Computer Graphics meets Estimation Theory: Parameter Estimation Lower Bounds for Plenoptic Imaging Systems

Abhinav V. Sambasivan
Department of ECE
Univ. of Minnesota - Twin Cities, MN
samba014@umn.edu

Richard G. Paxman
Sensor Programs
MAXAR - Ypsilanti, MI
rick.paxman@maxar.com

Jarvis D. Haupt
Department of ECE
Univ. of Minnesota - Twin Cities, MN
jdhaupt@umn.edu

***Abstract*—This work focuses on assessing the information-theoretic limits of parameter estimation in plenoptic imaging systems, which are capable of providing substantially more information about a given scene than conventional cameras. We present a framework to compute lower bounds for parameter estimation from noisy plenoptic observations, and our particular focus is on indirect imaging problems, where the observations do not contain line-of-sight (LOS) information about the parameter(s) of interest. Using computer graphics rendering software to synthesize the (often complicated) dependence among parameter(s) of interest and observations, we numerically evaluate the Hammersley-Chapman-Robbins bound to establish fundamental lower limits on the variance of any unbiased estimators of the unknown parameters. We demonstrate the utility of our proposed framework on a few canonical estimation tasks.**

***Index Terms*—Plenoptic Imaging, Cramer-Rao bound, Hammersley-Chapman-Robbins bound**

I. Introduction

Conventional imaging systems are modeled after human vision and provide information about a scene via a two-dimensional image, where each pixel described by a tristimulus value (e.g., RGB). However, the plenoptic function [1] captures the intensity of light at every location in space, over all possible angles and wavelengths, at all time instances, providing much more information about a given scene than a conventional camera. The resulting plenoptic function takes the form

$$\mathbf{L} = \mathbf{L}(\mathbf{r}, \boldsymbol{\varphi}, \lambda, t),$$

where the $\mathbf{r} \in \mathbb{R}^3$ refers to the location at which the plenoptic observation is made, $\boldsymbol{\varphi} \in [0, 2\pi) \times [0, \pi)$ denotes the angle pair (direction of arrival of light) in polar coordinates, λ is the wavelength, and t is the time index. The domain over which the plenoptic function is measured typically changes with the imaging modality used for a particular application. Plenoptic imaging has been used for a wide range of applications including stereoscopy [2], [3], microscopy [4], [5], and Non-Line-Of-Sight (NLOS) imaging [6]–[8].

Our focus here is on assessing the fundamental limitations of plenoptic imaging problems, with an eye towards indirect (NLOS) imaging, from an estimation-theoretic standpoint. Combining analytical tools from classical estimation theory [9]–[12] with computer graphics rendering engines [13], we develop a framework for computing lower bounds for estimating "scene parameters" from noisy plenoptic observations. In doing so, we are able to provide a benchmark for optimality against which we could compare various estimation strategies, and also obtain useful insights on where information is localized in a given set of observations.

The remainder of the paper is organized as follows. We begin by providing an overview of the *rendering equation* [14], which explains the forward model for our settings in Section II. Using this, the problem setup is formalized in Section III. The proposed framework and the information-theoretic tools used appear in Section IV. We demonstrate the utility of our framework to compute lower bounds for object location and size estimation problems in Section V, and finish with some concluding remarks about future work in Section VI.

II. Rendering Equation and the Forward Model

For a given scene with some unknown (deterministic) parameter of interest $\boldsymbol{\theta}^* \in \Theta \subseteq \mathbb{R}^d$, we denote the set of plenoptic samples (obtained by sampling the full plenoptic field) associated with the scene by, $\mathbf{L}^{\Omega}_{\boldsymbol{\theta}^*} \triangleq \{\mathbf{L}_{\boldsymbol{\theta}^*}(\boldsymbol{\omega})\}_{\boldsymbol{\omega} \in \Omega}$, where we introduce the shorthand $\boldsymbol{\omega} := [\mathbf{r}, \boldsymbol{\varphi}, \lambda, t]$, to denote the arguments of the plenoptic function, and Ω is the observation space, which is a subset of the domain over which the plenoptic function is defined.

Information-theoretic treatment of this estimation problem requires understanding the *forward model* $\boldsymbol{\theta}^* \mapsto \mathbf{L}^{\Omega}_{\boldsymbol{\theta}^*}$ that captures the functional dependence of the plenoptic observations on the scene parameters. This can be mathematically explained by the rendering equation [14], which models the behavior of light rays as they originate from a light source, bounce off the objects in the scene and ultimately reach the detector (an imaging device).

For fixed (λ, t), light *incident* on a surface point $\mathbf{r}$ of a scene, along direction $-\boldsymbol{\varphi}_i$, denoted by $\mathbf{L}^{\text{in}}_{\boldsymbol{\theta}^*}(\mathbf{r}, -\boldsymbol{\varphi}_i)$, is typically scattered along all directions, as shown in Fig. 1a. The proportion of this light reflected along $\boldsymbol{\varphi}_o$ is determined

The authors graciously acknowledge support from the DARPA REVEAL program, Contract No. HR0011-16-C-0024.

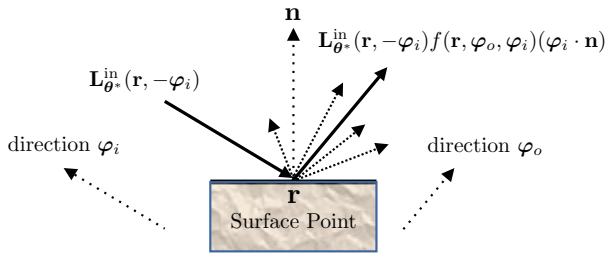

(a) The proportion of incident light coming in from direction $\boldsymbol{\varphi}_i$ that gets reflected along direction $\boldsymbol{\varphi}_o$ is determined by the BRDF of the surface.

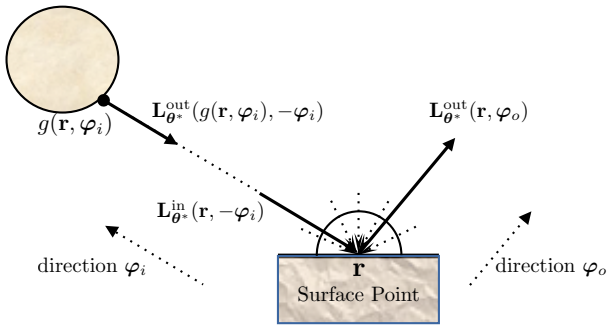

(b) Light *incident* on a surface point $\mathbf{r}$, can be seen as light *leaving* from another point in the scene $g(\mathbf{r}, \boldsymbol{\varphi}_i)$, $\Rightarrow$ $\mathbf{L}^{\text{in}}_{\boldsymbol{\theta}^*}(\mathbf{r}, -\boldsymbol{\varphi}_i) = \mathbf{L}^{\text{out}}_{\boldsymbol{\theta}^*}(g(\mathbf{r}, \boldsymbol{\varphi}_i), -\boldsymbol{\varphi}_i)$.

Fig. 1: The rendering equation, explained graphically.

by a surface-dependent quantity known as the bi-directional reflectance distribution function (BRDF) $f(\mathbf{r}, \cdot, \cdot)$.

The overall radiance *leaving* a surface point $\mathbf{r}$, along direction $\boldsymbol{\varphi}_o$, denoted by $\mathbf{L}^{\text{out}}_{\boldsymbol{\theta}^*}(\mathbf{r}, \boldsymbol{\varphi}_o)$, can hence be obtained by summing up the contributions of $\mathbf{L}^{\text{in}}_{\boldsymbol{\theta}^*}(\mathbf{r}, -\varphi_i)$ from all possible incident directions $-\boldsymbol{\varphi}_i$, and also any light emitted by the surface (if it is an emitter). Mathematically,

$$\mathbf{L}^{\text{out}}_{\boldsymbol{\theta}^*}(\mathbf{r}, \boldsymbol{\varphi}_o) = \mathbf{L}^{\text{e}}_{\boldsymbol{\theta}^*}(\mathbf{r}, \boldsymbol{\varphi}_o) + \int_{S^2_+(\mathbf{r})} \mathbf{L}^{\text{in}}_{\boldsymbol{\theta}^*}(\mathbf{r}, -\boldsymbol{\varphi}_i) f(\mathbf{r}, \boldsymbol{\varphi}_o, \boldsymbol{\varphi}_i)(\boldsymbol{\varphi}_i \cdot \mathbf{n}) d\boldsymbol{\varphi}_i,$$

where $\mathbf{L}^{\text{e}}_{\boldsymbol{\theta}^*}(\mathbf{r}, \boldsymbol{\varphi}_o)$ is the emitted radiance at $\mathbf{r}$ along the direction $\boldsymbol{\varphi}_o$, $\mathbf{n}$ is the surface normal at $\mathbf{r}$, and $S^2_+(\mathbf{r})$ is the unit hemisphere at $\mathbf{r}$ containing all outgoing directions.

From Fig. 1b, we can see that $\mathbf{L}^{\text{in}}_{\boldsymbol{\theta}^*}(\mathbf{r}, -\varphi_i) = \mathbf{L}^{\text{out}}_{\boldsymbol{\theta}^*}(g(\mathbf{r}, \boldsymbol{\varphi}_i), -\boldsymbol{\varphi}_i)$, where $g(\mathbf{r}, \boldsymbol{\varphi}_i)$ is a scene geometry dependent operator that essentially finds the first surface point reached when traveling outward from $\mathbf{r}$, along direction $\boldsymbol{\varphi}_i$. Using this we can obtain the rendering equation as

$$\mathbf{L}^{\text{out}}_{\boldsymbol{\theta}^*}(\mathbf{r}, \boldsymbol{\varphi}_o) = \mathbf{L}^{\text{e}}_{\boldsymbol{\theta}^*}(\mathbf{r}, \boldsymbol{\varphi}_o) + \int_{S^2_+(\mathbf{r})} \mathbf{L}^{\text{out}}_{\boldsymbol{\theta}^*}(g(\mathbf{r}, \boldsymbol{\varphi}_i), -\boldsymbol{\varphi}_i) f(\mathbf{r}, \boldsymbol{\varphi}_0, \boldsymbol{\varphi}_i)(\boldsymbol{\varphi}_i \cdot \mathbf{n}) d\boldsymbol{\varphi}_i, \quad (1)$$

where $\mathbf{L}^{\text{out}}_{\boldsymbol{\theta}^*}$ is the plenoptic function of interest in the forward model alluded to above.

Equation (1) is a Fredholm integral equation of the second kind, and is difficult to solve in closed form for all but simplest of settings. To overcome this hurdle, ray-tracing/rendering engines were developed to approximately solve the integral equation in (1) using Monte Carlo methods. Ray-tracing engines are widely used in the computer graphics community to generate photo-realistic images by tracing the path of light rays as they interact with objects in a scene. Here, we rely upon this (mature) technology for our purposes; for a given parameter value $\boldsymbol{\theta}^*$, we use ray-tracing packages to approximately solve (1) and synthesize the plenoptic observations $\mathbf{L}^{\Omega}_{\boldsymbol{\theta}^*}$.

III. Problem Statement

We model the noisy observations $\mathbf{Y}_{\boldsymbol{\omega}}$, as independent draws from a known class of probability distributions (whose parameters depends on $\boldsymbol{\theta}^*$ through the plenoptic function). Letting $\mathbf{Y}_\Omega \triangleq \{\mathbf{Y}_{\boldsymbol{\omega}}\}_{\boldsymbol{\omega}\in\Omega}$ be shorthand for the entire collection of noisy plenoptic observations, the likelihood of the observations can be written as

$$p(\mathbf{Y}_\Omega; \boldsymbol{\theta}^*) = \prod_{\boldsymbol{\omega}\in\Omega} p(\mathbf{Y}_{\boldsymbol{\omega}}; \boldsymbol{\theta}^*) \triangleq \mathbb{P}_{\boldsymbol{\theta}^*}. \quad (2)$$

We are interested in determining the fundamental limits of imaging a given scene with unknown parameter $\boldsymbol{\theta}^* \in \Theta$, from noisy plenoptic measurements $\mathbf{Y}_\Omega \in \mathcal{Y}$. Specifically, we seek to evaluate the performance of estimators, $\widehat{\boldsymbol{\theta}}(\mathbf{Y}_\Omega) \triangleq \widehat{\boldsymbol{\theta}} : \mathcal{Y} \to \Theta$, which have a finite measure w.r.t $\mathbb{P}_{\boldsymbol{\theta}^*}$, as a function of the true underlying parameter of interest $\boldsymbol{\theta}^*$. The accuracy (or lack thereof) of an estimator $\widehat{\boldsymbol{\theta}}$ in estimating the true scene parameter $\boldsymbol{\theta}^*$ can be measured in terms of the Mean Squared Error (MSE),

$$MSE_{\,\widehat{\boldsymbol{\theta}}}(\boldsymbol{\theta}^*) = \mathbb{E}\left[\left|\left|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\right|\right|_2^2\right], \quad (3)$$

where the expectation is taken with respect to the randomness in the noisy observations $\mathbf{Y}_\Omega$.

We restrict our attention to *unbiased estimators*, for which MSE reduces to variance. Hence we seek lower bounds on the covariance matrix $\text{Cov}_{\boldsymbol{\theta}^*}(\widehat{\boldsymbol{\theta}}) \triangleq \text{Cov}(\boldsymbol{\theta}^*)$ in a semi-definite ordering sense, or on $\text{Var}_{\boldsymbol{\theta}^*}(\widehat{\boldsymbol{\theta}}) \triangleq \text{Var}(\boldsymbol{\theta}^*) = \text{tr}\left\{\text{Cov}_{\boldsymbol{\theta}^*}(\widehat{\boldsymbol{\theta}})\right\}$.

IV. Renderer-Enabled Computation of Lower Bounds

Our approach employs the Hammersley-Chapman-Robbins lower bound (HCR-LB) [11], [12], which provides lower bounds on the variance of unbiased estimators. HCR-LB and its variants have been used in constrained parameter estimation problems for sensor array signal processing, like bearing estimation [15]–[17], frequency estimation [18], [19] and many more. In this section, we will state the HCR-LB and discuss some of its salient aspects which make it ideally suited to this problem setting.

Lemma 1 (HCR Lower Bound)**.** *Let $\boldsymbol{\theta}^* \in \Theta \subseteq \mathbb{R}^d$ be any deterministic but unknown parameter, and let $\mathbf{Y}_\Omega$ denote a set of noisy observations of the unknown parameter $\boldsymbol{\theta}^*$. Then the variance of any unbiased estimator of $\boldsymbol{\theta}^*_i$ obeys,*

$$\mathrm{Var}(\boldsymbol{\theta}_i^*) \geq \sup_{\substack{\boldsymbol{\Delta}\neq\mathbf{0} \\ \boldsymbol{\theta}^*+\boldsymbol{\Delta}\in\Theta}} \frac{\boldsymbol{\Delta}_i^2}{\mathbb{E}_{\boldsymbol{\theta}^*}\left[\frac{p(\mathbf{Y}_\Omega;\boldsymbol{\theta}^*+\boldsymbol{\Delta})}{p(\mathbf{Y}_\Omega;\boldsymbol{\theta}^*)}-1\right]^2}, \forall\ i=1,\ldots,d, \tag{4}$$

where $p(\mathbf{Y}_\Omega;\boldsymbol{\theta}^*)$, $p(\mathbf{Y}_\Omega;\boldsymbol{\theta}^*+\boldsymbol{\Delta})$ denote the pdfs (or pmfs) of the observations parametrized by $\boldsymbol{\theta}^*$, $\boldsymbol{\theta}^*+\boldsymbol{\Delta}$ respectively.

Here the denominator in the RHS of (4) is the so-called χ^2-divergence of $p(\mathbf{Y};\boldsymbol{\theta}^*+\boldsymbol{\Delta})$ from $p(\mathbf{Y};\boldsymbol{\theta}^*)$ which essentially measures changes in the probability distribution functions when the true parameter $\boldsymbol{\theta}^*$ is perturbed by $\boldsymbol{\Delta}$. When a small change in the unknown parameter results in distinctly different sets of observations, resulting in a large χ^2-divergence between the likelihoods, the HCR-LB is small, suggesting that the estimation problem could be easy and vice versa.

It is worth commenting on the relationship of the HCR-LB framework with the well known Cramer-Rao lower bound (CR-LB) [9], [10], [20]. If we let $\boldsymbol{\Delta}_i \to \mathbf{0}$, the expression inside the supremum of (4) converges to the CR-LB (if the corresponding limit exists). Thus when both bounds exist, HCR-LB is at least as tight as the CR-LB. Unlike the CR-LB, HCR-LB makes no "regularity" assumptions on the noise likelihood function and hence is applicable for a wider range of problems. In particular, CR-LB requires computing derivatives of the log-likelihood function (w.r.t $\boldsymbol{\theta}$) and is not well defined in scenarios where the log-likelihood is not differentiable (e.g., when the parameter space is a countable set). Even for simple scenes, due to the presence of occluding barriers and edges, sharp "transition regions" may occur in the true plenoptic intensities $\mathbf{L}^{\Omega}_{\boldsymbol{\theta}^*}$ as the underlying scene parameter $\boldsymbol{\theta}^*$ varies smoothly resulting in the log-likelihood to be non-differentiable.

On the other hand, HCR-LB doesn't require explicitly computing derivatives of the log-likelihood (which is hard to obtain in our problem setting). For any given noise model, this can be done by rendering or synthesizing $\mathbf{L}^{\Omega}_{\boldsymbol{\theta}^*+\boldsymbol{\Delta}}$ for a suitably large collection of possible values of $\boldsymbol{\theta}^*$ and $\boldsymbol{\Delta}$ using a ray-tracing engine and then evaluating the functional form of the χ^2-divergence.

A. HCR Lower Bound under Poisson Noise

In this section, we instantiate the HCR bound for a specific noise model. We consider noisy observations $\mathbf{Y}_{\boldsymbol{\omega}}$, which are drawn independently from a Poisson distribution with rates given by the true plenoptic intensities $\mathbf{L}_{\boldsymbol{\theta}^*}(\boldsymbol{\omega})$, i.e. $\mathbf{Y}_{\boldsymbol{\omega}} \overset{ind}{\sim} \mathrm{Poisson}(\mathbf{L}_{\boldsymbol{\theta}^*}(\boldsymbol{\omega}))$, $\forall \boldsymbol{\omega}\in\Omega$. The Poisson distribution is commonly used to characterize noise which is discrete or quantized in nature, e.g. when the imaging device counts the number of photons incident on the detector over a certain window of time. If we specialize the HCR bound given in (4) for this setting, we get

$$\mathrm{Var}(\boldsymbol{\theta}_i^*) \geq \sup_{\substack{\boldsymbol{\Delta}\neq\mathbf{0} \\ \boldsymbol{\theta}^*+\boldsymbol{\Delta}\in\Theta}} \frac{\boldsymbol{\Delta}_i^2}{\exp\left(\sum_{\boldsymbol{\omega}\in\Omega}\frac{\left(\mathbf{L}_{\boldsymbol{\theta}^*+\boldsymbol{\Delta}}(\boldsymbol{\omega})-\mathbf{L}_{\boldsymbol{\theta}^*}(\boldsymbol{\omega})\right)^2}{\mathbf{L}_{\boldsymbol{\theta}^*}(\boldsymbol{\omega})}\right)-1}, \quad \forall i=1,2,\ldots,d. \tag{5}$$

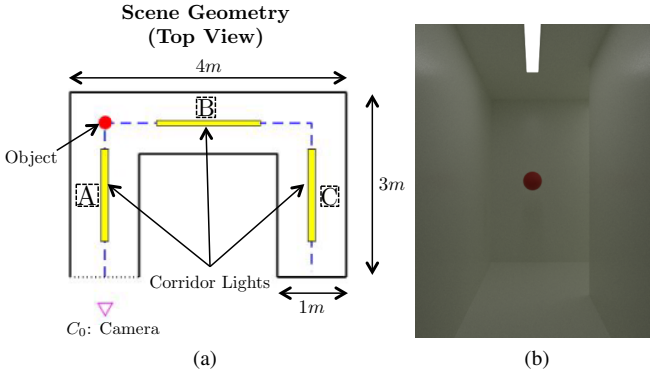

(a) (b)

Fig. 2: An illustrative example: (a) Layout of a Π-shaped hallway with dimensions marked. Corridors A, B, and C are 2.5m, 3m, and 2.5m long respectively, and 2m tall. The hallway is illuminated with white ceiling lights with a radiance 3 $\mathrm{Wsr}^{-1}\mathrm{m}^{-2}$. A red spherical ball is placed at the intersection of corridor A and B. (c) If we let the the ball location and/its radius to be the unknown parameter $\boldsymbol{\theta}^*$, a nominal RGB image of the scene $\mathbf{L}^{\Omega}_{\boldsymbol{\theta}^*}$, rendering using Mitsuba [13] is shown in (b).

where all the true plenoptic intensities in (5) can be obtained using the rendering software.

V. Experimental Evaluation: Estimating Object Location and Size

In this section, we demonstrate the efficacy of the proposed framework using two illustrative problems: object localization and size estimation. We begin by describing the scene layout and geometry and then discuss some results and their implications.

A. Experimental Setup

Let us consider a scene consisting of a Π-shaped hallway as shown in Fig. 2. A red spherical ball which is the object of interest is placed somewhere in the hallway. It is further assumed that the location of the object is confined to the dotted lines shown in Fig. 2a which is centered between the corridor walls, floor, and ceiling. Walls are painted with a diffuse eggshell paint, and their BRDF is modeled using the `roughplastic` plugin in Mitsuba [13].

We further assume that noisy plenoptic measurements are made by an imaging device located at C_0 which collects multispectral images with 30 uniformly spaced spectral channels in the $368-830$nm range (as opposed to the RGB observations illustrated in Fig. 2b). With this setup, we study the fundamental limits of two different *scalar* estimation problems:

- Estimating location of the ball (with a fixed radius of 10cm) along the Π-shaped hallway.
- Estimating radius (size) of the ball located at the center of Corridor B.

For these estimation problems, we assume that the scene geometry and the BRDFs of the surfaces are known, i.e., the only unknown in the scene is either the location of the ball or its radius. We discretize the 1-D Π-shaped manifold along which the ball could be located to 790 equispaced points (each

1cm apart). Ball locations are numbered clockwise from 1 (bottom-left of the Π-shaped manifold) to 790 (bottow-right of the manifold). For size estimation problem, we consider 101 possible values for ball radii in the range 5cm to 15cm (with 0.1cm increments). We use the Mitusba renderer [13] to synthesize physically accurate plenoptic samples $\mathbf{L}^{\Omega}_{\boldsymbol{\theta}^*}$, for different values of the unknown parameter $\boldsymbol{\theta}^*$ (and $\boldsymbol{\Delta}$), and numerically evaluate the HCR bound (5).

B. Results and Discussion

Using the rendered plenoptic data, we numerically evaluate the HCR-LB (5) for the problems described above. We then use finite-difference methods on the rendered data to (approximately) compute (pixelwise) Fisher Information, which we call Finite Difference-Fisher Information (FD-FI).

Figures 3a to 3c depict the HCR-LB for estimating the ball location. The HCR-LB is (unsurprisingly) small ($\ll 1$) when the ball is in corridor A (in LOS). The lower bound suddenly drops when the ball translates horizontally (w.r.t the camera) in corridor B (see Fig. 3a). This is expected, as detecting an object is easier when it moves horizontally as compared to when it moves away from the camera. HCR-LB increases by a few orders of magnitude when ball goes away from LOS – location 320 $\mapsto$ location 321 (see Fig. 3b), and continues to increase sharply for ball locations further down corridor B (see Fig. 3c). In corridor C, the actual value of the bound gets higher $\sim 10^4\text{cm}^2$ which indicates that accurately locating the ball around two corners is (again, unsurprisingly) a much more difficult estimation problem.

Some recent works [7], [8] have proposed that sharp edges and occlusions in a scene can help recover NLOS imagery. These sharp occlusions act like a pin-hole camera, projecting information about NLOS regions of a scene on to visible regions, e.g., walls and floors. From looking at the FD-FI images, we believe that some of the dips in the HCR-LB in corridor C (in Fig. 3c) result from such effects, and in a way corroborates those findings.

Pixelwise FD-FI images in Fig. 4a provide valuable insights on where and how information about a parameter of interest is distributed amongst the observed samples. As one might expect, when in LOS, the edges of the ball provide most information about its location. When the ball moves away from LOS, information about its location is conveyed by shadows cast by the ceiling light in corridor B on the floor, and other *indirect* photons from the back wall and the floor. These subtle details are however not visible in the raw plenoptic data (or the nominal RGB images in bottom-row of Fig. 4).

The HCR-LB for estimating the radius of a ball placed at the center of corridor B (location: 400) is shown in Fig. 3d. We observe that the lower bound decreases as the radius increases, suggesting that estimating radius of a larger ball could be relatively easier. Pixelwise FD-FI images for radius estimation problem in Fig. 4b, show that regions of information are mainly on the back wall which arise from multi-bounce photons, and the number of informative pixels increases for larger radii. We can see that the regions of high information differ significantly between Fig. 4a and Fig. 4b. We conclude that these informative regions/observations depend on both the parameter of interest at large (see difference between Figs. 4a and 4b), and also on the *particular value* of the parameter (see differences amongst images in the top-row of each of those figures).

VI. Conclusion and Future Work

We presented a framework to compute information-theoretic lower bounds for estimating scene parameters from noisy plenoptic data. Our approach employed HCR-LB over commonly used CR-LB, as the former is: (a) more amenable to our settings where computing partial derivatives (w.r.t parameters of interest) might be infeasible, (b) applies to a wider range of problems, and (c) yields a bound that is at least as good as CR-LB. Using computer graphics rendering packages, we overcome the difficulty of having to solve the forward model in closed form, by numerically evaluating the HCR bound.

Our renderer-enabled lower-bounding framework has been used to compute lower bounds for a few illustrative NLOS imaging problems. Additionally, we are able compute pixelwise Fisher information (using finite differences). This FD-FI data provides useful insights, especially for NLOS imaging problems as they tell us which of the indirect photons/observations convey more information about the parameter(s) of interest. We believe that these insights and tools can be used to develop novel adaptive sensing strategies for scene parameter estimation. Although we explore only classical multi-spectral imaging systems in this work, our estimation-theoretic framework readily generalizes to accommodate additional dimensions of the plenoptic function, for e.g., lenslet-array camera systems.

The potential benefits of HCR-LB framework come at a cost of increased computational requirements. Computing HCR-LB involves finding the supremum of the expression in (4) for general noise models, or (5) for the Poisson case. In this work, we compute the supremum by exhaustive search since the parameter space is small. However, if one is interested in problems with multiple parameters ($d \gg 1$), the computational time for exhaustive search would increase exponentially making it prohibitive to apply in high-dimensional settings. For such settings, we could potentially explore either derivative free (zeroth-order) optimization methods, or recently developed differentiable renderers like [21], [22] to compute derivatives and evaluate the HCR-LB (5). We defer these investigations to future work.

References

[1] E. H. Adelson, J. R. Bergen, et al., "The plenoptic function and the elements of early vision," 1991.

[2] E. H Adelson and J. Y. A. Wang, "Single lens stereo with a plenoptic camera," *IEEE Trans. on Pattern Analysis & Machine Intelligence*, , no. 2, pp. 99–106, 1992.

[3] M. Levoy and P. Hanrahan, "Light field rendering," in *Proc. of the 23rd Annual Conf. on Comp. graphics and interactive techniques*. ACM, 1996, pp. 31–42.

[4] R. Prevedel et al., "Simultaneous whole-animal 3d imaging of neuronal activity using light-field microscopy," *Nature methods*, vol. 11, no. 7, pp. 727, 2014.

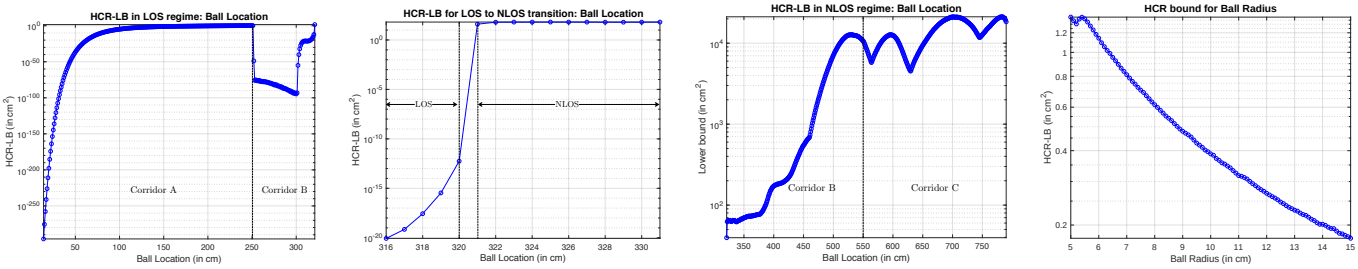

(a) HCR-LB drops significantly when ball starts moving in Corridor B.

(b) Sharp increase in the HCR-LB when the ball moves away from LOS.

(c) The HCR-LB in the NLOS region is much higher, implying the potential hardness of the estimation problem.

(d) The HCR-LB decreases with increasing size (radius) of the ball.

Fig. 3: HCR-LB for ball location estimation in (a) LOS region, (b) Transition from LOS to NLOS, and (c) NLOS region; (d) Ball radius estimation, With the help of these curves, one can quantify *how difficult* the problem of estimating NLOS parameters can be.

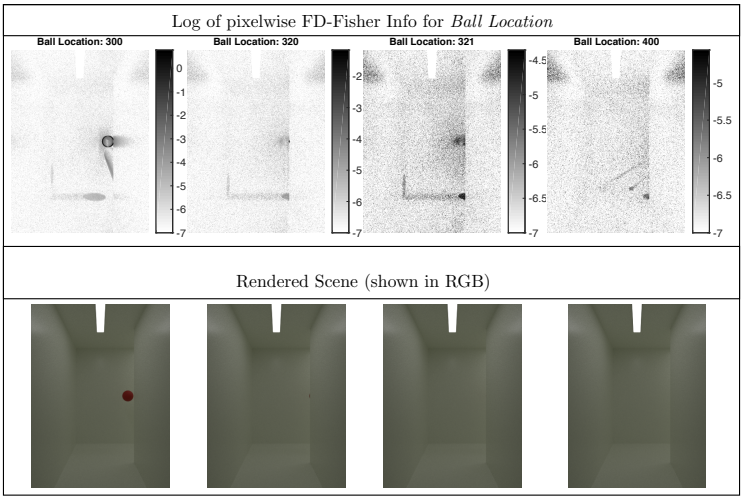

(a) Pixelwise FD-FI for 4 different ball locations: (top-row from left to right) completely in LOS, just *inside* LOS, just *outside* LOS, center of corridor B. Notice that different regions in the scene are more informative than others for different ball locations.

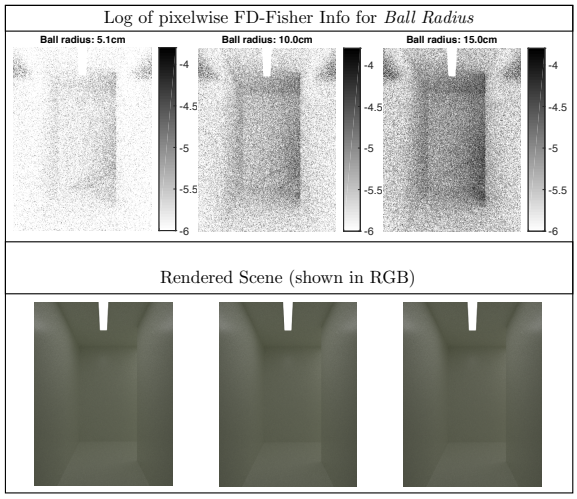

(b) Pixelwise FD-FI images show where information about *ball radius* is localized. Notice that regions of information differ from the FD-FI images for ball location in Fig. 4a.

Fig. 4: Pixelwise FD-Fisher Information (FD-FI), obtained by aggregating contributions from all 30 spectral channels. Darker regions ⇒ more informative, Pixelwise FD-FI shows where and how information about the parameter of interest is localized in our observations. These images highlight subtle details about the scene parameters which are not visible from the nominal RGB images (bottow-row) which are obtained from the rendering software.

[5] M. Levoy et al., "Light field microscopy," *ACM Trans. on Graphics (TOG)*, vol. 25, no. 3, pp. 924–934, 2006.

[6] T. Sasaki and J. R. Leger, "Light-field reconstruction from scattered light using plenoptic data," in *Unconventional and Indirect Imaging, Image Reconstruction, and Wavefront Sensing 2018*. International Society for Optics and Photonics, 2018, vol. 10772, p. 1077203.

[7] K. L. Bouman et al., "Turning corners into cameras: Principles and methods," in *Proc. of the IEEE Intl. Conf. on Comp. Vision*, 2017, pp. 2270–2278.

[8] M. Baradad et al., "Inferring Light Fields from Shadows," in *The IEEE Conf. on Comp. Vision and Pattern Recog.*, June 2018.

[9] H. Cramer, *Mathematical Methods of Statistics*, Princeton, NJ: Princeton Univ. Press, 1946.

[10] C. R. Rao, "Minimum variance and the estimation of several parameters," in *Mathematical Proc. of the Cambridge Philosophical Society*. Cambridge University Press, 1947, vol. 43, pp. 280–283.

[11] J. M. Hammersley, "On estimating restricted parameters," *Journal of the Royal Stats. Society. Series B (Methodological)*, vol. 12, no. 2, pp. 192–240, 1950.

[12] D. G. Chapman and H. Robbins, "Minimum variance estimation without regularity assumptions," *The Ann. of Mathematical Stats.*, pp. 581–586, 1951.

[13] W. Jakob, "Mitsuba renderer," 2010, http://www.mitsuba-renderer.org.

[14] J. T. Kajiya, "The rendering equation," in *ACM Siggraph Computer Graphics*. ACM, 1986, vol. 20, pp. 143–150.

[15] S. K. Chow and P. Schultheiss, "Delay estimation using narrow-band processes," *IEEE Trans. on Acoustics, Speech, and Sig. Process.*, vol. 29, no. 3, pp. 478–484, 1981.

[16] J. Tabrikian and J. L. Krolik, "Barankin bounds for source localization in an uncertain ocean environment," *IEEE Trans. on Sig. Process.*, vol. 47, no. 11, pp. 2917–2927, 1999.

[17] R. McAulay and E. Hofstetter, "Barankin bounds on parameter estimation," *IEEE Trans. on Info. Theory*, vol. 17, no. 6, pp. 669–676, 1971.

[18] L. Knockaert, "The Barankin bound and threshold behavior in frequency estimation," *IEEE Trans. on Sig. Process.*, vol. 45, no. 9, pp. 2398–2401, 1997.

[19] J. D. Gorman and A. O. Hero, "Lower bounds for parametric estimation with constraints," *IEEE Trans. on Info. Theory*, vol. 36, no. 6, pp. 1285–1301, Nov 1990.

[20] C. R. Rao, "Information and the accuracy attainable in the estimation of statistical parameters," in *Bulletin of the Calcutta Mathematical Society*, 1945, vol. 37, pp. 81–91.

[21] T. M. Li, M. Aittala, F. Durand, and J. Lehtinen, "Differentiable monte carlo ray tracing through edge sampling," in *SIGGRAPH Asia 2018 Technical Papers*. ACM, 2018, p. 222.

[22] M. M. Loper and M. J. Black, "Opendr: An approximate differentiable renderer," in *European Conf. on Comp. Vision*. Springer, 2014, pp. 154–169.