

Reconstruction of Photon Distributions in Liquid Scattering Media

Benjamin T. Cecchetto
Queen's University
25 Union Street
K7L 2N8
Kingston, ON, Canada
cecchett@cs.queensu.ca

James Stewart
Queen's University
25 Union Street
K7L 2N8
Kingston, ON, Canada
jstewart@cs.queensu.ca

ABSTRACT

We present an experimental setup to capture photon distributions from a liquid scattering medium and to reconstruct the flux density throughout the medium. The capture mechanism moves an occluder through the medium, which is illuminated by a laser. For each position of the occluder, the exiting light intensity is recorded at selected positions on the boundary of the medium; this produces a measured distribution of intensities blocked by the occluder. The differences between blocked and unblocked intensities are used to solve for sets of photon paths that match the observed exiting light intensity. From these photon paths, we calculate the photon distribution within the medium. This work is validated with experiments using two occluder sizes in different concentrations of milk-water solutions, as well as different incoming and outgoing light poses.

Keywords

Scattering, Physics, Photon Distributions, Optimization, Simulation, Computer Graphics, Rendering, Physically-based Modeling, Imaging

1 INTRODUCTION

This paper presents an experimental setup and algorithm to determine the distribution of photons within a physical scattering medium. In what follows, an input pose (i.e. a position and direction) and an output pose, both on the boundary of the medium, will be fixed. Then the paths of only those photons that enter the medium from the input pose and exit the medium from the output pose will be considered.

In a purely transmissive medium, the photon path is a straight line with all of the light travelling along the line. In a scattering medium, there are many photon paths for a given pair of input and output poses, so the flux through each position will vary across the medium.

The output of this algorithm is, in two dimensions, a *photon distribution image* in which each pixel stores the average flux density through points in the medium corresponding to that pixel, normalized by the output magnitude for a given entry and exit pose pair.

In addition to a fixed entry and exit pose, we consider an additional constraint in order to further study the photon paths. We fix an internal point that the light must travel through, and examine the distribution of photon paths with this additional constraint. Simulated examples of these distributions are depicted in Figure 1.

Knowledge of photon distributions may be of importance to rendering and reconstruction, because the distribution provides probabilistic weightings in the path matrix used in *ART* or *SART* reconstruction [1, 2]. However, it is difficult to capture these distributions using a simple, low-cost setup. Our goal is to capture a photon distribution image of desired resolution so that it may be used in the aforementioned reconstruction techniques.

To obtain the photon distribution image, a laser is directed into the medium and a camera records the output intensity on the other side of the medium. An occluder is then moved through the medium. At each position of the occluder, the measured difference between the output intensity with the occluder and that without the occluder will provide the power of the photons intersecting the occluder, provided the occluder has an albedo of zero. But the occluder must be much larger than a single pixel to obtain a sufficient signal-to-noise ratio in the difference of intensity. To obtain the power at individual pixels, further processing is required.

Photon paths are computed that produce the same differences of output intensity as were measured by the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

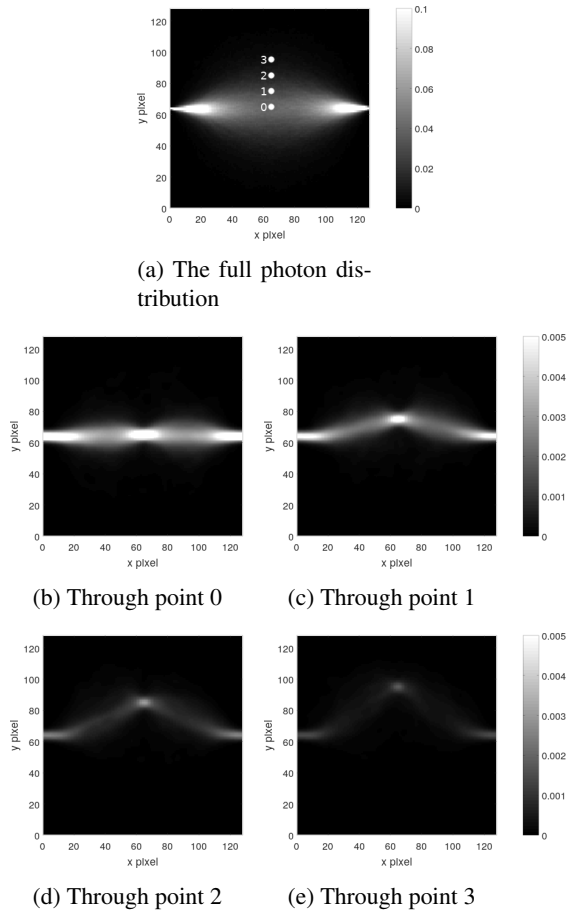


Figure 1: Photon distributions through different points in the medium.

large occluder. These paths can be used to provide an estimated power through any region of the medium, including a pixel region, which is smaller than the occluder.

The use of such paths was motivated by the paths analytically derived by Feng *et al.* [3] for a diffusion approximation. Each of their paths travels from an input pose to an output pose that is a varying distance away from a central “most probable path”. Their derivation also provides an “approximate boundary” of the distribution. Therefore, we propose a similar numerical approach.

We use an active contour-based method, which starts with an initial set of template photon paths originating at the input pose (i.e. the laser position) and terminating at the output pose (i.e. a pixel in a camera image). The simulated output from these photon paths is compared to the physically captured output, and the paths are warped to correct the error between the two. The use of photon paths permits arbitrary distribution shapes to be fit to the captured output. The density of the paths at a point within the medium is proportional to the flux density at that point. Compared to Monte

Carlo methods, this approach provides a very efficient simulation of light transport.

2 RELATED WORK

Our work builds on the distribution derived by Feng *et al.* [3]. The derivation depended on a diffusion model [4] for photon migration. A similar experimental setup has captured photon distributions, but was limited by the pixel resolution [5]. Feng *et al.* built on their work to perform an analysis for an object blocking the photon paths [6] as well. The diffusion model has been used to perform backprojection [7] to solve for the interior of a scattering region. With the actual photon distributions, it may be possible to perform optical tomography using a variety of techniques [8].

There are many different scattering approximations that can be used to render a scene. Real-time systems need further approximations to obtain the necessary performance [9], as ground truth Monte-Carlo simulations cannot yield real-time performance. Real-time systems use depth of materials, or approximate scattering in texture space. One of the more important contributions is by Jensen *et al.* [10], in which a diffusion approximation is used to derive a local model for incoming and outgoing light.

Light can be sampled through a scattering medium along paths to determine whether the light paths are occluded [11]. Sampling methods use parts of physical models to render pleasing, but physically inaccurate images. Gkioulekas *et al.* have done much work with reconstruction and simulation of scattering media [12, 13]. They have also shown the types of subtle variations possible through changing the phase function, which isn’t captured in approximated models [14].

3 APPROACH

3.1 Experimental Setup

The experimental setup is shown as a schematic in Figure 2 and as a photograph in Figure 3.

A glass container with a square cross-section holds the medium. Light from a laser enters at one edge. A two-dimensional plotter robot moves an occluder through the medium. For each position of the occluder, a camera image is captured (see Figure 4) at another edge of the medium. The whole setup is draped with light-blocking fabric.

For a fixed laser pose and a fixed camera pose, each pixel, $p = (u, v)$, in the camera image corresponds to an output pose at the boundary of the medium. The radiance recorded at p varies with the position (x, y) of the occluder. Let $I_p^w(x, y)$ be the radiance at p when the occluder of width w is at position (x, y) . Let I_p^{absent} be the radiance of p without the occluder present. Then, the




Figure 2: Photon distribution scanner setup.




Figure 3: Experimental setup




Figure 4: A camera image with the occluder in the front middle of the glass container.




Figure 5: The path the two-dimensional plotter robot takes over the medium. The *on* positions for the laser are marked with a solid line, the *off* positions are marked with a dotted line. The occluder is represented by the grey box following the arrow path. There are a total of 128 vertical paths, sampled at 128 positions along each vertical line to generate a 128×128 image.

photon distribution for a given pixel p with an occluder of width w at position (x, y) is

$$D_w^p(x, y) = I_p^{\text{absent}} - I_p^w(x, y).$$

While $D_w^p(x, y)$ is a difference of radiances, it is proportional to the difference of fluxes through p for every (x, y) , as the pixel area and its solid angle are constant. Thus, $D_w^p(x, y)$ is proportional to the flux passing through the occluder at position (x, y) .

The approximate flux density at (x, y) can be obtained by dividing by half of the occluder's surface area inside the volume, excluding the bottom surface of the occluder, which scrapes along the bottom of the container. As this surface area is constant, $D_w^p(x, y)$ is also approximately proportional to the flux density at (x, y) .

We will discuss using the raw data $I_p^w(x, y)$ to obtain the most probable photon paths in Section 3.2. Our path-based approach has the novelty of determining the photon distributions even for regions of the medium that are smaller than the occluder used in the experiment, yielding more accuracy in reconstruction or rendering. Example distribution images are shown in Section 4.1.

3.1.1 Scan Process

To accelerate the data acquisition, the camera captures a video during the experiment. While the camera records, the plotter moves continuously in the y direction from the start boundary to the end boundary. Then the plotter moves in the x axis by a small amount depending on the desired capture resolution, and the process repeats. This is illustrated in Figure 5. The laser is turned off between each continuous movement. Doing so facilitates the later (automatic) processing of the video.

3.2 Iterative Path Algorithm

Scanning the medium as described above allows us to use various occluder widths. However, a better signal-to-noise ratio is obtained with a larger occluder, since the measured difference between the occluder and no occluder is farther from the noise floor of the camera. The size of the occluder that maximizes the signal-to-noise is experiment-dependent. Performing a scan using a small occluder may obtain a better resolution, but we wish to obtain the same resolution with larger occluders. Thus, we propose a novel algorithm that obtains the most probable photon paths, and is relatively insensitive to the occluder size.

From these paths we can generate a simulated **absorption image** $I_p^w(x, y)$ for any occluder width w . To clarify, the absorption image $I_p^w(x, y)$ is not an image captured by the camera, nor is it a photon distribution image, but is an array of values of the intensity difference seen by the camera through a particular pixel, p , for different occluder locations, (x, y) .

Let the entry and exit poses be denoted by $\Pi_0 = (\mathbf{x}_0, \mathbf{d}_0)$ and $\Pi_1 = (\mathbf{x}_1, \mathbf{d}_1)$ with positions \mathbf{x}_0 and \mathbf{x}_1 , and directions \mathbf{d}_0 and \mathbf{d}_1 . Photons travel in many piecewise linear paths from pose Π_0 to Π_1 , which can be simulated using the radiative transfer equation. However, for better performance, we assume that the most probable photon paths that match the light's overall physical behaviour are smooth. This approximation yields reasonable results, as discussed in Section 4.2. We propose an iterative approach, where we start with an initial set of paths and generate a simulated absorption image. Then we compare the simulated absorption image with the corresponding physically measured absorption image, and update the paths to reduce the error between the two. This process is repeated until we converge on a solution. This approach is similar to active contour models, but solves for many contours at once [15].

3.2.1 Template Paths

We generate an initial set of smooth template paths with boundary points \mathbf{x}_0 and \mathbf{x}_1 as follows. Let the midpoint of the line connecting \mathbf{x}_0 and \mathbf{x}_1 be $\mathbf{M} = \frac{1}{2}(\mathbf{x}_0 + \mathbf{x}_1)$. Let \mathbf{n} be the unit normal of this line. Let d_{\min} be the distance from the midpoint to the closest medium boundary along \mathbf{n} .

The template paths are generated by fitting a quadratic curve between \mathbf{x}_0 , a middle point $\mathbf{M}_i = d_i \mathbf{n} + \mathbf{M}$, and endpoint \mathbf{x}_1 for each path i and $d_{\text{path}}^i \in [-d_{\min}, d_{\min}]$. If there are paths generated that leave the medium, we remove the same number of paths on both sides such that they are contained in the medium. The distances, d_{path}^i , are sampled using a Gaussian distribution so that they are more dense near zero. The analytical distribution curves are also quadratic, of the form

$$y \approx (2 \ln 2) \frac{x(r_{\text{prob}} - x)}{\kappa r_{\text{prob}}} \quad (1)$$

for a given y position, radius of the distribution probability r_{prob} , and the inverse diffusive absorption distance $\kappa = \sqrt{3\sigma_a\sigma_s(1-g)}$ [3]. Since the material parameters are unknown, so we set the quadratic curves to be initially non-intersecting in order to allow the warp to move the curves into the correct positions.

This is depicted in Figure 6. These paths do not intersect except at the fixed entry and exit points. This is necessary to prevent degenerate cases where overlapping paths cannot be separated with a warp. These template paths will be non-linearly warped at each path point such that the absorption image simulated from these paths matches the captured absorption image, as described in the following sections.

3.2.2 Distribution Image Generation

Let the simulated absorption image for the set of paths \mathbb{P} and known occluder width w be $G_{\mathbb{P}}^w(x, y)$. Given a




Figure 6: The template paths chosen for our iterative path algorithm given the start and end points \mathbf{x}_0 and \mathbf{x}_1 respectively. The dashed lines indicate the paths that do not contribute to the simulated output intensity at \mathbf{x}_1 , as they get absorbed by the occluder.

set of paths, to generate an absorption image we calculate the percentage of paths that are occluded by an occluder of width w centered at each position (x, y) in the medium. These positions match the distribution scanner positions as discussed earlier. This is a very efficient way to approximate the radiative transfer simulation, as it is simply an intersection test between every path and a square. We will compare this simulated absorption image with a captured absorption image with the same occluder width. This approach allows us to create photon distributions with a finer resolution than the original captured occluder width.

3.2.3 Warp Calculation

Given a target absorption image $I_p^w(x, y)$ for a known camera pixel, p , and occluder width, w , and set of paths, \mathbb{P} , we minimize $\|I_p^w(x, y) - G_{\mathbb{P}}^w(x, y)\|$ through warping \mathbb{P} . A warp field is defined as a vector field $\mathbf{v}(x, y) = (s(x, y), t(x, y))$ that encodes a displacement vector \mathbf{v} for every (x, y) position in the image. We evaluated three different approaches to generate a vector field between two absorption images: the gradient of the difference image $\nabla(I_p^w(x, y) - G_{\mathbb{P}}^w(x, y))$, the Lucas-Kanade optical flow field from $G_{\mathbb{P}}^w(x, y)$ to $I_p^w(x, y)$, and a combination of the two by alternating every iteration.

The convergence for the three methods is shown in Figure 7. The optical flow technique does not perform very well. The combination optical flow/gradient technique converges faster, but converges to a result with a higher error. The gradient approach performs the best, so we use this approach for the remainder of the paper.

3.2.4 Path Warping

For each path point \mathbf{x}_k over all paths, bilinear interpolation is applied to the four closest vectors in the warp field to determine the vector $\mathbf{v}_k = (s_k, t_k)$ at the given point. Then \mathbf{v}_k is projected onto the path's normal \mathbf{n}_k at point \mathbf{x}_k , and update the point

$$\mathbf{x}_k \leftarrow \mathbf{x}_k + h \cdot \text{proj}_{\mathbf{n}_k} \mathbf{v}_k, \quad (2)$$




Figure 7: The root mean squared error (*RMSE*) between simulated absorption image $G_{\mathbb{P}}^w(x, y)$ at iteration i and target absorption image $I_p^w(x, y)$ using a gradient approach, using Lucas-Kanade optical flow (LK), and using an alternating gradient and Lukas-Kanade optical flow approach. The *RMSE* is unitless as it is a difference of percentages.




Figure 8: Sampling the vector field \mathbf{v} to apply a warp to the path point \mathbf{x}_k . Then \mathbf{v}_k is projected onto the path's normal \mathbf{n}_k .

for a small $h > 0$. For a gradient field, the displacement equals a fraction of the directional derivative in the direction \mathbf{n}_k . This step is visualized in Figure 8.

3.2.5 Constraints

Applying the warp may introduce undesirable properties such as intersecting paths, non-smoothness, or broken boundary conditions. We introduce constraints to remove each of these properties as follows.

To force the paths not to intersect, the intersection points are projected outwards from the middle path, as follows. The middle path is compared to one of its neighbouring paths. The intersection points between the two paths is determined (if it exists). If any intersection points exist, the path points of the outer path are projected away from the middle path by some small distance. This process is repeated outwards until the boundary path for that side is reached. The same process is repeated for the other neighbouring side.

To obtain the correct direction at the entrance and exit poses, the pose positions are set, then nearby points are forced in line with the pose direction.

Paths are smoothed using a Gaussian filter. Finally, the path points are resampled at regular intervals using cubic splines. This ensures a good sampling for the next iteration of the warp algorithm. The whole process is repeated multiple times. This is similar to constraints using Verlet integration and other physics solvers [16].

3.2.6 Convergence

The above steps are repeated until the difference in error between consecutive iterations falls below a threshold. The algorithm outputs the set of paths from the last iteration. Finally, a *photon distribution image* is generated by counting the paths that intersect each pixel in that image.

4 RESULTS

4.1 Raw Capture Results for Milk-Water Solutions

Images were captured as video with an Olympus E-M1 Camera at 30 frames per second. The camera used a lens at 60mm focal length, f/2.8 aperture, and 1600 *ISO*. In order to capture these dark images, a high *ISO* and open aperture were required. The camera was focussed at the boundary between the solution and the glass. We assumed the amount of light arriving at the camera from the medium due to refraction and reflection events was minimal.

The distribution scan was performed for multiple different milk-water solutions. Figure 9, shows results from a 75% milk-water solution (that is: 75% milk and 25% water). The figure shows different scans for different output positions on the left edge and a fixed input position on the right edge. Figure 10 shows a similar image, except the camera is now on the bottom edge. Different distributions are shown for varying output pixel positions. Figure 11 shows the result for varying milk-water dilution ratios. All ratios use the same output pixel position on the left of each distribution. The intensity at each pixel is the amount of light absorbed by the occluder for the position in physical space. The darkest two regions occur where the occluder is blocking all the light. One position is in front of the pixel in question which blocks all exiting light. The other position is where the laser enters the solution, blocking all incoming light.




Figure 9: Distribution scans for varying pixel positions at 75% milk-water solution with the laser (right side) on the opposite side as the camera (left side).

4.2 Validation of Algorithm

The algorithm was validated through a simulated absorption image as well as captured absorption images from the scanner. For each run of the algorithm, 512 photon





Figure 10: Distribution scans for varying pixel positions at 75% milk-water solution with the laser (right side) on the adjacent side to the camera (bottom side).



(a) 12.5% milk-water dilution (b) 50% milk-water dilution (c) 100% milk-water dilution

Figure 11: Distribution scans for a single output pixel position with varying milk-water solutions. The laser (right side) is on the opposite side of the camera (left side).

paths were used in the initial and subsequent iterations. For visualization here, only a subset of these paths are shown.

Figure 12 depicts the output of the algorithm on a *virtual* set of target paths using a 10mm occluder. These target paths have a nonlinear pinching closer towards the end points to approximate a narrow band phase function. This provides a set of paths that incorporates pose orientation, as well as position. The algorithm performs quite well when given an image generated by a set of these known paths. Even though the error using the template paths is large, it converges on a valid result, as the final simulated absorption image is similar to the “ground truth” captured absorption image. Also, the root mean squared error (*RMSE*) shows a large improvement, from 0.114 with an image generated from the template paths to 0.014 with the final warped paths. This test was repeated with a 2mm occluder and obtained a similar *RMSE* improvement, from 0.18 to 0.03.

Figures 13 depict examples of the output of the algorithm with captured absorption images *from the physical experiment*. The three datasets are a 50% milk-water solution with output pose straight across from the laser, output pose at an offset from the laser on the opposite side, and a 75% milk-water solution with the output pose on the adjacent side relative to the laser. We denote these as the *straight*, *offset*, and *side* experiments respectively. Note that for each of the three experiments, the 2mm occluder captured images has a lower background intensity than using the 10mm occluder. This suggests that the 10mm captured images yields a higher signal-to-noise ratio since there should

be a uniform intensity outside the photon distribution boundary.

The errors between the reconstructed images and the captured ones are depicted in the fourth row of Figure 13. The template and final errors are shown in Table 1. The template errors correspond to the starting template paths. The algorithm works well for the virtual captures with both occluders. For physical captures, the algorithm performs better using the 10mm than the 2mm occluder captures. This is expected as the 2mm occluder capture images have more noise, hence the active contour algorithm could not find a better solution.

Other deviations may be caused by stray light, background noise, reflections and other physical phenomena not captured by our model. The diffusion model used also assumes that the medium is semi-infinite or infinite. Thus the paths derived do not interfere with the boundaries.

The fifth row of Figure 13 depicts the final resulting photon distributions generated with a 1mm occluder for a 128×128 resolution image. The scale is modified to be between 0 and 0.2 to enhance the darker areas.

Table 1: The final root mean squared error (*RMSE*) between capture and simulated reconstruction images. The *RMSE* is unitless, corresponding to the fraction of the flux over the medium arriving at the output pose.

| Experiment Type | Occluder Size (mm) | Template <i>RMSE</i> | Final <i>RMSE</i> |
|-----------------|--------------------|----------------------|-------------------|
| Virtual Offset | 2 | 0.180 | 0.030 |
| Virtual Offset | 10 | 0.114 | 0.014 |
| Straight | 2 | 0.078 | 0.078 |
| Offset | 2 | 0.070 | 0.069 |
| Side | 2 | 0.101 | 0.100 |
| Straight | 10 | 0.062 | 0.023 |
| Offset | 10 | 0.125 | 0.043 |
| Side | 10 | 0.063 | 0.016 |

5 CONCLUSION

In this experimental setup, we have shown that it is possible to scan a scattering medium with a relatively large occluder, in order to obtain photon distributions through the medium. We have also presented an iterative path algorithm that can efficiently determine a set of photon paths that produce a simulated absorption image that matches the captured absorption image.

It may be interesting to use multiple occluder sizes and different volumes of liquid. This may be useful to determine the limit for the flux density as the occluder size and liquid depth approach zero. Furthermore, a translucent occluder also might be a better posed problem in




Figure 12: Reconstruction from a *virtual* experiment. (a) the “ground truth” photon paths in the virtual medium. (b) the simulated absorption image from the ground truth paths. (c) the initial template paths. (d) the absorption image from the initial template paths. (e) the final computed paths. (f) the absorption image from the final computed paths.

terms of reconstruction, as each path may be able to encode the percentage of intensity absorbed instead of a binary amount.

Photon distributions are a generalization of the matrix used for *ART* and *SIRT*. It would be interesting to see how well the distributions perform using the linear reconstruction techniques for a non-linear simulated phenomena such as scattering. In addition to reconstruction, it would be useful to determine the accuracy of using the photon paths for rendering simulation. An example would be to use them to do efficient approximations for occluders in scattering media. The path shapes could determine the type and amount of scattering whereas the path attenuations could depend on the medium in question. This is an approximation of the Monte-Carlo process, but it would be beneficial to validate the performance versus accuracy tradeoffs.

6 REFERENCES

- [1] R. Gordon, R. Bender, and G. T. Herman, “Algebraic reconstruction techniques (art) for three-dimensional electron microscopy and x-ray pho-

- tography,” *Journal of theoretical Biology*, vol. 29, no. 3, pp. 471–481, 1970.
- [2] A. H. Andersen and A. C. Kak, “Simultaneous algebraic reconstruction technique (sart): a superior implementation of the art algorithm,” *Ultrasonic imaging*, vol. 6, no. 1, pp. 81–94, 1984.
- [3] S. C. Feng, F. Zeng, and B. Chance, “Monte Carlo simulations of photon migration path distributions in multiple scattering media,” in *Photon Migration and Imaging in Random Media and Tissues*, vol. 1888, pp. 78–90, International Society for Optics and Photonics, 1993.
- [4] M. S. Patterson, B. Chance, and B. C. Wilson, “Time resolved reflectance and transmittance for the noninvasive measurement of tissue optical properties,” *Applied optics*, vol. 28, no. 12, pp. 2331–2336, 1989.
- [5] W. Cui, C. Kumar, and B. Chance, “Experimental study of migration depth for the photons measured at sample surface,” in *Time-Resolved Spectroscopy and Imaging of Tissues*, vol. 1431, pp. 180–191, International Society for Optics and Photonics, 1991.
- [6] S. Feng, F.-A. Zeng, and B. Chance, “Photon migration in the presence of a single defect: a perturbation analysis,” *Applied optics*, vol. 34, no. 19, pp. 3826–3837, 1995.
- [7] S. Colak, D. Papaioannou, M. van der Mark, H. Schomberg, J. Paasschens, J. Melissen, N. Van Asten, *et al.*, “Tomographic image reconstruction from optical projections in light-diffusing media,” *Applied optics*, vol. 36, no. 1, pp. 180–213, 1997.
- [8] S. R. Arridge, “Optical tomography in medical imaging,” *Inverse problems*, vol. 15, no. 2, p. R41, 1999.
- [9] S. Green, “Real-time approximations to subsurface scattering,” *GPU Gems*, vol. 1, pp. 263–278, 2004.
- [10] H. W. Jensen, S. R. Marschner, M. Levoy, and P. Hanrahan, “A practical model for subsurface light transport,” in *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pp. 511–518, 2001.
- [11] B. Langsdorf, “Gpu programming exposed: The naked truth behind NVIDIA’s demos,” tech. rep., Technical report, NVIDIA Corporation, 2005.
- [12] I. Gkioulekas, S. Zhao, K. Bala, T. Zickler, and A. Levin, “Inverse volume rendering with material dictionaries,” *ACM Transactions on Graphics (TOG)*, vol. 32, no. 6, pp. 1–13, 2013.
- [13] I. Gkioulekas, A. Levin, and T. Zickler, “An evaluation of computational imaging techniques for




Figure 13: The *straight*, *offset*, and *side* experiments for 2mm and 10mm occluders. Depicted are the raw captured absorption images, reconstructed paths, simulated absorption images, error absorption images, and final photon distribution images. The light enters the left side and exits the right side for the straight and offset experiments. The light enters the right side and exits the top side for the side experiments.

heterogeneous inverse scattering,” in *European Conference on Computer Vision*, pp. 685–701, Springer, 2016.

[14] I. Gkioulekas, B. Xiao, S. Zhao, E. H. Adelson, T. Zickler, and K. Bala, “Understanding the role of phase function in translucent appearance,” *ACM Transactions on graphics (TOG)*, vol. 32, no. 5, pp. 1–19, 2013.

[15] M. Kass, A. Witkin, and D. Terzopoulos, “Snakes: Active contour models,” *International Journal of Computer Vision*, vol. 1, no. 4, pp. 321–331, 1988.

[16] L. Verlet, “Computer” experiments” on classical fluids. i. thermodynamical properties of lennard-jones molecules,” *Physical Review*, vol. 159, no. 1, p. 98, 1967.