

EDGE-RESOLVED TRANSIENT IMAGING: PERFORMANCE ANALYSES, OPTIMIZATIONS, AND SIMULATIONS

Charles Saunders¹, William Krska¹, Julián Tachella², Sheila W. Seidel¹, Joshua Rapp³,
John Murray-Bruce⁴, Yoann Altmann⁵, Stephen McLaughlin⁵, and Vivek K Goyal¹

Boston University¹, University of Edinburgh², Stanford University³,
University of South Florida⁴, Heriot-Watt University⁵

ABSTRACT

Edge-resolved transient imaging (ERTI) is a method for non-line-of-sight imaging that combines the use of direct time of flight for measuring distances with the azimuthal angular resolution afforded by a vertical edge occluder. Recently conceived and demonstrated for the first time, no performance analyses or optimizations of ERTI have appeared in published papers. This paper explains how the difficulty of detection of hidden scene objects with ERTI depends on a variety of parameters, including illumination power, acquisition time, ambient light, visible-side reflectivity, hidden-side reflectivity, target range, and target azimuthal angular position. Based on this analysis, optimization of the acquisition process is introduced whereby the illumination dwell times are varied to counteract decreasing signal-to-noise ratio at deeper angles into the hidden volume. Inaccuracy caused by a coaxial approximation is also analyzed and simulated.

Index Terms— direct time-of-flight, non-line-of-sight imaging, single-photon detection, Skellam distribution

1. INTRODUCTION

Non-line-of-sight (NLOS) optical imaging is a rapidly growing field inspired by both serious applications (e.g., sensing for autonomous vehicles) and the delight in defying human perception [1]. Since light scatters in all directions at a rough (i.e., not mirror-like) surface, when looking at a rough surface the contributions have come from all directions. Though this makes NLOS imaging seem impossible, unmixing of light originating at different locations is possible computationally when the set of measurements includes a diverse enough collection of mixtures.

Initial methods for NLOS imaging used hardware for direct time of flight (DTOF) measurement to produce combinations defined by equal light travel distances [2, 3]. The visible-surface point at which a pulsed laser is focused becomes a virtual source, and the visible-surface point at which a single-photon detector is focused becomes a virtual detector. Through varying the locations of virtual sources and virtual detectors, the diversity of measurements could be sufficient for inverse problem solutions. Making virtual sources and virtual detectors almost colocated simplifies the inverse problem and has become the prevalent paradigm [4–6]. Imaging with this approach is limited to the volume opposite to the area scanned by the virtual source and detector, in addition to having resolution and accuracy

This work was supported in part by the US Defense Advanced Research Projects Agency (DARPA) REVEAL Program under contract number HR0011-16-C-0030 and the US National Science Foundation under grant number 1955219.

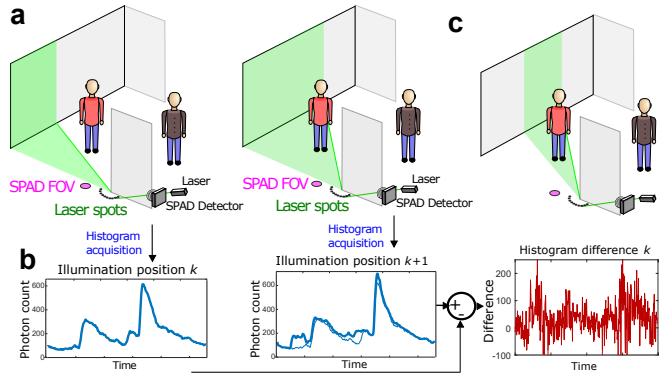


Fig. 1. Imaging scenario. **a.** Positions on a semi-circle on the floor, centered on the occluding wall edge, are illuminated by a pulsed laser. **b.** A histogram of photon arrival times is collected for each laser illumination spot. **c.** The difference between two sequential histograms provides a ‘difference histogram’ containing, on average, photon arrivals from a small wedge in the scene.

limited by the number of virtual sources and detectors and the received signal strength.

A wholly different NLOS imaging paradigm uses only ordinary digital cameras. The presence of known or estimable occluding objects result in shadows cast onto the measurement region, which produce enough diversity such that inversion can be achieved without requiring time-of-flight information [7–14]. The shape and position of the occluder relative to a camera field of view (FOV) determines the computational FOV [12]. For example, the *corner camera* [7] with a single vertical edge makes azimuthal resolution much better than radial resolution [14]. Passive methods have limited resolution, especially for static portions of the hidden volume.

Edge-resolved transient imaging (ERTI) [15] uses aspects of these approaches to achieve NLOS imaging through a small aperture using a small number of illumination positions. The most analogous previous works estimate only moving parts of the hidden scene [16] or projections thereof [17].

2. OVERVIEW OF ERTI

ERTI Data Acquisition. As illustrated in Fig. 1a, a pulsed laser sequentially illuminates points on the floor along a semi-circle with small radius (1.5 cm), centered on the point where a wall edge meets the floor. The laser light reflects in all directions from the Lambertian floor. Some of this light enters the hidden room, reflects off of the

hidden surfaces and returns to the FOV on the floor of the single-photon avalanche diode (SPAD) sensor. Together, the SPAD and a time-correlated single photon counting module record the arrival times of photons reaching the sensor. Arrival times relative to laser pulse emission are histogrammed over many repetitions of the laser (Fig. 1b). Here, a histogram is a vector of length $N = t_r/t_b$ where t_r is the repetition period of the laser, and t_b is the resolution of the timing hardware. Each entry is the number of arriving photons counted within a particular time interval.

Neglecting for now any sources of noise (e.g., shot noise), the histogram measured at illumination point k , denoted \mathbf{m}_k , contains photon count contributions from the hidden scene denoted \mathbf{h}_k . There are also contributions from the laser light reflecting back from any surfaces on the observer's (i.e., visible) side, as well as ambient and dark counts, which we combine together in \mathbf{b} . Note that \mathbf{b} is constant for every illumination position, assuming the radius of the semi-circle is reasonably small. As the position on the semi-circle moves further away from the wall, i.e., k increases, a larger portion of the hidden scene is illuminated. This implies that $\mathbf{h}_{k+1} = \mathbf{h}_k + \mathbf{u}_k$, where \mathbf{u}_k is the component of the histogram contributed by the part of the scene that is illuminated from spot $k+1$ but not from spot k .

Consider taking the difference between the histograms collected from neighboring illumination positions:

$$\mathbf{y}_k = \mathbf{m}_{k+1} - \mathbf{m}_k \approx \mathbf{b} + \sum_{j=0}^k \mathbf{u}_j - \mathbf{b} - \sum_{j=0}^{k-1} \mathbf{u}_j = \mathbf{u}_k. \quad (1)$$

Doing so provides a difference histogram \mathbf{y}_k containing only the response from the part of the hidden scene that was illuminated from spot $k+1$ but not spot k (the green wedge in Fig. 1c), as the rest of the terms cancel. Hence, the difference histograms $\{\mathbf{y}_k\}_{k=1, \dots, n-1}$ represent photons arriving from disjoint wedges in the hidden area.

Reconstruction Procedure. The presence of the occluding edge and the subtraction of histograms provides azimuthal resolution, while the arrival times of the photon detections provide longitudinal resolution within these wedges.

Difference histograms are treated as noisy observations of idealized parametric responses from the hidden scene wedge. Planar surfaces with a certain reflectivity, height and orientation about the floor-to-ceiling axis are fit via parametric estimation to the difference histogram for each wedge separately, thus providing a 2.5-dimensional recovery of the hidden area. An approximate, closed form solution to the temporal response from such surfaces (including occlusions between them) is derived in [15] and allows efficient evaluation of the forward model. The fitting of these planar surfaces is achieved in the original ERTI paper using Bayesian inference through a Markov chain Monte Carlo (MCMC) sampling approach that accounts for prior beliefs about scene structure and only requires evaluations of the forward model [18, 19].

3. DATA MODELING AND ANALYSIS

Let Δ_θ represent the azimuthal angular extent of one wedge in radians, or equivalently the angular spacing of illumination spots. For acquisitions of the full π radians of hidden volume, $\Delta_\theta = \pi/(n-1)$, where n is the number of illumination points. While all our experimental data collection was for this setting, one might use ERTI with attention on a limited region of interest or with variable angular resolution. Analyses and simulations in this section do not include variation of the rotation angles of facets; facet normal vectors are horizontal and point toward the imaging edge.

3.1. Signal strength dependence on range

For a facet of fixed size in linear units (e.g., 1 m by 1 m, contrasting with angular units), assuming Lambertian reflection of the illumination on the floor, the amount of laser light reaching the facet scales as $1/z^2$, where z is the distance between the facet and the illumination spot. Lambertian reflection from the facet then implies another $1/z^2$ scaling for the amount of light that reaches a detector FOV of fixed area near the edge. Hence, signal strength for a fixed-size facet scales as $1/z^4$. Note that this is a total signal strength from the facet, which may be split across multiple angular wedges.

Instead of considering a fixed-size facet, one may also naturally be interested in a vertical facet of fixed height that occupies the full Δ_θ angular extent of a wedge. Then, the amount of laser light reaching the facet scales as $1/z$. The $1/z^2$ solid angle extent of the detector FOV from the vantage point of the facet remains unchanged from the argument of the previous paragraph. Hence, signal strength for a fixed-height facet of full angular extent scales as $1/z^3$.

3.2. Signal-to-noise ratio analysis

For the analysis of this section, we omit height variations and concentrate on variations of signal and clutter strengths with angular position θ at the center of the wedge with index j . Seeking a fundamental scaling law, we consider facet recovery in one wedge (i.e., solved from one histogram pair, which are differenced). While the MCMC method exploits scene priors and joint processing of all data, here we consider one wedge in isolation to simplify our problem to a hypothesis test. Since radial distance is mapped to TOF, we reduce the problem to determination of the presence or absence of a fixed-height facet of full angular extent at one specified radial distance z .

Since histogram entries are independent Poisson random variables, histogram differences have Skellam-distributed entries. We parametrize these Skellam distributions by mean and variance because this is sufficient for the analysis. Our problem is to decide between

$$\text{No facet } H_0: m_j - m_{j-1} \sim \text{Skellam}(0, \sigma_{j-1}^2), \quad \text{and}$$

$$\text{Facet } H_1: m_j - m_{j-1} \sim \text{Skellam}(\mu_j, \sigma_{j-1}^2 + \mu_j),$$

where m_j is one entry of interest in the histogram measured for the j th illumination position, μ_j is the mean counts added by a facet present in wedge j , and σ_{j-1}^2 is the Skellam variance due to ambient light, visible-side scene elements, and hidden-side scene elements up to wedge $j-1$. If the distributions under the two hypotheses had the same variance, performance of the hypothesis test would simply depend on the difference of means after standardization. For simplicity, we take μ_j/σ_{j-1} as the approximate signal-to-noise ratio (SNR) and require it to be held approximately constant to attain approximately constant probability of error. The mean counts μ_j can be expressed as

$$\mu_j = P \Delta_\theta \alpha_j, \quad (2)$$

where P is the product of laser power and dwell time per spot and α_j is the effective reflectivity per azimuthal radian of the facet in wedge j , including distance-based falloff effects. The variance σ_{j-1}^2 can be expressed as

$$\sigma_{j-1}^2 = 2B + 2P \left(V + \Delta_\theta \sum_{k=0}^{j-1} \alpha_k \right), \quad (3)$$

where B represents ambient light contributions, V the visible-side reflections of laser light, and $\{\alpha_k\}_{k=0}^{j-1}$ the contributions of earlier

hidden-side wedges. By substituting (2) and (3), we get

$$\gamma = \frac{\mu_j}{\sqrt{\sigma_{j-1}^2}} = \frac{P\Delta_\theta\alpha_j}{\sqrt{2B + 2P(V + \Delta_\theta \sum_{k=0}^{j-1} \alpha_k)}} \quad (4)$$

as a figure of merit. To more easily draw some insights from this analysis, let us set $B \approx 0$ because optical filtering can make ambient light contributions negligible; set $V = \pi v$, where v represents visible-side reflectivity per azimuthal radian; and set $\alpha_k \approx h$ for $k = 0, 1, \dots, j-1$, representing constant hidden-side reflectivity. Then $\Delta_\theta \sum_{k=0}^{j-1} \alpha_k \approx \theta h$ and

$$\gamma \approx \frac{P\Delta_\theta h}{\sqrt{2P(\pi v + \theta h)}} = \sqrt{P\Delta_\theta} \frac{h}{\sqrt{2(\pi v + \theta h)}}. \quad (5)$$

An actionable use of (5) is to see how P should scale to maintain a constant facet-recovery success probability. Rearranging (5) gives

$$(\text{laser power})(\text{dwell time per spot}) \propto \frac{\theta/\pi + v/h}{\Delta_\theta^2 h}. \quad (6)$$

For example, without changing the laser power, to double the angular resolution (halve Δ_θ) requires $4 \times$ increase of dwell time per illumination spot, resulting in $8 \times$ increase of total acquisition time. Also, we see that increasing dwell time as θ increases would be beneficial.

4. ACQUISITION OPTIMIZATION

For our analysis in this section, we consider an approximately semi-circular room such that in each angular wedge there is a planar surface whose normal points towards the corner, at the same distance from the corner. The measurements from further around the corner hence contain more ‘clutter’, that is, the responses from all the surfaces at previous angles – as they are all localized at the same position in time and hence in the histogram. This is the worst-case scenario in terms of increasing measurement variance with angle. To simplify the derivations, we assume there is no background contribution and we define the visible scene contribution as the same as that of one wedge in the semi-circular hidden room, which we denote $\bar{\mu}$, in order to avoid the special case where the first measurement is entirely zero.

In the original ERTI paper [15], the laser illuminates each point for the same amount of time. Consequently, the later wedges become harder to reconstruct. To alleviate this issue, we consider allowing the dwell time to be different for each illumination point. Given a total dwell time budget T , we consider the resource allocation problem of assigning a dwell time to each illumination point. Optimality can be defined in many ways. Here we will consider:

1. Ensuring constant variance of each difference measurement, so that reconstruction performance does not degrade as the angle increases (as with fixed dwell times).
2. Minimizing sum of the variances of difference measurements.

Begin by denoting the dwell time at the k th illumination point by g_k . For these derivations it suffices to consider a single histogram time bin m_k where there is some response μ_k . Using a Gaussian approximation of the Poisson distribution,

$$m_k \sim \mathcal{N}(g_k \mu_k, g_k \mu_k), \quad (7)$$

and the difference measurement corresponding to the k th wedge is

$$y_k = \frac{m_{k+1}}{g_{k+1}} - \frac{m_k}{g_k} \sim \mathcal{N}\left(\mu_{k+1} - \mu_k, \frac{\mu_k}{g_k} + \frac{\mu_{k+1}}{g_{k+1}}\right).$$

In our semi-circular room, the response in each wedge is the same and is denoted $\bar{\mu}$, hence $\mu_{k+1} = \mu_k + \bar{\mu}$, and $\mu_k = k\bar{\mu}$. Then,

$$\sigma_k^2 = \frac{k\bar{\mu}}{g_k} + \frac{(k+1)\bar{\mu}}{g_{k+1}}. \quad (8)$$

Constant Variance. For the variances to be equal for all difference measurements we can require that $\sigma_k^2 = \sigma_{k+1}^2$:

$$\begin{aligned} \frac{k\bar{\mu}}{g_k} + \frac{(k+1)\bar{\mu}}{g_{k+1}} &= \frac{(k+1)\bar{\mu}}{g_{k+1}} + \frac{(k+2)\bar{\mu}}{g_{k+2}} \\ \Rightarrow g_{k+2} &= g_k \frac{k+2}{k}. \end{aligned} \quad (9)$$

We now have a recurrence relation for a sequence of dwell times that ensures the variance is constant for wedges at all angles. However, the initial conditions for g_1 and g_2 must be determined in order to generate and use the sequence. We can formulate an optimization problem to determine the best values of g_1 and g_2 for a fixed total dwell time T :

$$g_1, g_2 = \arg \min_{g_1, g_2} \frac{\bar{\mu}}{g_1} + \frac{2\bar{\mu}}{g_2} \quad \text{s.t. } \sum_{k=1}^n g_k = T \quad (10)$$

where, from (9),

$$g_k = \begin{cases} g_1 k, & k \text{ odd}; \\ \frac{1}{2} g_2 k, & k \text{ even}. \end{cases}$$

To solve (10), we use a gradient method with the Lagrangian form

$$\arg \min_{g_1, g_2} \frac{\bar{\mu}}{g_1} + \frac{2\bar{\mu}}{g_2} + \lambda \left(\sum_{k=1}^{n/2} (g_1(2k-1)) + \sum_{k=1}^{n/2} (g_2 k) - T \right).$$

Minimizing the Sum of Variances. Instead of maintaining a constant variance among wedges, dwell times can be computed to achieve the minimum sum of the variances of all difference histograms. To do so, one can solve

$$\arg \min_{g_1, \dots, g_n} \sum_{k=1}^{n-1} \left(\frac{k\bar{\mu}}{g_k} + \frac{(k+1)\bar{\mu}}{g_{k+1}} \right) \quad \text{s.t. } \sum_{k=1}^n g_k = T. \quad (11)$$

Fig. 2 shows the dwell times determined by solving (10) and (11) as well as the resulting difference measurement variances and SNR. To achieve constant variance, we find that the dwell time must increase as the angle around the corner increases, to counteract the noise contributed by the surfaces in preceding wedges. This is consistent with the conclusion drawn from (6).

5. DETECTOR OFFSET

The approximate, closed-form solution to the temporal response from an arbitrary planar surface allows for fast evaluation of the forward model, providing the measurements one would collect from any scene comprising such facets with high efficiency. Without this, evaluating the forward model requires significant rendering effort. However, one major assumption made in the derivation of this approximation is that the detector FOV is a point coinciding with the corner itself. This is problematic in practice because it would require extremely effective temporal gating, and the FOV needs to be able to receive light from the entire hidden scene. In the experiments in [15], the detector FOV is a small spot located about 20 cm away

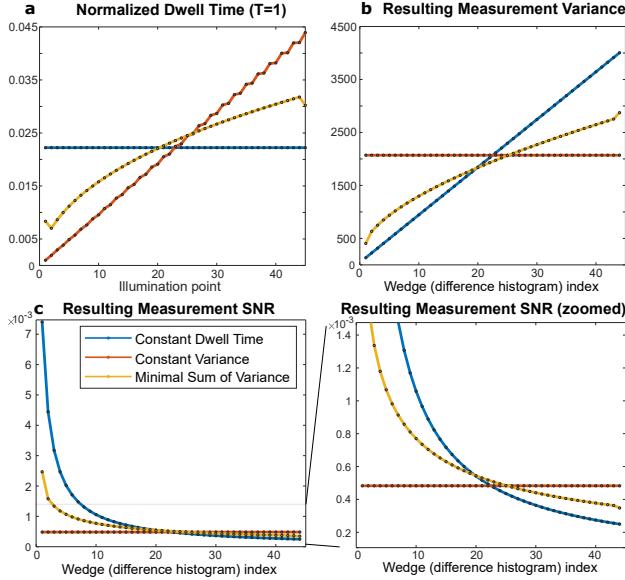


Fig. 2. Measurement variance due to different dwell time allotments. **a.** Proposed dwell time allotment schemes that sum up to a total dwell time of 1, for 45 illumination points and a semi-circular room with arbitrary $\bar{\mu} = 1$. **b.** The variance of the 44 difference measurements resulting from the different schemes. **c.** The SNR of the 44 difference measurements.

from the corner following the direction of the wall. This placement of the virtual detector allows it to detect light from all positions in the hidden scene. The reconstruction algorithm in [15] does not account for this discrepancy, assuming it to be negligible. Nevertheless, there is indeed an error caused by this model mismatch which we explore here.

Denote the vector from the corner to the true SPAD FOV position by \mathbf{d}_s and the vector from the corner to a point on a target surface by \mathbf{d}_t . The round-trip distance assumed in the original reconstruction algorithm is $r = 2\|\mathbf{d}_t\|_2$, whereas the true round-trip distance is given by $\hat{r} = \|\mathbf{d}_t\|_2 + \|\mathbf{d}_s - \mathbf{d}_t\|_2$. This gives a model mismatch error of:

$$e = 2\|\mathbf{d}_t\|_2 - (\|\mathbf{d}_t\|_2 + \|\mathbf{d}_s - \mathbf{d}_t\|_2) = \|\mathbf{d}_t\|_2 - \|\mathbf{d}_s - \mathbf{d}_t\|_2.$$

The error hence depends on the position of the target in the room, and the magnitude of the error is also upper-bounded by $\|\mathbf{d}_s\|_2$ from the triangle inequality. Fig. 3a shows the error as a heatmap within a hidden area for targets at floor height.

A simple improvement to the recovery algorithm is to compensate after the recovery is formed by subtracting the error to the central point on each estimated surface from the surface's radial distance. Fig. 3b shows the result of the reconstruction of a square, empty room performed using the MCMC algorithm from [15] with the detector FOV 30 cm offset from the corner. Displayed in green is the reconstruction with this compensation applied.

6. EXPERIMENTS

As discussed in Section 3.2, one way to estimate whether or not there is a surface in a particular wedge is considering the difference in Skellam likelihood between the case where there is a surface present (i.e., the mean is nonzero), or the case with zero mean. To determine

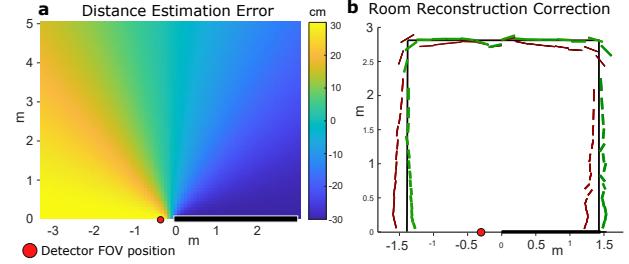


Fig. 3. Reconstruction distance error due to model mismatch and correction. **a.** Distance error at floor level with SPAD FOV 30 cm to the left of the corner. **b.** Bird's-eye view of ground truth room in black, MCMC reconstruction in red, corrected offset error in green.

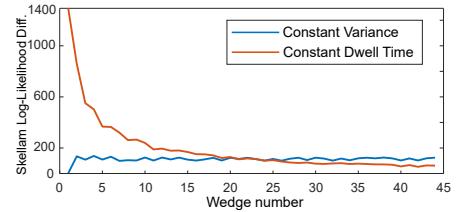


Fig. 4. The difference between the Skellam Log-Likelihood with mean as the response from a surface, or zero. Mean of 1000 trials, results shown for two dwell time schemes.

the effect of dwell-time optimization on such an estimation, we perform the following experiment. For each of 1000 trials, a realization of the noisy histograms that would be measured from a semi-circular room of radius 2 m and wall height 2.5 m is computed. Then, for each difference histogram, the difference between the logarithms of the Skellam likelihoods with mean given by the true underlying surface response with zero mean is calculated. If this value is above zero it implies a surface exists in that wedge, and below zero implies the wedge is empty.

This allows us to empirically confirm that the constant-variance dwell time allotment from the solution to (10) developed in Section 4 provides approximately constant performance between each wedge (i.e., it is equally easy to determine whether a surface is present at every angle). Indeed, we see in Fig. 4 that this is the case. There is a dip in the likelihood difference at the first angle. The dwell time allotted to the first position is very small and hence the number of counts recorded is so low that the Gaussian approximation no longer holds well. Therefore, it may be prudent in practice to include a lower bound on the expected number of total photon counts in the dwell time optimization, if required.

7. CONCLUSION

We have developed theory for performance analysis of edge-resolved transient imaging, aiming to show how the detection of hidden objects depends on various parameters. Furthermore, we propose and demonstrate a revised acquisition method with non-constant dwell times assigned to each illumination position to either minimize the sum of the measurement variances or achieve constant measurement variance for every wedge. We also explore the effect on scene estimation of the model mismatch between the modelled and true position of the detector's FOV.

8. REFERENCES

- [1] D. Faccio, A. Velten, and G. Wetzstein, “Non-line-of-sight imaging,” *Nat. Rev. Phys.*, vol. 2, no. 6, pp. 318–327, Jun. 2020.
- [2] A. Velten, T. Willwacher, O. Gupta, A. Veeraraghavan, M. G. Bawendi, and R. Raskar, “Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging,” *Nat. Commun.*, vol. 3, 2012.
- [3] A. Kirmani, H. Jeelani, V. Montazerhodjat, and V. K. Goyal, “Diffuse imaging: Creating optical images with unfocused time-resolved illumination and sensing,” *IEEE Signal Process. Lett.*, vol. 19, no. 1, pp. 31–34, Jan. 2012.
- [4] M. O’Toole, D. B. Lindell, and G. Wetzstein, “Confocal non-line-of-sight imaging based on the light-cone transform,” *Nature*, vol. 555, pp. 338–341, Mar. 2018.
- [5] D. B. Lindell, G. Wetzstein, and M. O’Toole, “Wave-based non-line-of-sight imaging using fast f - k migration,” *ACM Trans. Graph.*, vol. 38, no. 4, 2019, article 116.
- [6] X. Liu, I. Guillén, M. La Manna, J. H. Nam, S. A. Reza, T. H. Le, A. Jarabo, D. Gutierrez, and A. Velten, “Non-line-of-sight imaging using phasor-field virtual wave optics,” *Nature*, vol. 572, pp. 620–623, 2019.
- [7] K. L. Bouman, V. Ye, A. B. Yedidia, F. Durand, G. W. Wornell, A. Torralba, and W. T. Freeman, “Turning corners into cameras: Principles and methods,” in *Proc. 23rd IEEE Int. Conf. Computer Vision*, 2017, pp. 2270–2278.
- [8] M. Baradad, V. Ye, A. B. Yedidia, F. Durand, W. T. Freeman, G. W. Wornell, and A. Torralba, “Inferring light fields from shadows,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2018, pp. 6267–6275.
- [9] C. Thrampoulidis, G. Shulkind, F. Xu, W. T. Freeman, J. H. Shapiro, A. Torralba, F. N. C. Wong, and G. W. Wornell, “Exploiting occlusion in non-line-of-sight active imaging,” *IEEE Trans. Comput. Imaging*, vol. 4, no. 3, pp. 419–431, Sep. 2018.
- [10] F. Xu, G. Shulkind, C. Thrampoulidis, J. H. Shapiro, A. Torralba, F. N. C. Wong, and G. W. Wornell, “Revealing hidden scenes by photon-efficient occlusion-based opportunistic active imaging,” *Opt. Express*, vol. 26, no. 8, pp. 9945–9962, Apr. 2018.
- [11] S. W. Seidel, Y. Ma, J. Murray-Bruce, C. Saunders, W. T. Freeman, C. Yu, and V. K. Goyal, “Corner occluder computational periscopy: Estimating a hidden scene from a single photograph,” in *Proc. IEEE Int. Conf. Comput. Photography*, Tokyo, Japan, May 2019, pp. 25–33.
- [12] C. Saunders, J. Murray-Bruce, and V. K. Goyal, “Computational periscopy with an ordinary digital camera,” *Nature*, vol. 565, pp. 472–475, 24 Jan. 2019.
- [13] A. B. Yedidia, M. Baradad, C. Thrampoulidis, W. T. Freeman, and G. W. Wornell, “Using unknown occluders to recover hidden scenes,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2019, pp. 12231–12239.
- [14] S. W. Seidel, J. Murray-Bruce, Y. Ma, W. T. Freeman, C. C. Yu, and V. K. Goyal, “Two-dimensional non-line-of-sight scene estimation from a single edge occluder,” *IEEE Trans. Comput. Imaging*, vol. 7, pp. 58–72, 2021.
- [15] J. Rapp, C. Saunders, J. Tachella, J. Murray-Bruce, Y. Altmann, J.-Y. Tourneret, S. McLaughlin, R. M. A. Dawson, F. N. C. Wong, and V. K. Goyal, “Seeing around corners with edge-resolved transient imaging,” *Nat. Commun.*, vol. 11, no. 5929, Nov. 2020.
- [16] G. Gariepy, F. Tonolini, R. Henderson, J. Leach, and D. Faccio, “Detection and tracking of moving objects hidden from view,” *Nature Photon.*, vol. 10, no. 1, pp. 23–26, 2016.
- [17] C. A. Metzler, D. B. Lindell, and G. Wetzstein, “Keyhole imaging: Non-line-of-sight imaging and tracking of moving objects along a single optical path,” *IEEE Trans. Comput. Imaging*, vol. 7, 2021.
- [18] S. Brooks, A. Gelman, G. Jones, and X.-L. Meng, *Handbook of Markov chain Monte Carlo*. CRC press, 2011.
- [19] J. Tachella, Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, S. McLaughlin, and J.-Y. Tourneret, “Bayesian 3D reconstruction of complex scenes from single-photon lidar data,” *SIAM J. Imaging Sci.*, vol. 12, no. 1, pp. 521–550, Jan. 2019.