Visible Occluders as Opportunistic Apertures for Wide Field of View Non-line-of-sight 3D Imaging

Robinson Czajkowski and John Murray-Bruce
Bellini College of Artificial Intelligence, Cybersecurity and Computing, University of South Florida, Tampa, FL.
{czajkowski,murraybruce}@usf.edu

***Abstract*—In indoor environments, walls often limit visibility. For example, an observer outside a room has a limited view due to the surrounding walls. We illustrate that this observer, equipped with a digital camera, can opportunistically exploit the edges of the room's doorway—as a computational aperture—to create a 3D 180-degree (horizontal) field of view reconstruction of the interior without ever gaining a direct line of sight or using active illumination. Like other passive, occluder-aided non-line-of-sight imaging methods, our approach analyzes an intensity photograph of the shadow cast by the occluder. Still, the state-of-the-art is limited to 3D reconstructions that assume few hidden scene objects, each confined to a single depth. However, our approach enables far greater depth resolution, permitting accurate reconstructions even when hidden scene objects span substantial depth extents. To achieve this, we exploit all three edges of the doorway instead of just one or two orthogonal edges. This extension requires innovations in hidden scene discretization and representation based on Fisher information analyses, as well as advances in reconstruction algorithms to enable coherent, convergent, and accurate 3D imaging of real experimental scenes from ordinary 2D photographs.**

***Index Terms*—Non-line-of-sight imaging, computational photography, corner camera, Fisher information, information orthogonality**

I. Introduction

Non-line-of-sight (NLOS) imaging seeks to understand objects and scenes that are outside the observer's field of view (FOV). Without exploitable mirrors nearby, an emerging approach is to capture the intensity variations of a visible matte surface diffusely illuminated by the hidden scene.

From this, the ultimate goal is a final representation of the hidden scene that gives sufficient detail and high accuracy. This goal is challenging because the problem is ill-posed: light from the hidden scene scatters indiscriminately in all directions and combines at the visible surface. Thus, the relation between the hidden scene and the measured visible surface illumination is extremely weak. Moreover, seeking a 3D reconstruction from a single 2D photograph further exacerbates the difficulty. Prior works have demonstrated many optical techniques to address these challenges. These techniques fall into two categories: *(i) active methods* that measure back-reflected light from indirect illumination of the hidden scene using a controlled laser source, and *passive methods* that measure existing illumination on a nearby surface without using any controlled light source. In the former, the scene reconstructions rely on timing the roundtrip travel of the laser pulse bouncing off the observation surface, hitting the hidden scene object, and bouncing back off the observation surface for several illumination or detection positions [1]–[8]. Thus, they rely on expensive equipment and can have long acquisition times. In contrast, passive NLOS techniques are fast and use simple, cheap equipment. Accordingly, advancements in *passive* NLOS have typically involved exploiting occluders as functional apertures that encode information of the hidden scene in the shadows they project on a nearby visible surface [9]–[27].

Bouman et al. [10] first proposed using an edge occluder to estimate a 1D angular video of objects moving in the hidden scene. Seidel et al. extended this approach to enable 2D plan view imaging of a stationary scene relying on a polar coordinate representation of the hidden scene [16]. In the polar representation, the azimuthal resolution came from exploiting the vertical edge, while coarse range resolution came from subtle variations in intensity caused by radial falloff. Subsequent work combined two vertical edges to improve range accuracy in 2D plan view reconstructions [28]. Other passive, occluder-aided NLOS setups exploit occluders that are part of the hidden region. In such settings, the occluder is typically assumed to be known [13], [14], [23] or unknown and to be reconstructed [15], [26], [27]. Occluders have also been used in active imaging to improve reconstructions [6]–[8].

This work builds on previous work [29], which reconstructed three-dimensional scenes by exploiting an occluder with two orthogonal edges. In this case, each edge provides a mechanism for 2D angular resolution pivoting about the edge, while radial falloff once more provides weak range information. Thus, a two-step algorithm was used: estimating shape (parameterized by angles) in the first stage, followed by estimating ranges for a few clusters of pixels in the second stage, to account for the ill-conditioned range dimension and only estimated one range for each object in the scene. Unfortunately, this approach is limited to scenes where objects are well separated and each object is approximately at a single depth. Here, we exploit a third edge to provide resolution along the third dimension, enabling estimation of the entire 3D scene in a single step. Critically, our approach highlights the importance of coupling the hidden scene representation to the angular directions along which each occluding edge provides resolution. Representing and discretizing the hidden scene in conventional Cartesian coordinates fails, whereas our *information orthogonal* coordinate system yields accurate reconstructions.

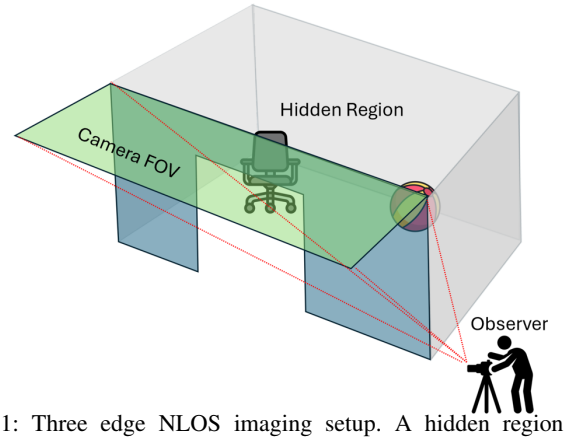

Fig. 1: Three edge NLOS imaging setup. A hidden region with unknown objects can be imaged by an observer who cannot see directly into the room. The observer captures an image of the ceiling, shown in green. The camera FOV contains shadows of all three edges of the doorframe.

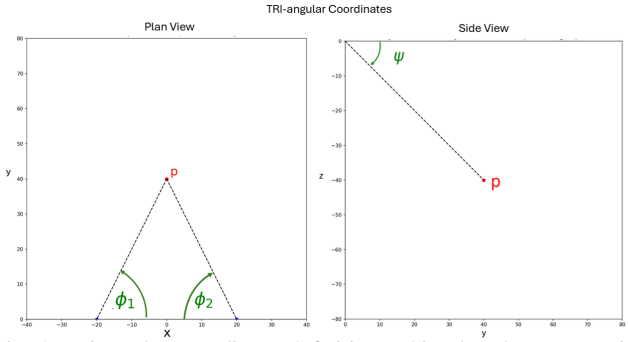

Fig. 2: Triangular coordinate definition. This plot shows two views of the 3D scene, from the top and side. A point p can be represented by Cartesian (x, y, z) or triangular (ϕ_1, ϕ_2, ψ) coordinates.

We organize the remainder of this manuscript as follows. Section II shows the formulation of the information orthogonal coordinate system, the computation of the physical model, and the method of solving the inverse problem. In Section III, we show a Cramer-Rao Bound analysis justifying our use of a three-edge occluder and present results of our reconstruction algorithm on various real experimental scenes. Finally, Section IV provides a brief discussion and concludes the paper.

II. Forward & Inverse Problem

For the NLOS imaging scenario of interest, Figure 1 shows the hidden region, observer, camera FOV, and three-edge doorway occluder. Objects in the hidden scene emit or reflect light in all directions; the walls block some light paths and let others through, creating a soft shadow of the doorway on the ceiling. The soft shadow, or penumbra, has an approximately trapezoidal shape with soft edges. An example is shown in Figure 3 (left). Each edge of the trapezoidal shadow corresponds to an occluding edge of the doorway, encoding information regarding the object's relative angular position. We demonstrate the methodology for extracting this information for complex scenes with multiple objects by formulating an imaging inverse problem.

A. Triangular Coordinate System

We parameterize the hidden region with coordinates (ϕ_1, ϕ_2, ψ) where ϕ_1 and ϕ_2 is the angle about the left and right edges respectively, and ψ is the angle that pivots about the horizontal edge. These coordinates are illustrated in Figure 2 for a point p in the hidden scene. This triangular coordinate system draws inspiration from the coordinate systems of [28], [29], which also employ a novel coordinate system that matches the occluder shape. The Fisher information (FI) matrix indicates the information in a penumbra photograph for each parameter and any correlation among them (in the off-diagonal terms). The FI matrix for estimating the position in Cartesian coordinates of a hidden scene target has appreciable off-diagonal terms, meaning the x, y, and z coordinates are strongly correlated. However, the FI matrix in a triangular coordinate parameterization is approximately diagonal, meaning the estimates are uncorrelated. Thus, errors in one parameter have a negligible impact on the others.

Converting to this *triangular* coordinate system from Cartesian is achieved via: $\tan(\phi_1) = y/(x+d)$, $\tan(\phi_2) = y/(x-d)$ and $\tan(\psi) = y/z$.

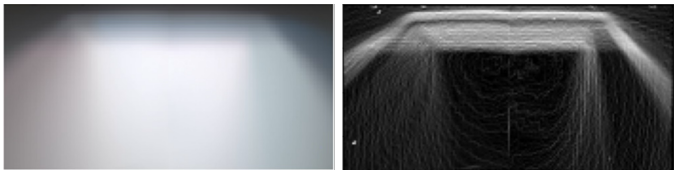

Fig. 3: Observed penumbra photograph $\mathbf{y}$ (left) and its absolute derivative $|\mathbf{D}\mathbf{y}|$ (right). The derivative image accentuates penumbra contributions and suppresses approximately constant background.

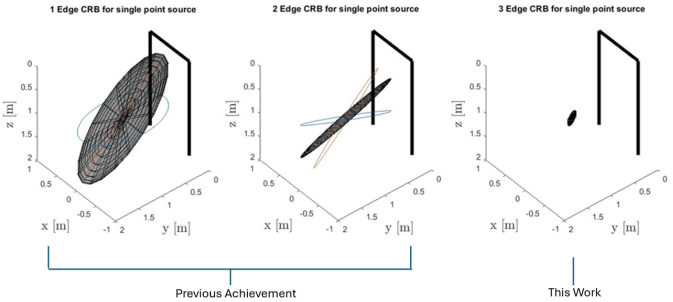

Fig. 4: Cramer-Rao Bound uncertainty regions for a point in the middle of the hidden scene. Computed for a one-edge occluder on the door head (left), a right and top edge occluder (middle), and all three edges (right).

B. Discrete Forward Model

The forward model is computed by simulating voxels at evenly spaced (ϕ_1, ϕ_2, ψ) locations in the hidden scene. We exclude degenerate positions from our discretization. For a voxel V with center v_0, the light contribution to an observation pixel P with center p_0 and size $k \times k$ is evaluated as

$$L(V,P) = cR(v_0,p_0)I(v_0,p_0)\int_{v\in V}\int_{p\in P} H(v,p)J\,\mathrm{d}p\,\mathrm{d}v, \quad (1)$$

where c is the brightness of the voxel, R models the radial falloff due to the distance between V and P, and J is the Jacobian for converting from Cartesian to triangular

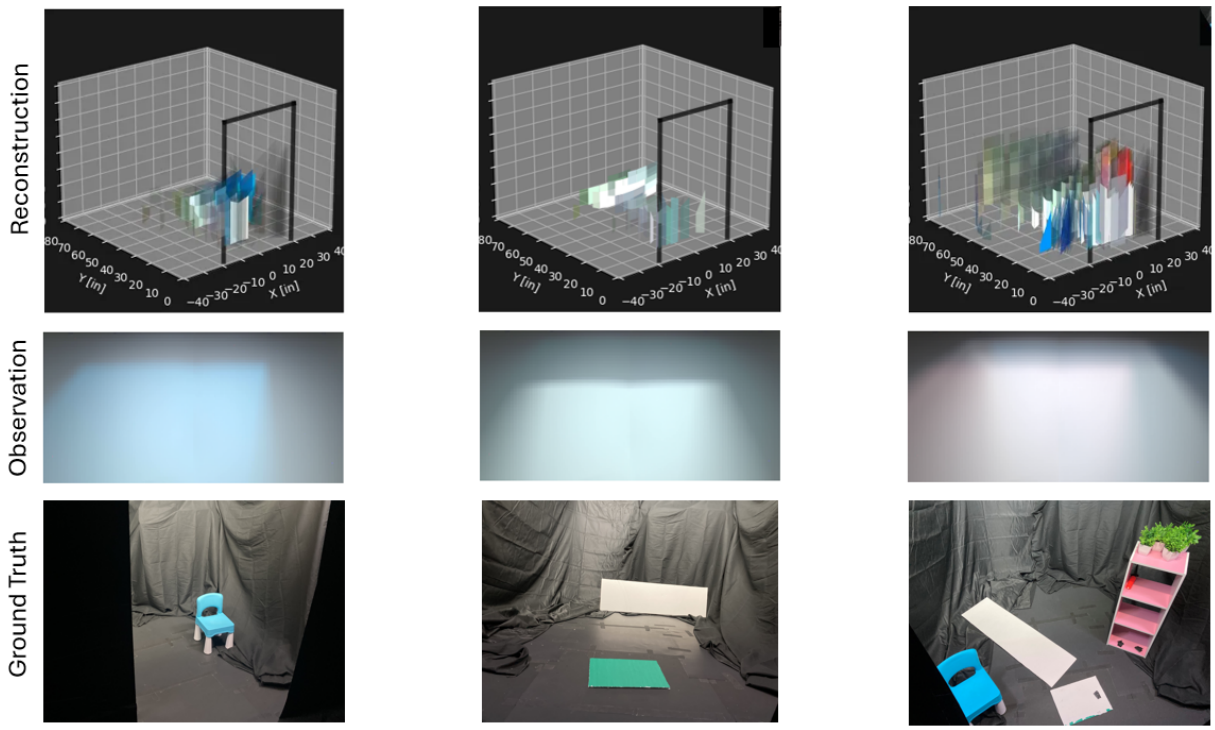

Fig. 5: Reconstructions of three scenes. Each reconstruction is a 36 by 36 by 18 discretization. The ground truth scene is shown in the top right of each plot, and the observation image of the ceiling captured is shown below.

coordinates and is omitted here because of its length. H represents the light transport matrix between the hidden scene and the ceiling, returning zero if the light path between v and p is blocked and one otherwise. By approximating P as an infinitely small patch with its center p_0, we have

$$L = ck^2 R(v_0, p_0) I(v_0, p_0) \int_{\phi_1^0}^{\phi_1^0} \int_{\phi_2^1}^{\phi_2^1} \int_{\psi^0}^{\psi^1} \left(\phi_1 - \operatorname{atan}\frac{b}{a+d}\right) \left(\phi_2 - \operatorname{atan}\frac{b}{a-d}\right) \left(\psi - \operatorname{atan}\frac{b}{h}\right) J \mathrm{d}\psi \mathrm{d}\phi_2 \mathrm{d}\phi_1, \quad (2)$$

where $p_0 = (a, -b, -h)$, h is the distance between the doorway head and ceiling, and $2d$ is the width of the doorway opening. For small fine discretizations, V has approximately straight edges, which gives the approximation:

$$L \approx ck^2 R(v_0, p_0) I(v_0, p_0) \int_{\phi_1^0}^{\phi_1^0} \left(\phi_1 - \operatorname{atan}\frac{b}{a+d}\right) \mathrm{d}\phi_1 \cdot \int_{\phi_2^1}^{\phi_2^1} \left(\phi_2 - \operatorname{atan}\frac{b}{a-d}\right) \mathrm{d}\phi_2 \cdot \int_{\psi^0}^{\psi^1} \left(\psi - \operatorname{atan}\frac{b}{h}\right) \mathrm{d}\psi, \quad (3)$$

$$L = ck^2 R(v_0, p_0) I(v_0, p_0) U\left(\phi_1 - \operatorname{atan}\frac{b}{a+d}\right) U\left(\phi_2 - \operatorname{atan}\frac{b}{a-d}\right) U\left(\psi - \operatorname{atan}\frac{b}{h}\right), \quad (4)$$

where U is the Heaviside step function. An observation measurement is then simulated at each evenly spaced voxel V_i and pixel P_i to form a forward model matrix $\mathbf{A}$. The relation between the hidden scene and the observation image follows the matrix-vector expression:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{b} + \boldsymbol{\epsilon}, \quad (5)$$

where $\mathbf{x}$ is a vector representing the luminosity of each voxel, $\mathbf{y}$ is the observation image, $\mathbf{Bb}$ represents the contribution from visible-side illumination not affected by occlusions, and ϵ is noise. Because visible-side illumination is slow-varying, applying the finite difference operator $\mathbf{D}$ approximately cancels its effect. Equivalently:

$$\mathbf{D}\mathbf{y} = \mathbf{D}\mathbf{A}\mathbf{x} + \mathbf{D}\mathbf{B}\mathbf{b} + \boldsymbol{\epsilon}, \quad (6)$$

To solve for $\mathbf{x}$, we formulate an optimization problem,

$$(\hat{\mathbf{x}}, \hat{\mathbf{b}}) = \underset{(\mathbf{x},\mathbf{b})}{\arg\min} \|\mathbf{D}(\mathbf{y} - (\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{b}))\|_2^2 + \lambda \|\mathbf{x}\|_1, \quad (7)$$

where λ is a tuning parameter that controls the intensity of the l_1 regularization.

III. Experiments and Discussion

A. Cramer-Rao Bound Analysis

We compute and compare the Cramer-Rao Bound (CRB) for estimating a single hidden target given a measurement that exploits one, two, and three edges. The CRB gives an upper bound on the information about a parameter in the measurement. Figure 4 shows the uncertainty regions for the three cases. The single-edge edge occluder forms a large elliptical uncertainty region. The two-edge occluder, exploited in [29], shows the ellipse almost flattening into a line. Previously, [29] accounted for this direction of low information, i.e., the length of the line, by including a clustering step that effectively

reduces the number of ranges to be computed. In the three-edge case, however, the CRB is a small localized region, indicating the information present is strong and directionally balanced. The trend is that increasing the number of edges encodes more information in the measurement, with the three-edge case having the most informative measurement.

B. Reconstructions

We evaluate our approach on several real scenes composed of combinations of objects, such as a chair, a shelf, white-boards, and a green board. All experimental reconstructions have a discretization size of 30 by 30 by 15, placing voxels along 30 ϕ_1 positions, 30 ϕ_2 positions, and 15 ψ positions. The angles ψ_1 and ϕ_2 span 180 degrees, while ϕ values span $\text{atan}(\frac{\text{ImageLength}}{h})$ to $\pi/2$. After forming the discretized grid, facets with asymptotically large sizes or those that fall outside the $2 \times 2 \times 2$ meter room were excluded. The scenes shown are captured from a lab setup, where objects are placed inside a room with black walls. The observation photograph is of diffuse white boards attached to the ceiling. Light sources placed near the ceiling within the scene illuminate the hidden area but do not contribute to penumbra because the fall in a region $\psi < \text{atan}(\frac{\text{size}(y)}{h})$, such that they do not contribute directly to the ceiling observation. The lab setup also has white walls on the visible side that contribute significantly to the ambient background illumination.

Reconstructions shown in Figure 5 are obtained using triangular coordinates and the difference operator for three example scenes. In all three cases, the reconstructions have accurate color and positional information with minimal clutter.

Furthermore, to demonstrate the importance of the triangular coordinate system, Figure 6 shows reconstructions obtained when using the proposed triangular coordinates and conventional Cartesian coordinates representation using the same measurements. The Cartesian reconstructions fail, showing many spurious elements and poor color accuracy. By comparison, the triangular reconstructions accurately reconstruct the objects in the scene with little clutter. We also compared the impact of using the difference operator to suppress background in the measurements. The background impression procedure had little effect on improving the reconstructions.

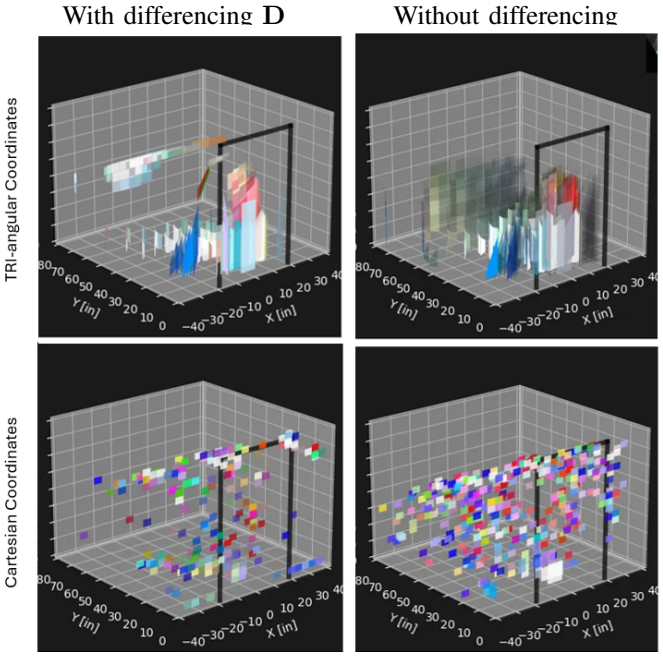

Fig. 6: Four reconstructions of the same scene using four different reconstruction algorithms: Triangular vs Cartesian coordinates, and with differencing operator $\mathbf{D}$ vs without differencing. The ground truth image for this scene is shown in the last row of column three in Figure 5.

IV. Conclusion

We presented a technique that resulted in an improvement to three-dimensional passive non-line-of-sight imaging. The information present in the penumbra drastically increased by expanding the complexity of the occluder to three edges. This expansion enabled a complete, single-step reconstruction algorithm that estimated the hidden scene's color, shape, and position information. Our reconstruction is a full 180° view reconstruction of the objects within the hidden room. The effectiveness of our approach was evaluated on real experimental scenes, where ambient light levels and model mismatch are prevalent, indicating the potential for use in real-world settings. The experiments shown here use black walls in the scene to reduce the amount of uninteresting light contributing to the observation. However, because this algorithm can localize the origin of light-emitting objects in three dimensions, it is feasible that non-black walls could also be recovered as a part of the reconstruction. Indeed, specular contributions from the back wall can be seen in a number of the reconstructions, despite attempts to reduce contributions from the room's inner walls.

We suspect that this work approaches a limit regarding passive NLOS with edge occluders, as adding a fourth edge will likely yield only marginal improvements. Similar to prior works, we rely on the Crámer-Rao bound analyses to tailor the reconstruction pipeline to the occluder shape, which here is the frame of an open doorway. While this simple occluding structure yields mathematically tractible expressions for the CRB, future work will likely entail formulating algorithms based on numerical CRB computations for more complex occluder structures with curved edges, volumetric structures, or unknown occluders. More generally, this work lays the foundation for determining optimal representations via CRB-based analyses in occluder-aided computational imaging with simple occluder geometries and may find utility in imaging inverse problems.

V. Awknowledgements

This material is based upon work supported in part by the US National Science Foundation CAREER award under Grant No. 2443049.

REFERENCES

[1] Ahmed Kirmani, Tyler Hutchison, James Davis, and Ramesh Raskar, "Looking around the corner using transient imaging," in *Proc. IEEE 12th Int. Conf. Computer Vision*, 2009, pp. 159–166.

[2] Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Moungi G Bawendi, and Ramesh Raskar, "Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging," *Nat. Commun.*, vol. 3, no. 745, 2012.

[3] Ahmed Kirmani, Haris Jeelani, Vahid Montazerhodjat, and Vivek 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, 2012.

[4] Matthew O'Toole, David B. Lindell, and Gordon Wetzstein, "Confocal non-line-of-sight imaging based on the light-cone transform," *Nature*, vol. 555, pp. 338–341, Mar. 2018.

[5] Xiaochun Liu, Ibón Guillén, Marco La Manna, Ji Hyun Nam, Syed Azer Reza, Toan Huu Le, Adrian Jarabo, Diego Gutierrez, and Andreas Velten, "Non-line-of-sight imaging using phasor-field virtual wave optics," *Nature*, vol. 572, no. 7771, pp. 620–623, 2019.

[6] Joshua Rapp, Charles Saunders, Julián Tachella, John Murray-Bruce, Yoann Altmann, Jean-Yves Tourneret, Stephen McLaughlin, Robin MA Dawson, Franco NC Wong, and Vivek K Goyal, "Seeing around corners with edge-resolved transient imaging," *Nature communications*, vol. 11, no. 1, pp. 1–10, 2020.

[7] Sheila Seidel, Hoover Rueda-Chacón, Iris Cusini, Federica Villa, Franco Zappa, Christopher Yu, and Vivek K Goyal, "Non-line-of-sight snapshots and background mapping with an active corner camera," *Nature Communications*, vol. 14, no. 1, pp. 3677, 2023.

[8] Christos Thrampoulidis, Gal Shulkind, Feihu Xu, William T Freeman, Jeffrey H Shapiro, Antonio Torralba, Franco NC Wong, and Gregory W Wornell, "Exploiting occlusion in non-line-of-sight active imaging," *IEEE Transactions on Computational Imaging*, vol. 4, no. 3, pp. 419–431, 2018.

[9] J. Grebe-Ellis and T. Quick, "Soft shadow images," *European Journal of Physics*, vol. 44, no. 4, pp. 045301, may 2023.

[10] Katherine L. Bouman, Vickie Ye, Adam B. Yedidia, Fredo Durand, Gregory W. Wornell, Antonio Torralba, and William T. Freeman, "Turning corners into cameras: Principles and methods," in *Proc. 23rd IEEE Int. Conf. Computer Vision*, 2017, pp. 2270–2278.

[11] Miika Aittala, Prafull Sharma, Lukas Murmann, Adam Yedidia, Gregory W. Wornell, William T. Freeman, and Frédo Durand, "Computational mirrors: Blind inverse light transport by deep matrix factorization," in *Advances in Neural Information Processing (NeurIPS)*, 2019, pp. 14311–14321.

[12] S. W. Seidel, Y. Ma, J. Murray-Bruce, C. Saunders, W. T. Freeman, C. C. Yu, and V. K. Goyal, "Corner occluder computational periscopy: Estimating a hidden scene from a single photograph," in *Proc. IEEE Int. Conf. Computational Photography*, May 2019, pp. 1–9.

[13] Charles Saunders, John Murray-Bruce, and Vivek K Goyal, "Computational periscopy with an ordinary digital camera," *Nature*, vol. 565, no. 7740, pp. 472–475, 2019.

[14] Charles Saunders, Rishabh Bose, John Murray-Bruce, and Vivek K Goyal, "Multi-depth computational periscopy with an ordinary camera," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Proc.* IEEE, 2020, pp. 9299–9305.

[15] Adam B. Yedidia, Manel Baradad, Christos Thrampoulidis, William T. Freeman, and Gregory W. Wornell, "Using unknown occluders to recover hidden scenes," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition*, June 2019.

[16] Sheila W. Seidel, John Murray-Bruce, Yanting Ma, Christopher Yu, William T. Freeman, and Vivek 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.

[17] Di Lin, Connor Hashemi, and James R. Leger, "Passive non-line-of-sight imaging using plenoptic information," *J. Opt. Soc. Am. A*, vol. 37, no. 4, pp. 540–551, Apr 2020.

[18] Takahiro Sasaki and James R. Leger, "Non-line-of-sight object location estimation from scattered light using plenoptic data," *J. Opt. Soc. Am. A*, vol. 38, no. 2, pp. 211–228, Feb 2021.

[19] Ruixu Geng; Yang Hu; Zhi Lu; Cong Yu; Houqiang Li; Hengyu Zhang; Yan Chen, "Passive non-line-of-sight imaging using optimal transport," *IEEE Transactions on Image Processing*, vol. 31, pp. 110–124, 2022.

[20] T. Swedish, C. Henley, and R. Raskar, "Objects as cameras: Estimating high-frequency illumination from shadows," in *Proc. IEEE/CVF Int. Conf. Computer Vision*, October 2021, pp. 2593–2602.

[21] Robinson Czajkowski and John Murray-Bruce, "Turning door frames into cameras for 3d non-line-of-sight imaging," in *2022 IEEE Research and Applications of Photonics in Defense Conference (RAPID)*, 2022, pp. 1–2.

[22] C. Henley, T. Maeda, T. Swedish, and R. Raskar, "Imaging behind occluders using two-bounce light," in *Proc. Eur. Conf. Computer Vision*. Springer, 2020, pp. 573–588.

[23] John Murray-Bruce, Charles Saunders, and Vivek K. Goyal, "Occlusion-based computational periscopy with consumer cameras," in *SPIE Wavelets and Sparsity XVIII*. International Society for Optics and Photonics, 2019, vol. 11138, pp. 286–297, SPIE.

[24] Antonio Torralba and William T. Freeman, "Accidental pinhole and pinspeck cameras: Revealing the scene outside the picture," *Int. J. Computer Vision*, vol. 110, no. 2, pp. 92–112, 2014.

[25] Adam Lloyd Cohen, "Anti-pinhole imaging," *Optica Acta*, vol. 29, no. 1, pp. 63–67, 1982.

[26] Fadlullah Raji and John Murray–Bruce, "Occlusion-aided 3d non-line-of-sight imaging without time-resolved measurements," in *2024 IEEE Research and Applications of Photonics in Defense Conference (RAPID)*, 2024, pp. 01–02.

[27] Fadlullah Raji and John Murray-Bruce, "Towards 3d computational persicopy with an ordinary camera: a separable non-linear least squares formulation," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 7475–7479.

[28] William Krska, Sheila W. Seidel, Charles Saunders, Robinson Czajkowski, Christopher Yu, John Murray-Bruce, and Vivek Goyal, "Double your corners, double your fun: The doorway camera," in *Proc. IEEE Int. Conf. Computational Photography*, 2022, pp. 1–12.

[29] Robinson Czajkowski and John Murray-Bruce, "Two-edge-resolved three-dimensional non-line-of-sight imaging with an ordinary camera," *To Appear*, 2024.