Eng. 45, (2006)

[Arvo93] J Arvo, "Transfer equations in global illumination", Global Illumination SIGGRAPH '93 Course Notes, 1993.

[Bamji01] C. Bamji, "CMOS-Compatible 3-Dim. Image Sensor IC," *USPatent*; no. US 6,323,942 B1; 2001.

[Blair99] Blair JB, Rabine DL, Hofton MA (1999) The Laser Vegetation Imaging Sensor (LVIS): A medium-altitude, digitization-only, airborne laser altimeter for mapping vegetation and topography. ISPRS Journal of Photogrammetry & Remote Sensing 56: 112-122.

[Busck04] J. Busck and H. Heiselberg, "Gated Viewing and High-Accuracy Three-dimensional Laser Radar," Appl. Opt. 43(24), 4705-4710 (2004).

[Campillo83] Campillo, A.; Shapiro, S., "Picosecond streak camera fluorometry--A review," Quantum Electronics, IEEE Journal of , vol.19, no.4, pp. 585-603, Apr 1983

[Canasta] http://www.canesta.com/

[Das93] B.B. Das, K. M. Yoo, and R. R. Alfano, "Ultrafast time-gated imaging in thick tissues: a step toward optical mammography", Optics Letters, 18(13), 1993.

[Davis03] J. Davis and H. Gonzalez-Banos, "Enhanced Shape Recovery with Shuttered Pulses of Light," *IEEE Workshop on Projector-Camera Systems*, 2003.

[Debevec00] Debevec P., Hawkins T., Tchou C., Duiker H.-P., Sarokin W., Sagar M.: Acquiring the Reflectance Field of a Human Face. In SIGGRAPH (2000)

[Farsiu07] S. Farsiu, J. Christofferson, B. Eriksson, P. Milanfar, B. Friedlander, A. Shakouri, and R. Nowak, "Statistical detection and imaging of objects hidden in turbid media using ballistic photons," Appl. Opt. 46, 5805-5822 (2007).

[Garg06] Gaurav Garg, Eino-Ville Talvala, Marc Levoy, Hendrik P. A. Lensch, Symmetric Photography: Exploiting Data-sparseness in Reflectance Fields, Proc. Eurographics Symposium on Rendering (EGSR) 2006.

[Gonzales-Banos04] Hector Gonzales-Banos, James Davis, “A method for computing depth under ambient illumination using multi-shuttered light”, *IEEE (CVPR),* 2004.

[Gvili03] R. Gvili and A. Kaplan and E. Ofek and G. Yahav, “Depth Keying,” *SPIE Electronic Imaging 2003 Conference*, 2003.

[Hamamatsu] http://sales.hamamatsu.com/en/products/system-division /ultra-fast/streak-systems.php

[Heckbert91] Heckbert, P. 1991. Simulating Global Illumination Using Adaptive Meshing. PhD thesis, University of California, Berkeley. 26.

[Hofton00] Hofton MA, Minster JB, Blair JB (2000) Decomposition of laser altimeter waveforms. IEEE Transactions on Geoscience and Remote Sensing 38 (4), 1989–1996.

[Iddan01] G. Iddan and G. Yahav, “3D Imaging in the Studio,” *Proc. Of SPIE Videometrics and Optical Methods for 3D Shape Measurement*, SPIE vol. 4298, pp. 48-55, 2001.

[Itatani02] J. Itatani, F. Quéré, G. L. Yudin, M. Yu. Ivanov, F. Krausz, and P. B. Corkum, “Attosecond Streak Camera“, Phys. Rev. Lett., 88(17), 2002.

[Jutzi06] Jutzi, B., Stilla, U. 2006. „Precise range estimation on known surfaces by analysis of full-waveform laser.” In Proceedings of Photogrammetric Computer Vision PCV 2006. International Archives of Photogrammetry and Remote Sensing Vol. 36.

[Kajiya86] Kajiya, J. T. 1986. “The Rendering Equation.“ In Proceedings of the Conference on Computer Graphics and Interactive Techniques. ACM SIGGRAPH 1986.

[Kamermann93] Kamermann, G. W. 1993. Laser Radar, Active Electrical Optical Systems, The Infrared & Electro-Optical Systems Handbook.

[Kawakita00]. Kawakita and K. Iizuka and T. Aida and H. Kikuchi and H. Fujikake and J. Yonai and K. Takizawa, “Axi-Vision Camera (real-time distance-mapping camera),” *Applied Optics* , Opt. Soc. America, vol.39, no.22, p.3931-9, Aug 2000.

[Kutulakos05] Kiriakos N. Kutulakos and Eron Steger, A theory of Refractive and Specular 3D Shape by Light-Path Triangulation, Proc. 10th Int. Conf. Computer Vision, Beijing, China, pp. 1448-1455, 2005

[Lange01] R. Lange and P. Seitz, “Solid-State Time-of-Flight Range Camera,” *IEEE Journal of Quantum Electronics* , vol. 37, no. 3, March 2001.

[Masselus03] Masselus V., Peers P., Dutré P., Willems Y. D.: Relighting with 4D Incident Light Fields. In SIGGRAPH '03 (2003), pp. 613-620.

[McLean95] E. A. McLean, H. R. Burris, Jr., and M. P. Strand, “Short-pulse range-gated optical imaging in turbid water”, Applied Optics 34(21), 1995.

[MESA] http://www.mesa-imaging.ch/

[Miyagawa97] R. Miyagawa and T. Kanade, “CCD-Based Range-Finding Sensor,” *IEEE Transactions on Electron Devices*, vol. 44, no.10, October 1997.

[Morris07] Nigel J. W. Morris and Kiriakos N. Kutulakos, Reconstructing the Surface of Inhomogeneous Transparent Scenes by Scatter Trace Photography, *Proc. 11th Int. Conf. Computer Vision*, 2007

[Narasimhan06] Narasimhan, S. G., Gupta, M., Donner, C., Ramamoorthi, R., Nayar, S. K., and Jensen, H. W. 2006. Acquiring scattering properties of participating media by dilution. In *ACM SIGGRAPH 2006*

[Nayar90] Nayar, S.K.; Ikeuchi, K.; Kanade, T., "Shape from interreflections," Computer Vision, 1990. Proceedings, Third International Conference on , vol., no., pp.2-11, 4-7 Dec 1990

[Nayar06] Nayar, S. K., Krishnan, G., Grossberg, M. D., and Raskar, R. 2006. Fast separation of direct and global components of a scene using high frequency illumination. *ACM Trans. Graph.* 25, 3 (Jul. 2006),

[PMD] http://www.pmdtec.com/

[Russell96] Russell, G. T., Bell, J. M., Holt, P. O., Clarke, S. J. 1996. “Sonar Image Interpretation and Modeling.” In Proceedings of the 1996 Symposium on Autonomous Underwater Vehicle Technology, AUV, 1996.

[Schroeder99] W. Schroeder and E. Forgber and S. Estable, “Scannerless laser range camera,” *Sensor Review*, MCB University Press, vol. 19, no. 4, pp. 285-291, 1999.

[Seitz05] Seitz, S. M., Matsushita, Y., and Kutulakos, K. N. A Theory of Inverse Light Transport. *IEEE international Conference on Computer Vision, ICCV,* 2005.

[Sen05] Sen, P., Chen, B., Garg, G., Marschner, S. R., Horowitz, M., Levoy, M., and Lensch, H. P. 2005. Dual photography. *ACM Trans. Graph.* 24, 3 (Jul. 2005)

[Vandapel04] Vandapel N, Amidi O, Miller JR (2004) Toward Laser Pulse Waveform Analysis for Scene Interpretation. IEEE International Conference on Robotics and Automation (ICRA 2004).

[Yahav00] G. Yahav and G. Iddan, “Optimal Ranging Camera,” *United States Patent*; no. US 6,057,909; May 2, 2000.

[Yu99] Yu, Y., Debevec, P., Malik, J., and Hawkins, T. 1999. Inverse global illumination: recovering reflectance models of real scenes from photographs. In *Proceedings of the 26th Annual Conference on Computer Graphics and interactive Technique, SIGGRAPH,* 1999.