

Recovering Three-dimensional Shape Around a Corner using Ultra-fast Time-of-Flight Imaging

Andreas Velten¹, Thomas Willwacher², Otkrist Gupta¹,
Ashok Veeraraghavan³, Mounqi G. Bawendi¹, Ramesh Raskar¹
¹MIT, ²Harvard University, ³Rice University

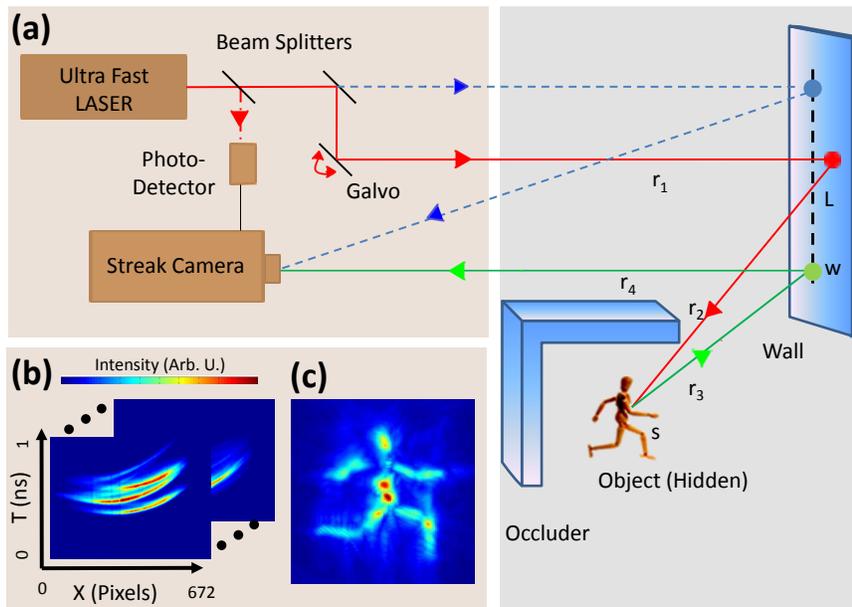


Figure 1: Experimental Setup to recover the 3D shape of a hidden object. (a) The capture process. We capture a series of images by sequentially illuminating a single spot on the wall with a pulsed laser and recording an image of the dashed line segment on the wall with a streak camera. The laser pulse strikes the wall at a point L , some of the diffusely scattered light strikes the hidden object (for example at s), returns to the wall (for example at w) and is collected by the camera. The position of the laser beam on the wall is changed by a set of galvanometer actuated mirrors. (b) An example of streak images sequentially collected. (c) The 2D projected view of the 3D shape of the hidden object, as recovered by the reconstruction algorithm.

Table of Contents Abstract

We demonstrate an incoherent ultra-fast imaging technique to recover 3D shapes of non-line-of-sight objects using their diffuse reflections from line-of-sight surfaces.

Abstract

One goal of scientific imaging is to image objects obscured due to scattering by exploiting, for example, coherence, ballistic photons or penetrating wavelengths.

Common methods depend on scattered light through semi-transparent occluders but fail for opaque occluders. Instead of scattered light through the occluder, we use diffuse reflections from neighboring visible surfaces. The light from hidden scene points is mixed after the diffuse reflections before reaching the image sensor. This mixing is difficult to decode using a traditional camera. We show a time-of-flight (ToF) technique in combination with computational reconstruction algorithms to untangle image information mixed in the diffuse reflections. In effect, we demonstrate a 3D range camera able to “look around a corner” and

demonstrate an experimental method which achieves up to sub-millimeter depth precision, and centimeter lateral precision over 40 cm x 40 cm x 40 cm of hidden space.

1 Introduction

The light detected on an image sensor is composed of light from objects that are in the line of sight as well as those beyond the line of sight. Light from objects outside the line of sight reaches the sensor via multiple diffuse reflections (or *bounces*) and is usually discarded. We demonstrate an incoherent ultra-fast imaging technique to recover 3D shapes of non-line-of-sight objects using those inter-reflections.

The reconstruction of an image from diffusely scattered light is of interest in a variety of fields. Change in spatial light distribution due to the propagation through a turbid medium is in principle reversible [Zahid Yaqoob and Yang 2008] and allows imaging through turbid media via computational imaging techniques [Dylov and Fleischer 2010; Popoff et al. 2010; Vellekoop et al. 2010]. Careful modulation of light can shape or focus pulses in space and time inside a scattering medium [Choi et al. 2011; Katz et al. 2011]. Images of objects behind a diffuse screen, such as a shower curtain, can be recovered by exploiting the spatial frequency domain properties of direct and global components of scattered light in free space [Nayar et al. 2006]. Our treatment of scattering is different but could be combined with many of these approaches.

Line-of-sight time of flight information is commonly used in LIDAR (LIght Detection And Ranging) [Kamerman 1993] and two dimensional gated viewing [Busck and Heiselberg 2004] to determine the object distance or to reject unwanted scattered light. By considering only the early “ballistic” photons from a sample, these methods can image through turbid media or fog [Wang et al. 1991]. We use diffuse reflections from neighboring surfaces to reconstruct and image *around* an occluder instead of analyzing volumetric scattering *through* an object.

Recent methods in computer vision and inverse light transport study multiple diffuse reflections in free-space. Dual Photography [Sen et al. 2005] shows one can exploit scattered light to recover 2D images of objects illuminated by a structured dynamic light source

and hidden from the camera. Time gated viewing using mirror reflections allows imaging around corners, for example from a glass window [Repasi et al. 2009; Sume et al. 2009; Chakraborty et al. 2010]. Three bounce analysis of a time-of-flight camera can recover hidden 1-0-1 planar barcodes [Kirmani et al. 2009] but the authors assume well separated isolated hidden patches with known correspondence between hidden patches and recorded pulses. In addition, a time-of-flight camera can recover black and white patterns on a hidden plane (i.e. a 2D image) if the position of the hidden plane is known [Kirmani et al. 2011]. Similar to these and other inverse light transport approaches [Seitz et al. 2005], we use a light source to illuminate one scene spot at a time and record the reflected light after its interaction with the scene. But, instead of a 2D image, our goal is to recover the 3D structure of a hidden scene. We show that the extra temporal dimension of the observations under very high temporal sampling rates makes the hidden 3D structure observable. With a single or a few isolated hidden patches, pulses recorded after reflections are distinct and can be easily used to find 3D position of the hidden patches. However, with multiple hidden scene points, the reflected pulses may overlap in time when they arrive at the detector. The loss of correspondence between 3D scene points and their contributions to the detected pulse stream is the main technical challenge. We present a computational algorithm based on back-projection to invert this process.

While most of the previous 2D imaging experiments are algebraically linear in nature, the reconstruction of a 3D surface shape is intrinsically a nonlinear algebraic problem. The 2D hidden image recovery can be viewed as a simplified special case of our problem statement.

Our main contributions are two-fold. (i) We introduce the new problem of recovering the 3D structure of a hidden object and we show that the 3D information is retained in the temporal dimension after multi-bounce interactions between visible and occluded parts. (ii) We present an experimental realization of the ability to recover the 3D structure of a hidden object, thereby demonstrating a 3D range camera able to “look around a corner”. The ability to record 3D shapes beyond the line of sight can potentially be applied in industrial inspection, endoscopic observations, disaster relief scenario or more generally in situations where direct imaging of a scene is impossible.

Imaging Process The experimental setup is shown in Figure 1. Our scene consists of a 40 cm high and 25 cm wide wall referred to as *diffuser wall*. We use an ultra-fast laser and a streak camera and both are directed at this wall. We place a target object hidden in the scene (mannequin in Figure 1). We block direct light paths between the object and the laser or the camera. Our goal is to produce three dimensional range data of the target object.

The streak camera records a streak image with one spatial dimension and one temporal dimension. We focus the camera on the dashed line segment on the diffuser wall shown in Figure 1(a). We arrange the scanning laser to hit spots on the wall above or below this line segment so that single bounce light does not enter the camera. Though the target object is occluded, light from the laser beam is diffusely reflected by the wall, reaches the target object, is reflected by multiple surface patches and returns back to the diffuser wall, where it is reflected again and captured by the camera. In a traditional camera, this image would contain little or no information about the occluded target object. In our experimental setup, the laser emits 50 femtosecond long pulses. The camera digitizes information in time intervals of 2 picoseconds. We assume the geometry of the directly visible part of the setup is known. Hence the only unknown distances in the path of the laser pulses are those from the diffuser wall to the different points on the occluded target object and back (paths r_2 and r_3 in Fig. 1). The 3D geometry of the occluded target is thus encoded in the streak images acquired by the camera and decoded using our reconstruction algorithm.

The recorded streak images lack correspondence information, i.e., we do not know which pulses received by the camera came from which surface point on the target object. Hence, a straightforward triangulation or trilateration to determine the hidden geometry is not possible.

Consider a simple scenario with a small hidden patch as illustrated in Figure 2a. It provides intuition on how the geometry and location of the target object are encoded in the streak images. The reflected spherical pulse propagating from the hidden patch arrives at the points on the diffuser wall with different time delays and creates a hyperbolic curve in the space-time streak image. We scan and successively change the position of the laser spot on the diffuser wall. The shape and position of

the recorded hyperbolic curve varies accordingly. Each pixel in a streak image corresponds to a finite area on the wall and a 2 picosecond time interval, a discretized *space time bin*. However, the effective time resolution is 15 picoseconds due to a finite temporal point spread function of the camera. A detailed description of image formation is included in Section 1 of the supplementary material.

The inverse process to recover the position of the small hidden patch from the streak images is illustrated in Figure 2b–e. Consider three pixels p, q and r in the streak image at which non zero light intensity is measured (Figure 2a). The possible locations in the world which could have contributed to a given pixel lie on an ellipsoid in Cartesian space. For illustration, we draw only a 2D slice of the ellipsoid, i.e., an ellipse, in Figure 2b. The individual ellipses from each of the three pixels p, q and r intersect at a single point. In the absence of noise, the intersection of three ellipses uniquely determines the location of the hidden surface patch that contributed intensity to the three camera pixels. In practice we lack correspondence, i.e., we do not know whether or not light detected at two pixels came from the same 3D surface point.

Therefore, we discretize the Cartesian space into voxels and compute the likelihood of the voxel being on a hidden surface. For each voxel, we determine all streak image pixels that could potentially have received contributions of this voxel based on the time of flight $r_1 + r_2 + r_3 + r_4$ and sum up the measured intensity values in these pixels. In effect, we let each pixel “vote” for all points on the corresponding ellipsoid. The signal energy contributed by each pixel is amplified by a factor of $r_2 r_3$ to compensate for the distance attenuation. This process of computing likelihood by summing up weighted intensities is called *backprojection*. The 3D scalar function on voxels is called a *heatmap*.

The summation of weighted intensities from all pixels in a single streak image creates an approximate heatmap for the target patch (Figure 2c). Repeating the process for many laser positions on the diffuser wall, and using pixels from the corresponding streak images provides a better approximation (Figure 2d). In practice, we use approximately 60 laser positions. Traditional backprojection requires a high-pass *filtering* step. We use the second derivative of the data as an effective filter and recover the hidden surface patch in Figure 2e.

The reconstruction is described in detail in Sections 2 and 3 of the supplementary material.

Algorithm The outline of our imaging algorithm is as follows.

- **Data Acquisition:** Sequentially illuminate a single spot on the diffuser wall with a pulsed laser and record an image of the line segment of the wall with a streak camera.
- **Voxel Grid Setup:** Estimate an oriented bounding box for the working volume to set up a voxel grid in Cartesian space (see below).
- **Backprojection:** For each voxel, record the summation of weighted intensity of all streak image pixels that could potentially have received contributions of this voxel based on the time of flight. Store the resulting three dimensional *heatmap* of voxels.
- **Filtering:** Compute a second derivative of the heatmap along the direction of the voxel grid facing away from the wall.
- **Postprocessing (optional):** Compute confidence value for each voxel by computing local contrast with respect to the voxel neighborhood in the filtered heatmap. To compute contrast, divide each voxel heatmap value by the maximum in the local neighborhood. For better visualization, apply a soft threshold on the voxel confidence value.

We estimate the oriented bounding box of the object in the second step by running the above algorithm at low spatial target resolution and with downsampled input data. Details of the reconstruction process and the algorithm can be found in the Methods Section as well as in the supplementary material in Section 3.

2 Results

We show results of the 3D reconstruction for multi-part objects in Figures 3 and 4. The mannequin in Figure 3 contains non-planar surfaces with variations in depth and occlusions. We accurately recover all major geometrical features of the object. Figure 3i shows the reconstruction of the same object in slightly different poses to demonstrate the reproducibility and stability of the method as well as the consistency in the captured

data. Note that the sporadic inaccuracies in the reconstruction are consistent across poses and are confined to the same 3D locations. The stop-motion animation in the supplementary movie shows the local nature of the missing or phantom voxels more clearly. Hence, the persistent inaccuracies are not due to signal noise or random measurement errors. This is promising as the voxel confidence errors are primarily due to limitations in the reconstruction algorithm and instrument calibration. These limitations can be overcome with more sophistication.

Figure 4 shows a reconstruction of multiple planar objects at different unknown depths. The object planes and boundaries are reproduced accurately to demonstrate depth and lateral resolution.

The reconstruction is affected by several factors such as calibration, sensor noise, scene size and time resolution. Below, we consider them individually.

The sources of calibration error are lens distortions on the streak camera that lead to a warping of the collected image, measurement inaccuracies in the visible geometry, and measurement inaccuracies in the center of projection of the camera and the origin of the laser. For larger scenes, the impact of static calibration errors is reduced.

The sensor introduces intensity noise and timing uncertainty, i.e., jitter. The reconstruction of 3D shapes is more dependent on the accuracy of the time of arrival than the signal to noise ratio (SNR) in received intensity. Jitter correction, as described in the Methods section, is essential, but does not remove all uncertainties. Improving the SNR is desirable because it yields better jitter correction and faster capture times. Similar to many commercial systems, e.g., LiDAR, the SNR could be significantly increased by using an amplified laser with a repetition rate in the kHz range and a triggered camera. The overall light power will not change, but fewer measurements for light collection should significantly reduce detection noise.

We can increase the scale of the system for larger distances and bigger target objects. By using a longer pulse one could build solutions without any change in the ratio of received and emitted energy, i.e., the link budget. When the distance r_2 between diffuser wall and the hidden object (see Fig. 1) is increased, the signal strength drops dramatically ($\propto 1/(r_2 r_3)^2$) and the

size of the hidden scene is therefore limited. A configuration where laser and camera are very far from the rest of the scene is, however, plausible. A loss in received energy can be reduced in two ways. The laser beam can be kept collimated over relatively long distances and the aperture size of the camera can be increased to counterbalance a larger distance between camera and diffuser wall.

The timing resolution, along with spatial diversity in the positions of spots illuminated and viewed by the laser and the camera affects the resolution of 3D reconstructions. Additional factors include the position of the voxel in Cartesian space and the overall scene complexity. The performance evaluation section of the supplementary material describes depth and lateral resolution. In our system, translation along the direction perpendicular to the diffuser wall can be resolved with a resolution of $400 \mu\text{m}$ – better than the full width half maximum (FWHM) time resolution of the imaging system. Lateral resolution in a plane parallel to the wall is lower and is limited to 0.5-1 cm depending on proximity to the wall.

3 Discussion

This paper’s goals are twofold: to introduce the new challenging problem of recovering the 3D shape of a hidden object and to demonstrate the results using a novel co-design of an electro-optic hardware platform and a reconstruction algorithm. We have demonstrated the 3D imaging of a non-trivial hidden three dimensional geometry from scattered light in free space. We compensate for the loss of information in the spatial light distribution caused by the scattering process by capturing ultra-fast time of flight information.

The application of imaging beyond the line of sight is of interest for sensing in hazardous environments such as machinery with moving parts, for monitoring highly contaminated areas such as the sites of chemical or radioactive leaks where even robots can not operate or need to be discarded after use [Blackmon et al. 1999]. Disaster response and search and rescue planning, as well as autonomous robot navigation can benefit from the ability obtain complete information about the scene quickly [Burke et al. 2004; Ng et al. 2004]

A very promising theoretical direction is in inference and inversion techniques that exploit scene priors, spar-

sity, rank, meaningful transforms and achieve bounded approximations. Adaptive sampling can decide the next-best laser direction based on a current estimate of the 3D shape. Further analysis will include coded sampling using compressive techniques and noise models for SNR and effective bandwidth. Our current demonstration assumes friendly reflectance and planarity of the diffuse wall. But our reconstruction method can easily include non-Lambertian behavior and a different model for the visible parts.

In the future, emerging integrated solid state lasers, new sensors and non-linear optics should provide practical and more sensitive imaging devices. Beyond 3D shape, new techniques should allow us to recover reflectance, refraction and scattering properties and achieve wavelength resolved spectroscopy beyond the line of sight. The formulation could also be extended to shorter wavelengths (e.g., x-rays) or to ultrasound and sonar frequencies. The new goal of hidden 3D shape recovery may inspire new research in the design of future ultra-fast imaging systems and novel algorithms for hidden scene reconstruction.

4 Methods

Capture Setup The light source is a Kerr lens mode-locked Ti:Sapphire laser. It delivers pulses of about 50 femtoseconds length at a repetition rate of 75 MHz. The laser wavelength is centered at 795 nm. The main beam is scanned across the diffuser wall via a system of two galvanometer actuated mirrors. A small portion of the laser beam is split off with a glass plate and is used to synchronize the laser and streak camera as shown in Figure 1.

For time jitter correction, another portion of the beam is split off, attenuated and directed at the wall as the *calibration spot*. The calibration spot is in the direct field of view of the camera and can be seen in Figure 11 in the Supplementary material.

The calibration spot serves as a time and intensity reference to compensate for drifts in the synchronization between laser and camera as well as changes in laser output power. It also helps in detecting occasional shifts in the laser direction due to, for example, beam pointing instabilities in the laser. If a positional shift is detected, the data is discarded and the system is re-calibrated.

The camera is a Hamamatsu C5680 streak cam-

era [Hamamatsu 2011] that captures one spatial dimension, i.e., a line segment in the scene, with an effective time resolution of 15 picoseconds and a quantum efficiency of about 10%. The position and viewing direction of the camera are fixed. The diffuser wall is covered with Edmund Optics NT83 diffuse white paint.

Section 4 and 5 of the Supplementary material describe in detail the experimental setup and its limitations.

Reconstruction Technique We use a set of Matlab routines to implement the backprojection-based reconstruction. Geometry information about the visible part of the scene, i.e., diffuser wall could be collected using our time of flight system. Reconstructing the three dimensional geometry of a visible scene using time of flight data is well known [Busck and Heiselberg 2004]. We omit this step and concentrate on the reconstruction of the hidden geometry. We use a FARO Gauge digitizer arm to measure the geometry of the visible scene and also to gather data about a sparse set of points on hidden objects for comparative verification. The digitizer arm data is used as *ground truth* for later independent verification of the position and shape of hidden objects as shown via transparent planes in Figure 4f.

We estimate the oriented bounding box around the hidden object using a lower resolution reconstruction. We reduce the spatial resolution to 8 mm/voxel and down-sample the input data by factor of 40. We can scan a 40 cm x 40 cm x 40 cm volume spanning the space in front of the wall in 2-4 seconds to determine the bounding box of a region of interest. The finer voxel grid resolution is 1.7 mm in each dimension. We can use the coarse reconstruction obtained to set up a finer grid within this bounding box. Alternatively we can set an optimized bounding box from the collected ground truth. To minimize reconstruction time, we used this second method in most of the published reconstructions. We confirmed that apart from the reconstruction time and digitization artefacts, both methods produce the same results. We compute the principal axis of this low resolution approximation and orient the fine voxel grid with these axes. During back-projection (and voting), we amplify the contribution of each pixel to a particular voxel by a factor of $r_2 r_3^\alpha$ to compensate for the attenuation due to spherical divergence. We experimented with various values of α and found that $\alpha = 1$ is a good choice and avoids excessive amplification and contribution of noisy pixels.

In the postprocessing step, we use a common approach to improve the surface visualization. We estimate the local contrast and apply a soft threshold. The confidence value for a voxel is $V' = \tanh(20(V - V_0))V/m_{loc}$, where V is the original voxel value in filtered heatmap and m_{loc} is a local maximum computed in a 20 x 20 x 20 voxel sliding window around the voxel under consideration. Division by m_{loc} normalizes for local contrast. The value V_0 is a global threshold and set to 0.3 times the global maximum of the filtered heatmap. The tanh function achieves a soft threshold.

5 Author Contributions

The method was conceived by A. Velten and R. Raskar, the experiments were designed by A. Velten, M. G. Bawendi and R. Raskar and performed by A. Velten. The algorithm was conceived by A. Velten, T. Willwacher, O. Gupta, and R. Raskar and implemented and optimized by A. Velten, T. Willwacher, O. Gupta, A. Veeraraghavan, and R. Raskar. All authors took part in writing the paper.

We would like to thank Amy Fritz for her help with acquiring and testing data and Chinmaya Joshi for his help with the preparation of the figures.

References

- BLACKMON, T. T., THAYER, S., TEZA, J., , OSBORN, J., HEBERT, M., AND ET. AL. 1999. Virtual reality mapping system for chernobyl accident site assessment. In *Proceedings of the SPIE*, vol. 3644, 338–345.
- BURKE, J. L., MURPHY, R. R., COOVERT, M. D., AND RIDDLE, D. L. 2004. Moonlight in Miami: A Field Study of HumanRobot Interaction in the Context of an Urban Search and Rescue Disaster Response Training Exercise. *Human-Computer Interaction 19*, 85–116.
- BUSCK, J., AND HEISELBERG, H. 2004. Gated viewing and high-accuracy three-dimensional laser radar. *Applied Optics 43*, 24, 4705–4710.
- CHAKRABORTY, B., LI, Y., ZHANG, J., TRUEBLOOD, T., PAPANDREOU-SUPPAPPOLA, A., AND MORRELL, D. 2010. Multipath exploitation with adaptive waveform design for tracking in urban terrain. In *IEEE ICASSP Conference*, 3894 – 3897.

- CHOI, Y., YANG, T. D., FANG-YEN, C., KANG, P., LEE, K. J., DASARI, R. R., FELD, M. S., AND CHOI, W. 2011. Overcoming the diffraction limit using multiple light scattering in a highly disordered medium. *Phys. Rev. Lett.* 107, 2 (Jul), 023902.
- DYLOV, D. V., AND FLEISCHER, J. W. 2010. Non-linear self-filtering of noisy images via dynamical stochastic resonance. *Nature Photonics* 4, 323–328.
- HAMAMATSU, 2011. Hamamatsu Streak Camera. <http://learn.hamamatsu.com/tutorials/java/streakcamera/>.
- KAMERMAN, G. W., 1993. Laser Radar [M]. Chapter 1 of Active Electro-Optical System, Vol. 6, The Infrared and Electro-Optical System Handbook.
- KATZ, O., SMALL, E., BROMBERG, Y., AND SILBERBERG, Y. 2011. Focusing and compression of ultrashort pulses through scattering media. *Nature Photonics* 5, 6 (Jul), 372–377.
- KIRMANI, A., HUTCHISON, T., DAVIS, J., AND RASKAR, R. 2009. Looking around the corner using transient imaging. In *ICCV*.
- KIRMANI, A., VELTEN, A., LAWSON, E., HUTCHISON, T., GOYAL, V., BAWENDI, M. G., AND RASKAR, R., 2011. Reconstructing an image on a hidden plane using ultrafast imaging of diffuse reflections. In Submission.
- 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 (July), 935–944.
- NG, T. C., IBANEZ-GUZMAN, J., SHEN, J., GONG, Z., WANG, H., AND CHENG, C. 2004. Vehicle following with obstacle avoidance capabilities in natural environments. In *Robotics and Automation, 2004. Proceedings. ICRA '04. 2004 IEEE International Conference on*, vol. 5, 4283 – 4288 Vol.5.
- POPOFF, S., LEROSEY, G., FINK, M., BOCCARAI, A. C., AND GIGAN, S. 2010. Image transmission through an opaque material. *Nature Communications*.
- REPASI, E., LUTZMANN, P., ELMQVIST, O. S. M., GHLER, B., AND ANSTETT, G. 2009. Advanced short-wavelength infrared range-gated imaging for ground applications in monostatic and bistatic configurations. *Applied Optics* 48, 5956–5969.
- SEITZ, S. M., MATSUSHITA, Y., AND KUTULAKOS, K. N. 2005. A theory of inverse light transport. In *Proc. Tenth IEEE International Conference on Computer Vision ICCV 2005*, vol. 2, 1440–1447.
- SEN, P., CHEN, B., GARG, G., MARSCHNER, S. R., HOROWITZ, M., LEVOY, M., AND LENSCH, H. P. A. 2005. Dual photography. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*, ACM, New York, NY, USA, 745–755.
- SUME, A., GUSTAFSSON, M., JNIS, A., NILSSON, S., RAHM, J., AND RBOM, A. 2009. Radar detection of moving objects around corners. In *Proc. of SPIE*.
- VELLEKOOP, I. M., LAGENDIJK, A., AND MOSK1, A. P. 2010. Exploiting disorder for perfect focusing. *Nature Photonics* 4, 320–322.
- WANG, L., HO, P. P., LIU, C., ZHANG, G., AND ALFANO, R. R. 1991. Ballistic 2-d imaging through scattering walls using an ultrafast optical kerr gate. *Science* 253, 5021 (Aug), 769–771.
- ZAHID YAQOUB, DEMETRI PSALTIS, M. S. F., AND YANG, C. 2008. Optical phase conjugation for turbidity suppression in biological samples. *Nature Photonics* 2, 110–115.

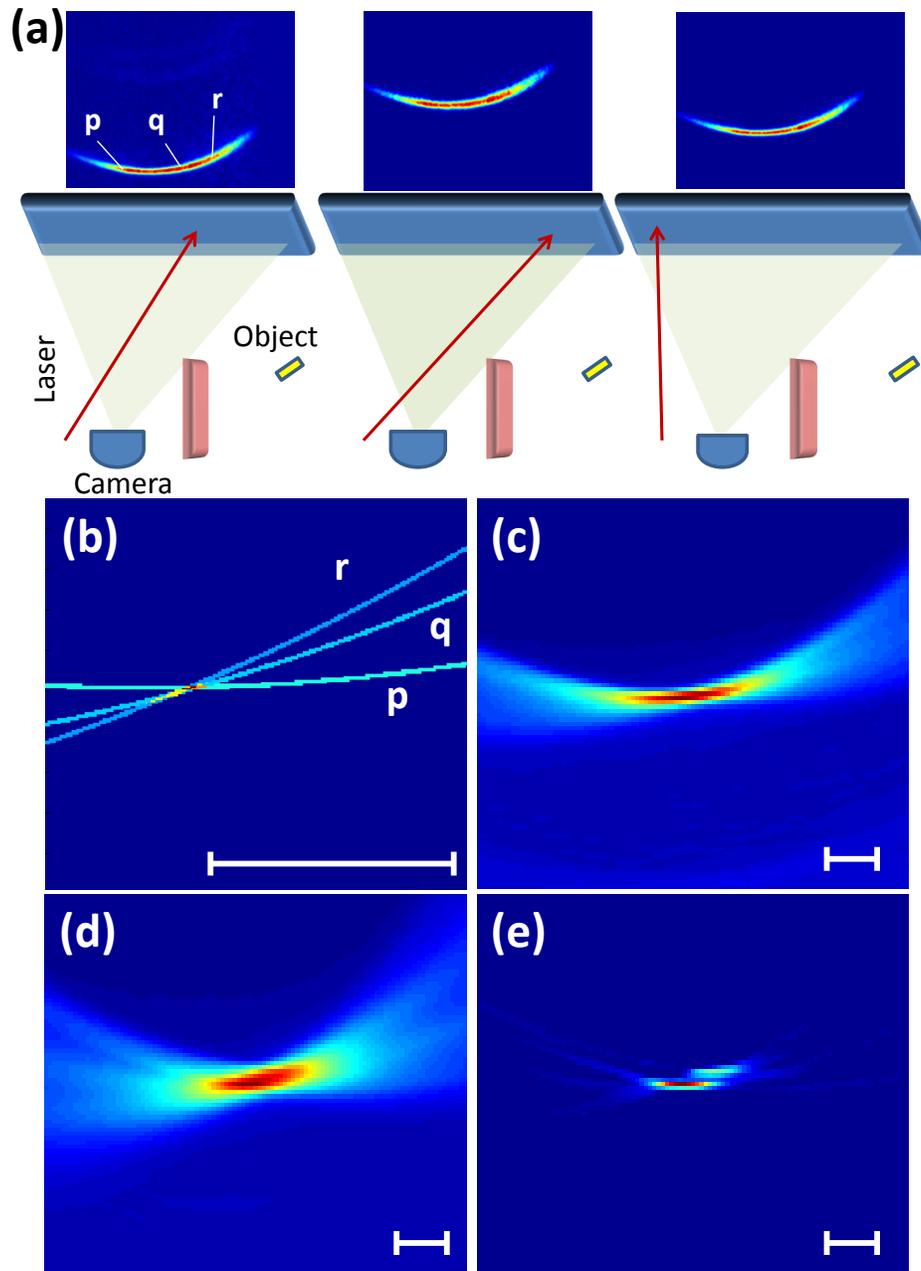


Figure 2: Reconstruction Algorithm An illustrative example of geometric reconstruction using streak camera images. (a) Data capture. The object to be recovered consists of a $2\text{ cm} \times 2\text{ cm}$ size patch beyond the line of sight (i.e. “hidden”). The captured streak images corresponding to each of the three different laser positions are displayed in the top row. Shapes and timings of the recorded response vary with laser positions and encode the position and shape of the hidden patch. (b) Contributing voxels in Cartesian space. For recovery of hidden position, consider the choices of contributing locations. The possible locations in Cartesian space that could have contributed intensity to the streak image pixels p , q , r are ellipses (ellipsoids in 3D). If there is a single world point contributing intensity to all 3 pixels, the corresponding ellipses intersect, as is the case here. The white bar corresponds to 2 centimeters in all sub-figures. (c) Backprojection and heatmap. We use a back-projection algorithm that finds overlaid ellipses corresponding to all pixels. Here we show summation of elliptical curves from all pixels in the first streak image. (d) Backprojection using all pixels in a set of 59 streak images. (e) Filtering. After filtering with a second derivative, the patch location and 2 centimeter lateral size are recovered.

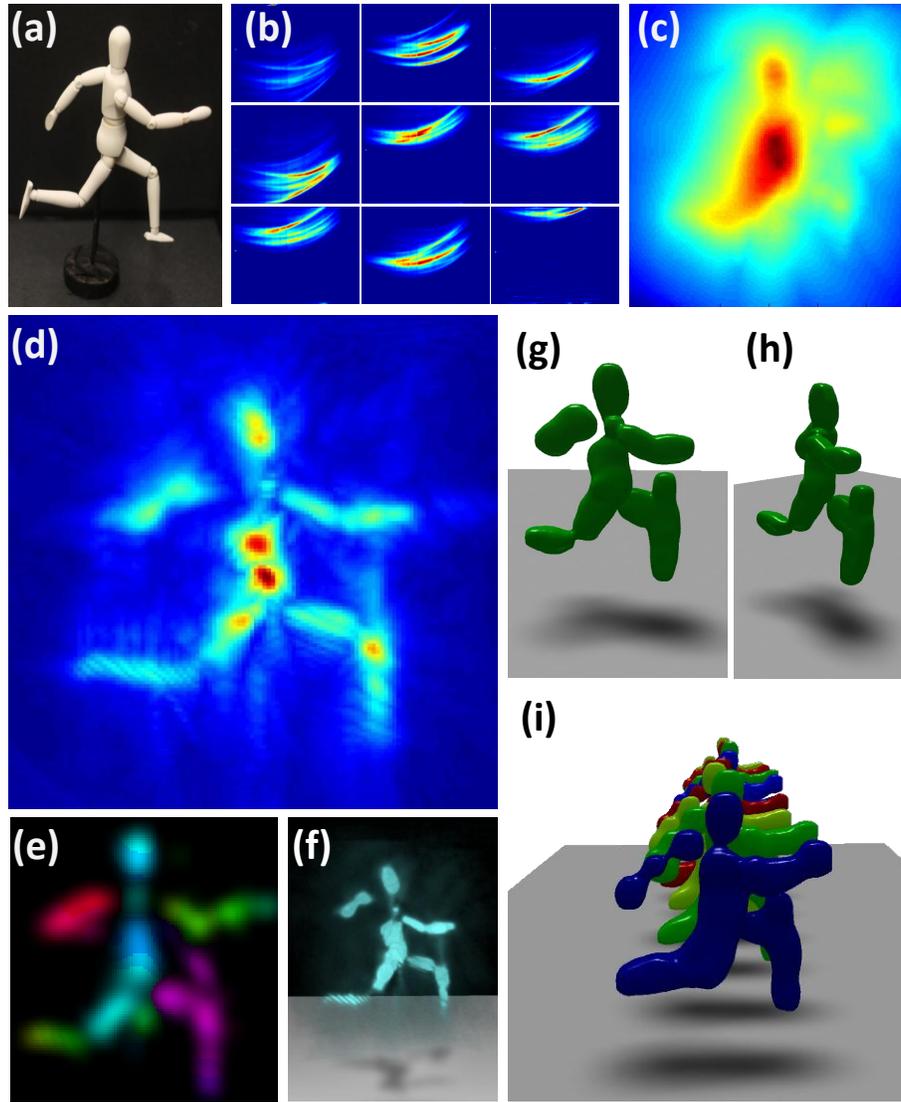


Figure 3: Complex Object Reconstruction in multiple poses. (a) Photo of the object. The mannequin is approximately 20 cm tall and is placed about 25 cm from the diffuser wall. (b) Nine of the 60 raw streak images. (c) Heatmap. Visualization of the heatmap after backprojection. The maximum value along the z direction for each x - y coordinate in Cartesian space. The hidden shape is barely discernible. (d) Filtering. The second derivative of the heatmap along depth (z) projected on the x - y plane reveals the hidden shape contour. (e) Depth map. Color encoded depth (distance from the diffuser wall) shows the left leg and right arm closer in depth compared to the torso and other leg and arm. (f) Confidence map. A rendered point cloud of confidence values after soft threshold. Images (g-h) show the object from different viewpoints after application of a volumetric blurring filter. (i) The stop-motion animation frames from multiple poses to demonstrate reproducibility. See the supplementary video for an animation. Shadows and the ground plane in images (f-i) have been added to aid visualization.

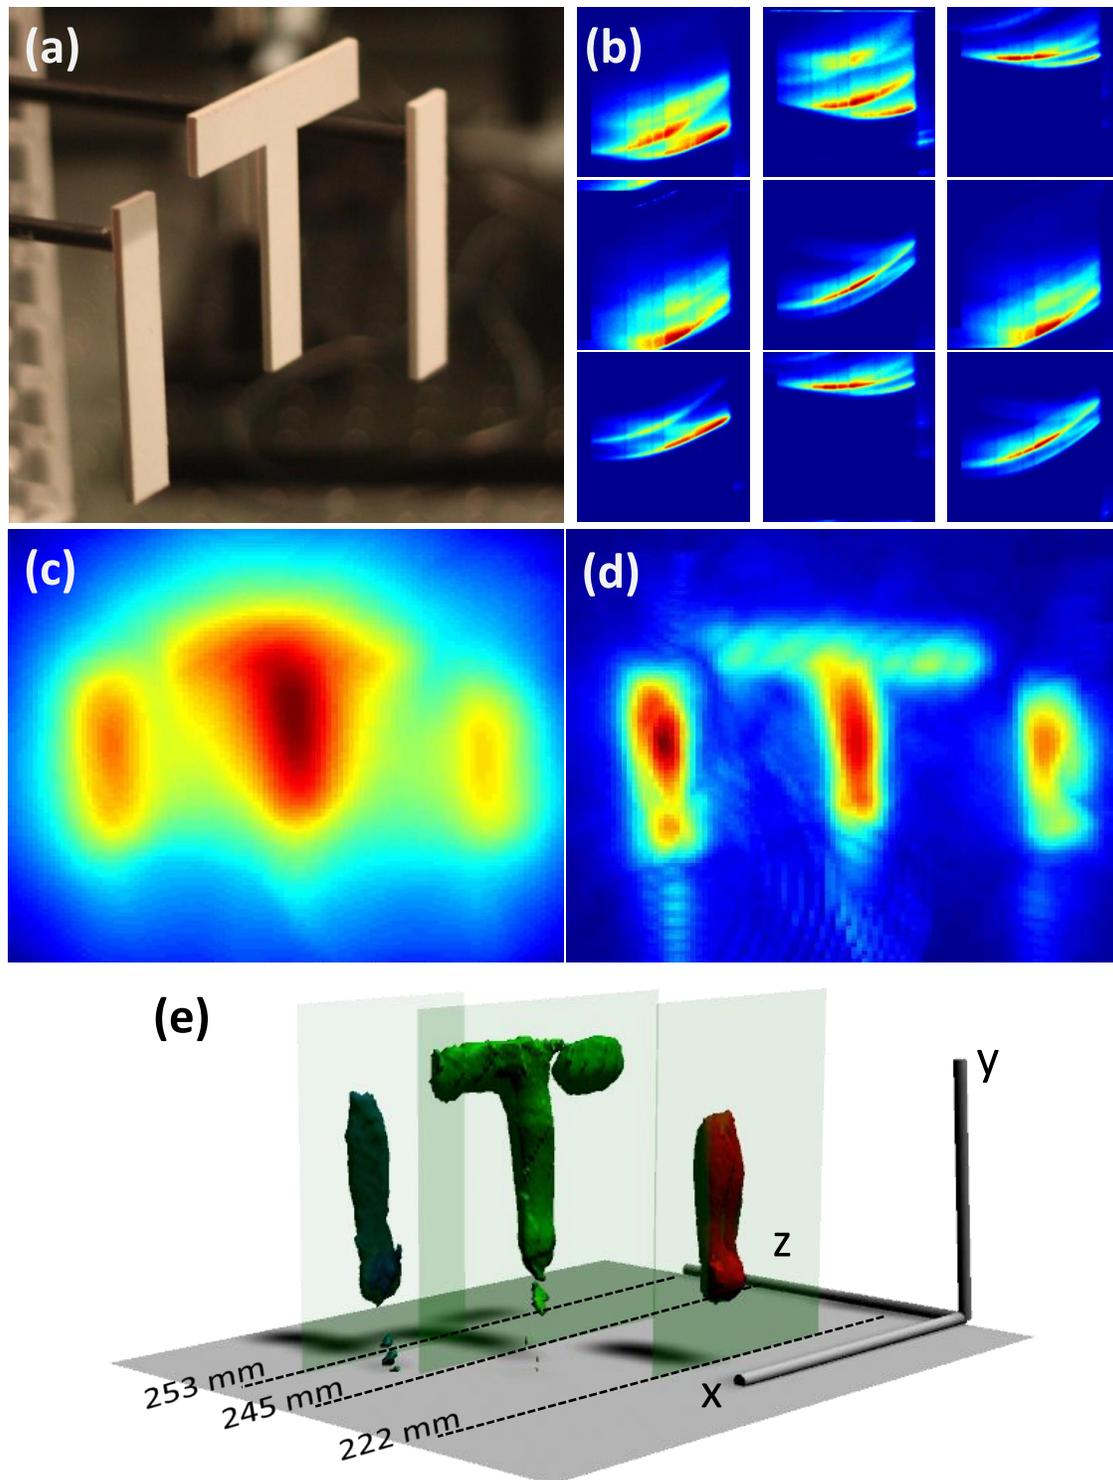


Figure 4: Depth in Reconstructions *Demonstration of the depth and lateral resolution. (a) The hidden object to be recovered are three letters, I, T, I at varying depths. The "I" is 1.5 cm in wide and all letters are 8.2 cm high. (b) 9 of 60 images collected by the streak camera. (c) Projection of the heatmap on the x-y plane created by the back projection algorithm. (d) Filtering after computing second derivative along depth (z). The color in these images represents the confidence of finding an object at the pixel position. (e) A rendering of the reconstructed 3D shape. Depth is color coded and semi-transparent planes are inserted to indicate the ground truth. The depth axis is scaled to aid visualization of the depth resolution.*