Homogeneous Codes for Energy-Efficient Illumination and Imaging

Matthew O'Toole
University of Toronto

Supreeth Achar
Carnegie Mellon University

Srinivasa G. Narasimhan
Carnegie Mellon University

Kiriakos N. Kutulakos
University of Toronto

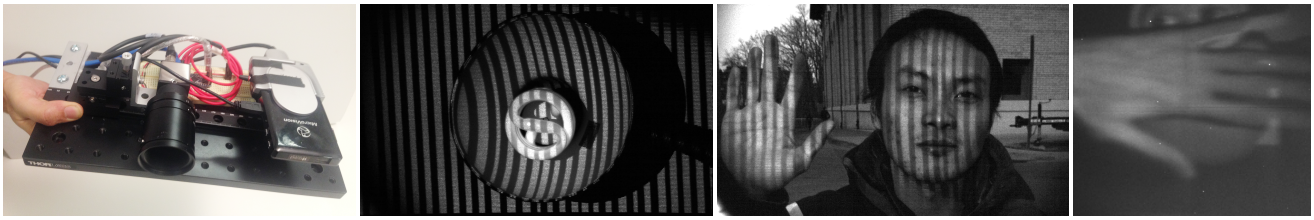

(a) (b) (c) (d)

Figure 1: *Imagine trying to acquire live video (30fps) of structured-light patterns as they are being projected onto a compact fluorescent bulb that has already been turned on (rated 1600 Lumens), or onto a face in bright sunlight (80 kLux)—with a 5-Lumen projector. One of our two prototypes, shown in (a), achieves this with an off-the-shelf laser projector and a CMOS camera with an ordinary lens. We then used it to capture the video frames in (b) and (c). We also show how to use our prototype to generate a live video feed from the projector's—rather than the camera's—point of view, shown in (d).*

Abstract

Programmable coding of light between a source and a sensor has led to several important results in computational illumination, imaging and display. Little is known, however, about how to utilize energy most effectively, especially for applications in live imaging. In this paper, we derive a novel framework to maximize energy efficiency by "homogeneous matrix factorization" that respects the physical constraints of many coding mechanisms (DMDs/LCDs, lasers, *etc.*). We demonstrate energy-efficient imaging using two prototypes based on DMD and laser illumination. For our DMD-based prototype, we use fast local optimization to derive codes that yield brighter images with fewer artifacts in many transport probing tasks. Our second prototype uses a novel combination of a low-power laser projector and a rolling shutter camera. We use this prototype to demonstrate never-seen-before capabilities such as (1) capturing live structured-light video of very bright scenes—even a light bulb that has been turned on; (2) capturing epipolar-only and indirect-only live video with optimal energy efficiency; (3) using a low-power projector to reconstruct 3D objects in challenging conditions such as strong indirect light, strong ambient light, and smoke; and (4) recording live video from a projector's—rather than the camera's—point of view.

CR Categories: I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Imaging geometry, radiometry

Keywords: energy efficiency, low-power imaging, coded illumination, coded exposure, computational photography, 3D scanning

1 Introduction

When we capture an image under controlled lighting, the power of the light source matters a lot: all things being equal, brighter sources will send more photons to the sensor during an exposure, producing a brighter and less noisy image. The brightness of the source, however, is just one way to control how much light reaches the sensor of a computational imaging system. Modern systems use an arrangement of devices to transport light from a source to the scene (or from the scene to sensor) and these devices are often programmable—galvanometers [Mertz et al. 2012], digital micromirror devices [Nayar et al. 2004; Hitomi et al. 2011], liquid-crystal panels [Raskar et al. 2006], phase modulators [Damberg and Heidrich 2015], *etc.* This brings up a natural question: how should we program the spatio-temporal behavior of these devices to maximize an arrangement's *energy efficiency*, *i.e.*, the energy that can be transmitted from the source to the sensor for a given imaging task, power, and exposure time?

Studies of this problem began in the 1960s for the special case of arrangements with just three "active" components: a light source that is always turned on, a light-blocking mask that is controlled by a binary code, and a sensor [Ibbett et al. 1968; Decker and Harwit 1969]. The optimal sequence of codes for this case is derived from the Hadamard matrix [Harwit and Sloane 1979] and enjoys widespread use [Schechner et al. 2007], mainly because most conventional computational imaging systems are arranged this way.

Despite the ubiquity of these arrangements, two general principles have emerged in recent years. On one hand, masks are inefficient because they waste energy whenever they block a photon [Hoskinson et al. 2010]. On the other hand, masks confer unique abilities when arranged in layers along an optical path and programmed to change repeatedly in a single exposure. This is because photons can be blocked far more selectively this way, enabling light field displays [Lanman et al. 2010; Wetzstein et al. 2012], indirect-only photography [O'Toole et al. 2012], and several other imaging functionalities [O'Toole et al. 2014].

Unfortunately, neither the original Hadamard multiplexing theory nor its recent extensions [Cossairt et al. 2012; Mitra et al. 2014a; Mitra et al. 2014b] apply to multi-layer arrangements, or to arrangements that avoid masks altogether (*e.g.*, laser-based projectors [Damberg et al. 2014; Gupta et al. 2013]). As a result, the problem of computing energy-efficient codes for multi-layer ar-

rangements is poorly understood—and even less is known about how to design them for optimal energy efficiency. These problems are especially relevant for live imaging, where short exposures and low-power restrictions leave no room for wasting light.

In this paper we present, for the first time, a mathematical framework to derive energy-efficient codes for such arrangements and a novel experimental testbed for live energy-efficient imaging. We specifically focus on the family of two-layer arrangements, with one layer controlling illumination and the other controlling a sensor mask (Figure 2). Our work makes the following contributions:

Energy-efficient codes for probing light transport: We give a principled way to compute code sequences for existing structured-light-transport cameras [O'Toole et al. 2014], and transport probing systems more generally [O'Toole et al. 2012]. These arrangements use two programmable masks for imaging but their code sequences have so far been constructed heuristically. In contrast, our energy-efficient codes transmit more photons to the camera for a given light source power (*i.e.*, brighter images, all else being equal) and produce superior images for a given number of photons transmitted (*i.e.*, fewer artifacts, all else being equal). We confirm these gains experimentally with the DMD-based prototype of O'Toole *et al.* [2014].

Live energy-efficient probing with lasers and rolling shutters: Although energy-efficient codes do improve image quality, far bigger gains are possible by optimizing the arrangement itself. To this end, we demonstrate highly efficient transport probing with a novel combination of a low-power laser projector and a rolling shutter camera. We use this second prototype to demonstrate several never-seen-before capabilities such as (1) capturing epipolar-only and indirect-only live video without wasting any photons; (2) capturing live epipolar structured-light video of very bright scenes—even a light bulb that has been turned on; (3) reconstructing 3D objects in challenging conditions such as strong indirect light, strong ambient light, and smoke; and (4) recording live video from a projector's—rather than the camera's—point of view.

Unified mathematical model for redistributive projection: We show that despite the significant differences between our two prototypes, it is possible to describe both of them with the same underlying model. This makes it easy to compare their energy efficiency, and to define an efficiency criterion for code sequences that applies to a whole spectrum of projection devices—from conventional DMD projectors to rapidly-steerable laser beams. Mathematically, our model is expressed as a bound on the ℓ_1-norm and the ℓ_∞-norm of the illumination patterns realizable by a given device; intuitively, the ratio of these two norm-bounds measures the device's "redistribution ratio," *i.e.*, its ability to direct all photons to just one pixel for a specific interval of time.

Energy-efficient codes by homogeneous factorization: We show that finding energy-efficient codes under our model is equivalent to solving a *homogeneous* constrained matrix factorization problem. In this setting, the matrix to be factorized is task-dependent and known *a priori*, the constraints capture the physical constraints on masks and illumination patterns, and individual factors hold the energy-optimized code sequence of individual layers. We solve this factorization problem by minimizing an objective function that includes the projective tensor norm; we use the fast local optimizer of Haeffele *et al.* [2014] in our implemenetation, to compute energy-efficient codes for our DMD-based prototype.

Impulse illumination is globally optimal: We show that the global minimum of our objective function—across all two-layer code sequences and redistribution ratios—can be derived in closed form and is actually very simple: it is a sequence of impulse illuminations produced by a projector having the smallest possible redistribution ratio for the exposure time, *i.e.*, it emits all photons into a tight beam that can be redirected very quickly. A key corollary of this result is that our laser-based prototype, whose illumination sequence is fixed by the hardware and not under our control, is globally optimal for the tasks of epipolar-only and indirect-only imaging. Moreover, global optimality holds even in the presence of ambient light. For the other tasks we consider, however, a significant gap from global optimality still exists. This can only be closed with yet-unavailable hardware for more flexible laser projection and electronic masking.

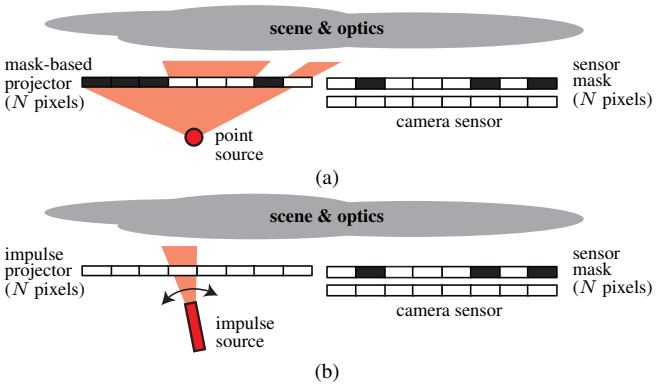

Figure 2: *Two-layer arrangements considered in this work. (a) At the low end of the efficiency spectrum are projectors that always spread light to all pixels on a mask (*i.e.*, redistribution ratio is N). These projectors waste energy because they create patterns by blocking light—and that light is not used for imaging. (b) At the other extreme are "impulse" projectors, able to concentrate all their energy to just one pixel (*i.e.*, redistribution ratio is 1). In between are arrangements that can concentrate light partially, which we also analyze.*

Energy-efficient codes for redistributive projectors: Of course, between the two extremes of ideal impulse projection (minimum redistribution ratio) and conventional mask-based projection (maximum redistribution ratio) lies a whole spectrum of arrangements that redistribute light partially [Hoskinson et al. 2010; Mertz et al. 2012; Gupta et al. 2013; Damberg and Heidrich 2015]. We explore this spectrum briefly through simulations, noting that a rapid "phase transition" seems to occur in our energy-efficient codes, with pure impulses on one end switching to dense codes on the other.

2 Lights, Masks and Energy Efficiency

We begin by revisiting the oft-used concepts of an illumination pattern and a mask, expressing them in terms of physical units and constraints. These lead to a definition of energy efficiency for two-layer code sequences, for the purpose of transport probing.

2.1 Redistributive Projection and Sensor Masking

We assume that the projector's light source is always on and emits light at a constant rate Φ, measured in Watts. Illuminating a scene for an exposure time T means that the total energy generated by the source is ΦT and given in Joules.

We use illumination vectors $\mathbf{l}$ to represent actual patterns. Each element of $\mathbf{l}$ gives the energy emitted from a specific projector pixel during the exposure time. The sum of all elements of $\mathbf{l}$ cannot exceed the energy generated by the source:

$$0 \leq \mathbf{l}, \quad \|\mathbf{l}\|_1 \leq \Phi T, \tag{1}$$

where $\|\cdot\|_1$ denotes the ℓ_1-norm of a vector.

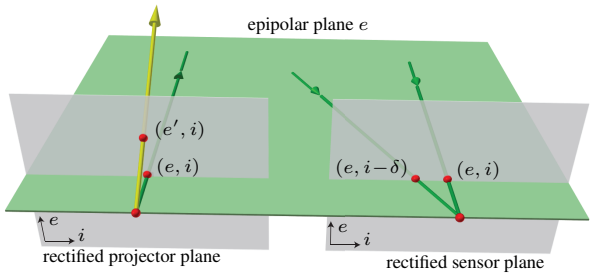

Figure 3: *Transport probing allows us to block, permit, or attenuate the energy transmitted from individual projector pixels to individual sensor pixels—regardless of scene content. This is specified by a probing matrix* $\mathbf{\Pi}$ *that has an element for every such pair. We adopt the rectified stereo representation of O'Toole* et al. *[2014] to parameterize individual pixels and the elements of* $\mathbf{\Pi}$*. In this representation, identical rows on the projector and sensor belong to the same epipolar plane; rays through pixels with the same coordinates are parallel in 3D; and intersecting rays always lie on the same epipolar plane and are assigned a non-zero disparity* δ*.*

The redistribution ratio of a projector A key measure of projector flexibility is how well it can "channel" to individual pixels all the energy generated by the source. We express this ability as an upper bound on the individual elements of $\mathbf{l}$:

$$\|\mathbf{l}\|_\infty \leq \Phi T/\sigma, \tag{2}$$

where $\|\cdot\|_\infty$ is the ℓ_∞-norm, giving the largest element of a vector, and σ is a projector-specific parameter we call the redistribution ratio, that is equal to the ratio of the bounds in Equations 1 and 2. For an N-pixel projector, this ratio takes values between 1 and N and models energy redistribution: the larger its value is, the lower the energy we can send through any one pixel, and the more energy we waste when projecting a pattern with just one pixel turned on (or just a few).

The specific value of σ depends on the projection technology. At the far end of the range, with $\sigma = N$, are conventional projectors (Figure 2a). These projectors use mirrors and lenses to distribute light evenly over spatial light modulator, which may then block all or part of it. The near end of the range, with $\sigma = 1$, represents an idealized projector that is perfectly efficient (Figure 2b). This projector can send all its light through just one pixel for the *entire* exposure time T but can also distribute it—without any blocking—according to an arbitrary illumination $\mathbf{l}$. Between these two extremes lies a whole spectrum of projection technologies that approach this ideal to a greater or lesser extent [Hoskinson et al. 2010; Gupta et al. 2013; Damberg et al. 2014; Damberg and Heidrich 2015].

Norm bound on illumination vectors The ℓ_1 and ℓ_∞ bounds on $\mathbf{l}$ can be written more concisely as

$$0 \leq \mathbf{l}, \quad \|\mathbf{l}\|_{\dagger\sigma} = \max\left(\frac{\sigma\|\mathbf{l}\|_\infty}{\Phi}, \frac{\|\mathbf{l}\|_1}{\Phi}\right) \leq T, \tag{3}$$

where $\|\cdot\|_{\dagger\sigma}$ is the max of two norms and therefore also a norm. These bounds are useful in three ways. First, we can optimize arrangements with very different light redistribution properties by adjusting the redistribution ratio. Second, the dependence on exposure time makes a distinction between systems that conserve energy and those that merely conserve power (but require on long exposures [Gupta et al. 2013]). Third, they explicitly account for timescale-dependent behavior (*e.g.*, raster-scan projectors, like the one used in our laser-based prototype, can act like a beam, light sheet, or point source depending on T).

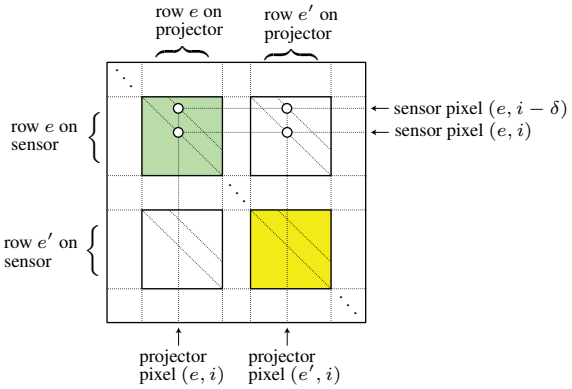

Figure 4: *The probing matrix* $\mathbf{\Pi}$*. For a raster-scan order of projector pixels and sensor pixels, the elements controlling energy transmission within an epipolar plane always occupy diagonal blocks (green for epipolar plane* e *and yellow for* e'*). Circles indicate the elements controlling transmission for the pixels in Figure 3.*

Norm bound on mask vectors Programmable masks can transmit, block, or attenuate light in a structured way [Takhar et al. 2006; Veeraraghavan et al. 2011; Hitomi et al. 2011; Gu et al. 2013]. Here we use mask vectors $\mathbf{m}$ to denote energy attenuation on the sensor. Each element of $\mathbf{m}$ is a unit-less scalar between zero and one describing attenuation at a specific sensor pixel:

$$0 \leq \mathbf{m}, \quad \|\mathbf{m}\|_\infty \leq 1. \tag{4}$$

2.2 Transport Probing

Masked sensors are clearly less energy efficient than mask-less ones because masks may block incident photons. Despite this apparent limitation, programmable masks come with an important advantage: by changing masks *and* illuminations repeatedly in a single exposure, we can capture images that are impossible to capture with an unmasked sensor in one shot. This general technique—called *transport probing* [O'Toole et al. 2012]—expands the capabilities of an imaging system at a cost of reduced energy efficiency. We briefly review this technique here, as it defines the tasks whose energy efficiency we wish to optimize.

Transport probing uses a coded sequence of illuminations and masks to control, in a scene-independent way, the energy transmitted from the projector to the sensor (Figure 3). This control is typically expressed as a *probing matrix* $\mathbf{\Pi}$ whose individual elements specify energy attenuation from a specific projector pixel to a specific sensor pixel (Figure 4). Just like masks, probing matrices contain unit-less scalars in the range $[0, 1]$. Intuitively, each matrix $\mathbf{\Pi}$ can be thought of as describing a specific imaging task, and is always provided. Figure 5 shows several examples of probing matrices we use in our experiments.

Energy efficiency of illumination-mask sequences Probing matrices describe attenuation in relative terms and do not tell us how much energy can actually be transmitted from the projector to the sensor. Indeed, the energy efficiency of transport probing depends on the precise sequence of illuminations and masks used.

More specifically, we can realize any rank-1 probing matrix $\mathbf{\Pi}$ using just one illumination and one sensor mask (Figure 6):

$$\gamma\,\mathbf{\Pi} = \mathbf{m}\,\mathbf{l}^{\mathrm{T}}, \tag{5}$$

where vectors $\mathbf{l}$ and $\mathbf{m}$ depend on matrix $\mathbf{\Pi}$ and satisfy all physical constraints, and γ is a scalar that is measured in Joules. Intuitively,

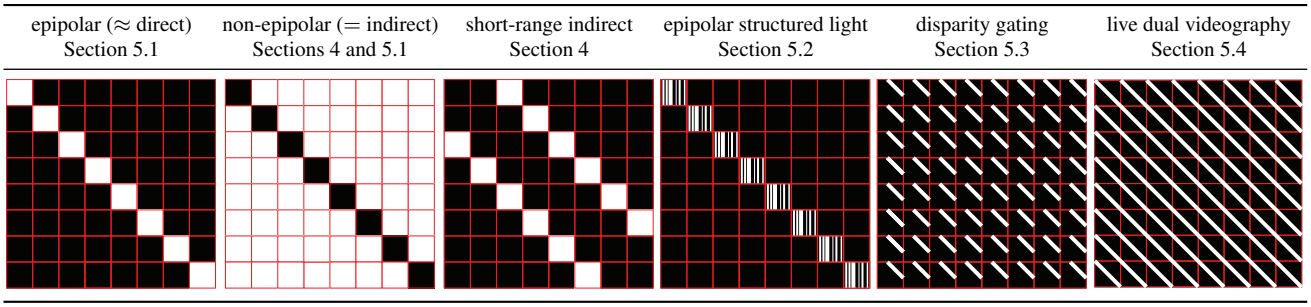

Figure 5: *Probing matrices used in our live video experiments. We follow the conventions of Figure 4 for a projector and a sensor that have just eight rows of pixels (but many columns). Individual elements range from zero (black) to one (white). Red lines indicate row boundaries in the raster-scan ordering of individual pixels. These lines partition* $\mathbf{\Pi}$ *into contiguous "epipolar blocks," with each block controlling energy transmission from projector pixels on a specific epipolar plane to sensor pixels on another epipolar plane.*

γ converts the unit-less elements of $\mathbf{\Pi}$ into actual energies that can be potentially transmitted during the exposure time. We call γ the *energy efficiency* of the illumination-mask pair $\mathbf{l}$ and $\mathbf{m}$.

Higher-rank probing matrices require changing illuminations and masks $K > 1$ times during the exposure time. Mathematically this is equivalent to expressing matrix $\mathbf{\Pi}$ as a sum of outer products, subject to the physical constraints on illuminations and masks:

$$\gamma\,\mathbf{\Pi} \;=\; \sum_{k=1}^{K} \mathbf{m}_k\,(\mathbf{l}_k)^{\mathrm{T}}, \tag{6}$$

$$0 \le t_k, \quad \sum_{k=1}^{K} t_k \;\le\; T, \tag{7}$$

$$0 \le \mathbf{m}_k, \quad \|\mathbf{m}_k\|_\infty \le 1, \quad 0 \le \mathbf{l}_k, \quad \|\mathbf{l}_k\|_{\dagger\sigma} \le t_k, \tag{8}$$

where the mask $\mathbf{m}_k$ and illumination $\mathbf{l}_k$ are used for an interval t_k of the total exposure time, and scalar γ is the energy efficiency of the whole sequence.

2.3 Homogeneous Factorization

The greater the energy efficiency of a sequence, the more total energy will flow from the projector to the camera in any given scene. We therefore seek illumination-mask sequences that maximize γ.

Re-writing Eq. (6) in the form of a matrix factorization equation

$$\gamma\,\mathbf{\Pi} \;=\; [\underbrace{\mathbf{m}_1\;\mathbf{m}_2\;\cdots\mathbf{m}_K}_{\text{masks }\mathbf{M}}][\underbrace{\mathbf{l}_1\;\mathbf{l}_2\;\cdots\mathbf{l}_K}_{\text{illuminations }\mathbf{L}}]^{\mathrm{T}} \tag{9}$$

leads to a homogeneous factorization problem whose goal is to find the masks $\mathbf{M}$, illuminations $\mathbf{L}$, and timeslices t_k that maximize energy efficiency:

$$\max_{\gamma,\mathbf{M},\mathbf{L},t_1,\ldots,t_K} \quad \gamma \tag{10}$$

$$\text{subject to} \quad \gamma\,\mathbf{\Pi} \;=\; \mathbf{M}\mathbf{L}^{\mathrm{T}} \tag{11}$$
$$0 \le \mathbf{m}_k, \quad \|\mathbf{m}_k\|_\infty \le 1$$
$$0 \le \mathbf{l}_k, \quad \|\mathbf{l}_k\|_{\dagger\sigma} \le t_k$$
$$0 \le t_k, \quad \sum_{k=1}^{K} t_k \le T.$$

The energy efficiency of a particular factorization—and thus the solution to this optimization problem—depends to a large degree on a projector's ability to redistribute light, *i.e.*, on its redistribution ratio. See Figure 6 for a detailed illustration.

3 Homogeneous Low-Rank Factorization

While the optimization in Eq. (10) is hard to perform directly, it is possible to relax it into the following form:

$$\min_{\mathbf{M},\mathbf{L}} \quad \|\mathbf{\Pi} - \mathbf{M}\mathbf{L}^{\mathrm{T}}\|_F^2 + \lambda \sum_{k=1}^{K} \|\mathbf{m}_k\|_\infty \|\mathbf{l}_k\|_{\dagger\sigma} \tag{12}$$
$$\text{subject to} \quad 0 \le \mathbf{m}_k, 0 \le \mathbf{l}_k$$

where λ is a regularization parameter that balances energy efficiency and the reproduction of $\mathbf{\Pi}$. Details of this relaxation, which absorbs timeslices and inverse energy efficiency γ^{-1} into the second term of Eq. (12), are in the supplemental materials. Its main intuition is to convert the homogeneous maximization of Eq. (10) into an inhomogeneous one by seeking the masks and illuminations with the *minimum* sum of norm products. Once this decomposition is found, we can "brighten" matrix $\mathbf{\Pi}$, and thus find γ, by re-scaling $\mathbf{M}$ and $\mathbf{L}$ to saturate their upper-bound constraints.

The relaxation in Eq. (12) has received some attention recently in computer vision and machine learning [Bach et al. 2008; Haeffele et al. 2014]. Here we focus on aspects of that work that are specific to our optimization problem.

The projective tensor norm By leaving the sequence length K unconstrained and dropping non-negativity constraints, Eq. (12) becomes equivalent to a minimization of the following general objective [Haeffele et al. 2014]:

$$\min_{\mathbf{X}} \quad \|\mathbf{\Pi} - \mathbf{X}\|_F^2 + \lambda h(\mathbf{X}) \tag{13}$$

where function $h(\mathbf{X})$ is the *projective tensor norm*, defined as

$$h(\mathbf{X}) \;=\; \min_{\mathbf{X}=\mathbf{M}\mathbf{L}^{\mathrm{T}}} \Big\{ \sum_{k=1}^{K} \|\mathbf{m}_k\|_p \|\mathbf{l}_k\|_q \Big\} \tag{14}$$

$$\;=\; \min_{\mathbf{X}=\mathbf{M}\mathbf{L}^{\mathrm{T}}} \Big\{ \frac{1}{2} \sum_{k=1}^{K} \|\mathbf{m}_k\|_p^2 + \|\mathbf{l}_k\|_q^2 \Big\} \tag{15}$$

with $p = \infty$ and $q = \dagger\sigma$ according to Eq. (12).

3.1 Implications of the Theory

Impulse illumination is globally optimal A closed-form solution to Eq. (12) exists for the special case of impulse projectors. This is because the norms $\|\mathbf{l}_k\|_1$ and $\|\mathbf{l}_k\|_{\dagger\sigma}$ in that equation always coincide when the projector's redistribution ratio is equal to

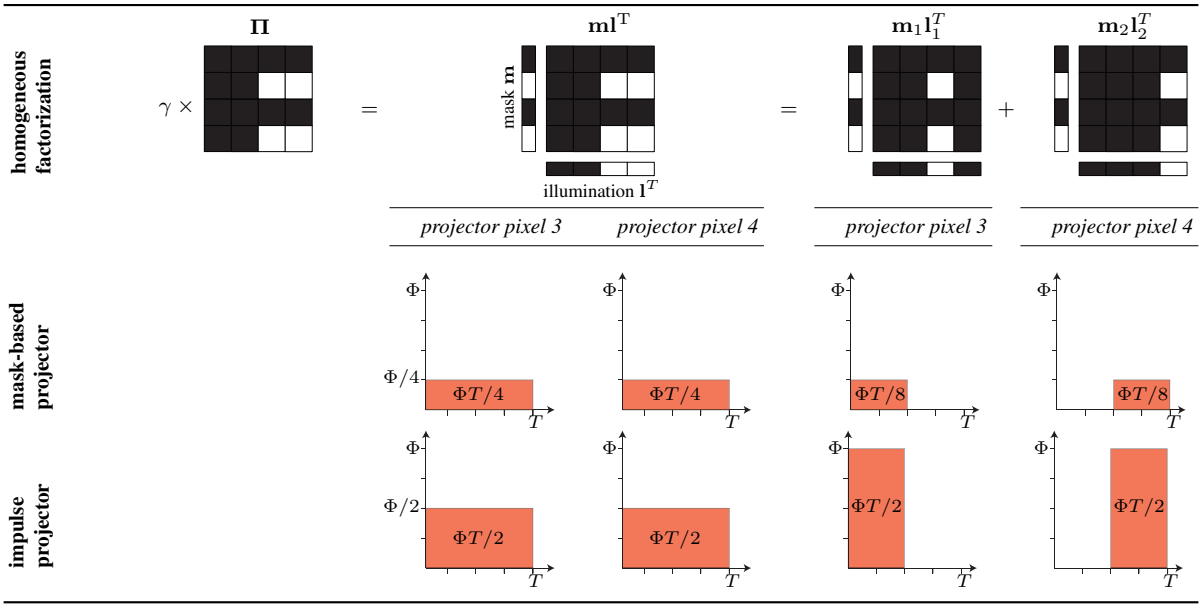

Figure 6: ***Top row:*** *Realizing a rank-1 probing matrix of size 4×4 in two ways—with a single illumination-mask pair (middle) and with a sequence of two illuminations and masks (right).* ***Middle row:*** *Plots of power versus time for the "active" pixels in a mask-based projector. The projector's redistribution ratio is 4 because no more than a quarter of the power can be emitted from any given pixel. Although the single illumination-mask pair is more efficient than the length-two sequence ($\gamma = \Phi T/4$ versus $\gamma = \Phi T/8$) neither makes full use of the source's energy.* ***Bottom row:*** *With an impulse projector, on the other hand, both are optimal ($\gamma = \Phi T/2$).*

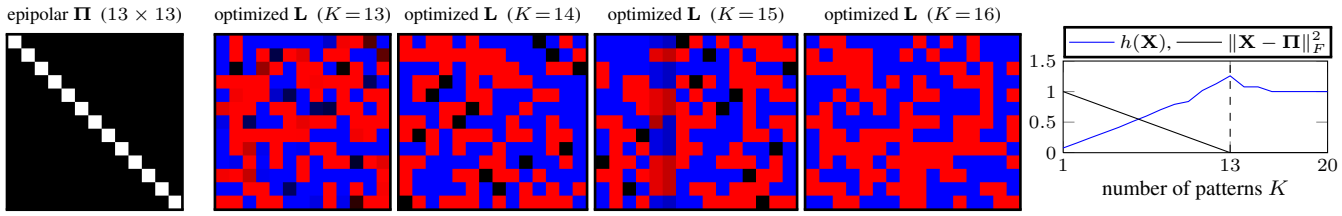

Figure 7: *Codes optimized for different sequence lengths K. Here we dropped the non-negativity constraint and computed codes* $\mathbf{M}$ *and* $\mathbf{L}$ *that require two photos to implement efficiently on a DMD-based projector* (i.e., *one for the positive and one for the negative elements of* $\mathbf{L}$*).* ***Left:*** *A probing matrix.* ***Middle:*** *Optimized illuminations for various sequence lengths (mask matrices omitted for brevity). Positive and negative elements are color-coded in red and blue, respectively.* ***Right:*** *Sequence length versus energy efficiency as measured by the projective tensor norm (lower is better). Note that energy efficiency starts improving once* $\mathbf{\Pi}$ *can be reconstructed exactly,* i.e., *for $K \geq 13$.*

one. In particular, Bach *et al.* [2008] showed that for $p = \infty$ and $q = 1$, the projective tensor norm—which is equal to the global minimum of Eq. (12)—is given by the sum of the ℓ_∞ norms of the mask vectors. The factorization corresponding to this minimum is

$$\mathbf{L} = \mathbf{I}_{N \times N}, \quad \mathbf{M} = \mathbf{\Pi}\,\mathbf{L}, \tag{16}$$

where $\mathbf{I}_{N \times N}$ is the identity matrix. In other words, the illumination vectors that maximize energy efficiency are *impulses*.

Epipolar illumination is globally optimal for epipolar and non-epipolar imaging The global optimality of impulse projectors has another, more practical implication: scanning-based laser projectors that can concentrate all their energy onto individual scanlines are also globally optimal for probing—but only when their scanlines are along the *epipolar planes*, and only for probing matrices whose elements do not vary within an epipolar block (*e.g.*, the first three matrices in Figure 5). In that case, each epipolar line can be treated as a "pixel" without loss of generality, and each epipolar block of $\mathbf{\Pi}$ can be treated as a single element. The globally-optimal solution is again given by Eq. (16), with "impulses" corresponding to a single epipolar line turned on.

Epipolar illumination and epipolar masking confer robustness to ambient light A second practical implication of Eq. (16) concerns probing matrices $\mathbf{\Pi}$ that are permutations of the identity matrix: since each pixel is unmasked exactly once in $\mathbf{M}$, the globally-optimal solution also globally minimizes the time that individual pixels are exposed on the sensor. In other words, *impulse illumination and impulse masking* simultaneoulsy maximize the energy efficiency of probing and minimize the impact of ambient light. Similarly, *epipolar illumination and epipolar masking* are optimal for epipolar probing in the presence of ambient light—a principle we directly exploit with our laser-based prototype (Section 5).

Code optimization algorithm for non-impulse projectors No closed-form solution is known for the projective tensor norm when σ is greater than one. We use the *structured low-rank matrix factorization* algorithm of Haeffele *et al.* [2014] to locally optimize this norm with convergence guarantees. The base requirement of their algorithm is the ability to evaluate the *proximal operator*. For a norm $\|\cdot\|_p$ this operator is defined as

$$\mathbf{prox}_{\lambda\|\mathbf{x}\|_p}(\mathbf{v}) = \mathrm{argmin}_{\mathbf{x}} \left(\|\mathbf{x}\|_p + \frac{1}{2\lambda}\|\mathbf{x} - \mathbf{v}\|_2^2 \right). \tag{17}$$

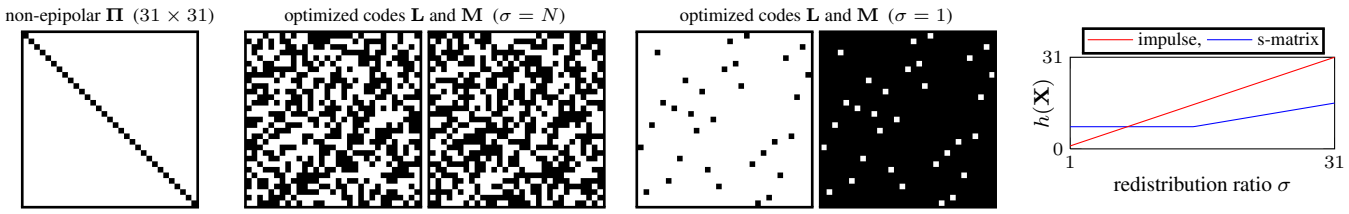

Figure 8: *Optimized codes for different values of the redistribution ratio σ. For a DMD-based projector ($\sigma = 31$), our algorithm returns masks $\mathbf{M}$ and illuminations $\mathbf{L}$ that are complementary S-matrices. For the impulse projector ($\sigma = 1$) we get impulses. The projective tensor norm plot of the S-matrix and impulse solutions is shown on the right for different values of σ (lower is better).*

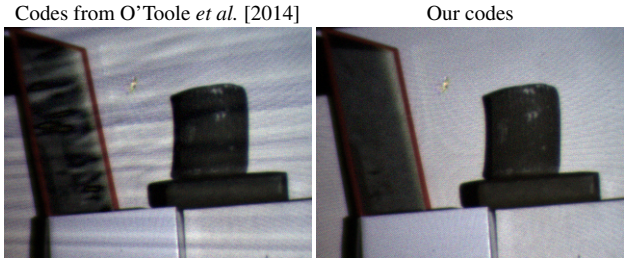

Figure 9: *Epipolar imaging with the DMD-based prototype in O'Toole* et al. *[2014]. The scene has a mirror on the left and a translucent candle on top of a box on the right, all in front of a white wall (see Figure 12(a) for a conventional photo). We generated codes using their randomized algorithm and ours for $K = 512$. Both codes are equally energy efficient since they produce images of similar brightness. Our codes, however, are not just efficient, they approximate $\mathbf{\Pi}$ well too. This yields images with far fewer artifacts (note the bright streaks on the left image).*

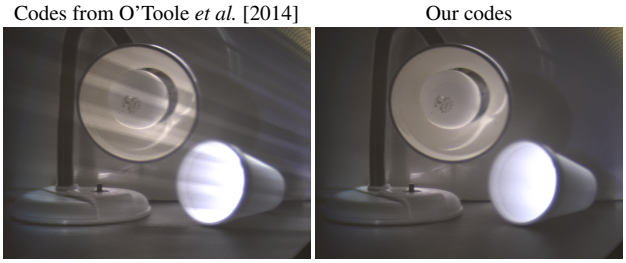

Figure 10: *Non-epipolar imaging using our codes versus those of O'Toole* et al. *[2014]. The scene consists of a styrofoam cup and a lamp (turned off). We generated a short sequence of codes ($K = 32$) using their randomized algorithm and ours, and captured the images shown above. Again, both codes have similar energy efficiency but ours produce no discernible artifacts (note the bright diagonal streaks on the left image).*

The proximal operator for $\|\cdot\|_{\dagger\sigma}$ has a simple form when $\sigma = N$ and a very efficient algorithm exists for computing it [Parikh and Boyd 2014]. We use this approach to compute optimized code sequences for our DMD-based prototype (Section 4). For projectors with redistribution ratios between 2 and $N-1$, we found no simple solution. We evaluate Eq. (17) explicitly in such cases by solving a constrained minimization problem. To impose non-negativity constraints on masks and patterns we replace each negative component of the proximal operator's output with zero [Xu and Yin 2013; Parikh and Boyd 2014].

4 DMDs for Live Energy-Efficient Imaging

DMDs offer flexible masking but have two important limitations: they are very inefficient when used for light projection ($\sigma = N$) and they can only support short code sequences for live imaging (*e.g.*, $K \leq 96$ for the Texas Instruments LightCrafter DMD). This makes the design of energy-efficient codes all the more important, so that images are as bright and artifact-free as possible.

Optimized codes for epipolar imaging Probing with the epipolar matrix $\mathbf{\Pi}$ in Figure 5 yields an image that is almost exclusively due to direct surface reflections [O'Toole et al. 2014]. This is because $\mathbf{\Pi}$ allows energy transmission only between pixels on the same epipolar plane, and this transmission is predominantly caused by direct surface reflection in most scenes.

The algorithm of O'Toole *et al.* [2014] uses random binary illuminations for matrix $\mathbf{L}$ and derives the masks $\mathbf{M}$ from $\mathbf{L}$ and $\mathbf{\Pi}$. This guarantees that the approximation error $\mathbf{\Pi} - \mathbf{M}\mathbf{L}^{\mathrm{T}}$ goes to zero as $K \to \infty$ but can produce poor results for practical values of K.

Figure 7 shows the illumination codes from our algorithm, using Eq. (12). These codes are far from random: for probing matrices whose size is divisible by four, we obtain illuminations $\mathbf{L}$ similar to the Hadamard matrix [Harwit and Sloane 1979]. Moreover, $\mathbf{M}\mathbf{L}^{\mathrm{T}}$ converges to the probing matrix as K approaches $\mathbf{\Pi}$'s size. This has a clear, positive impact on image quality (Figure 9). We observe the same convergence behavior for other sizes too (where Hadamard matrices do not exist); we do find, however, that energy efficiency can be improved even further by computing code sequences that are *longer* than $\mathbf{\Pi}$'s size (Figure 7c-e).

Optimized codes for non-epipolar imaging Figure 10 illustrates the visual quality improvement from our optimized codes in this task. The task requires probing with the non-epipolar matrix of Figure 5 which guarantees that no direct light reaches the sensor. Short illumination-mask sequences from the algorithm of O'Toole *et al.* [2014] result in strong visual artifacts (bright diagonal streaks along epipolar lines). These occur because the energy that is transmitted from projector to sensor is attenuated non-uniformly over the sensor plane (*i.e.*, $\mathbf{M}\mathbf{L}^{\mathrm{T}}$ is a poor approximation of $\mathbf{\Pi}$). Our optimized codes, on the other hand, yield bright and artifact-free images for a short sequence of the same length. These codes are shown in Figure 8.

Optimized codes for high-rank probing Many matrices useful for probing have rank much higher than the maximum attainable value of K. In such cases, the matrix $\mathbf{\Pi}$ can only be approximated by a rank-K approximation of $\mathbf{\Pi}$. Figure 11 compares our codes to those obtained by the approach of O'Toole *et al.* [2012] for the hitherto not studied task of *short-range indirect imaging*. This task involves probing with the corresponding matrix from Figure 5, which blocks all energy transmissions except those occurring between any two nearby epipolar planes.

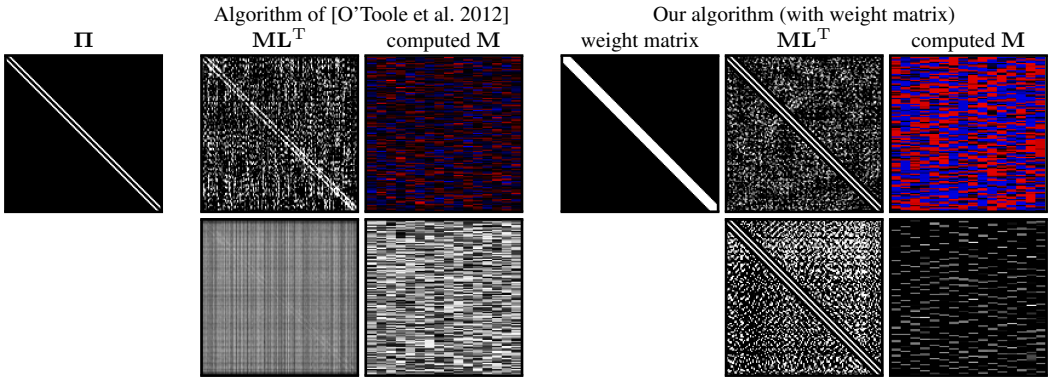

Figure 11: *Computing a $K = 16$ code sequence for a rank-128 short-range indirect probing matrix. The top row shows the mask matrix computed without enforcing non-negativity whereas the bottom row shows the matrix with those constraints enfored. We can use this matrix to capture images with only "short-range indirect" contributions,* i.e.*, light that contains neither direct reflections nor contributions transported between distant epipolar planes. We are aware of no techniques to capture such images, either by factorizing $\mathbf{\Pi}$ or by other means. In particular, the randomized algorithm of O'Toole* et al. *[2012; 2014] fails to produce usable codes because $\mathbf{\Pi}$'s off-diagonals are reconstructed poorly. By incorporating the weight matrix shown into the Frobenius norm of Eq.* (12)*, our algorithm produces a low-rank approximation of $\mathbf{\Pi}$ whose off-diagonals are reconstructed very well. Use of these weights was partly inspired by [Lanman et al. 2010]. Figures 12c shows images captured with our DMD-based prototype using this code.*

Figure 12 shows captured images for several high-rank probing tasks. These tasks were impossible to perform previously because no efficient low-rank factorizations of $\mathbf{\Pi}$ were known (*i.e.*, with enough energy efficiency for a useful image signal).

Implementation details The projective tensor norm is a prior that naturally favors binary codes. Nevertheless, the codes computed by our algorithm are not binary in general. We implement non-binary codes by temporally dithering each code with the DMD, *i.e.*, projecting a sequence of binary codes that achieve the desired intensity value. This increases the number of codes but does not affect energy efficiency.

Ignoring non-negativity constraints in Eq. (12) can significantly increase the energy efficiency of code sequences. Codes with both positive and negative values, however, are not physically realizable. We follow the approach of O'Toole *et al.* [2012; 2014] to implement such codes: we convert a length-K illumination-mask sequence into two such sequences, one for their negative and one for their non-negative elements, capture one image for each sequence, and subtract them in software.

Energy-efficient codes for redistributive projectors Finding globally-optimal codes for general redistributive projectors remains an open problem. Nevertheless, we can draw interesting conclusions about the characteristics of (locally-optimized) energy-efficient codes for different redistribution ratios and different probing matrices. Take, for instance, the case of non-epipolar $\mathbf{\Pi}$. For $\sigma = N$, our algorithm automatically returns illuminations $\mathbf{L}$ similar to an S-matrix [Schechner et al. 2007] (Figure 8). At the other extreme of $\sigma = 1$ it returns the impulse illuminations predicted by Eq. (16). Between these two extremes, we found no solution that improves upon either the S-matrix or the impulse illuminations. This suggests a rapid flip in the energy efficiency of these two code sequences.[1]

[1]One can draw an analogy to the work of Cossairt *et al.* [2012] who observed similar behavior in their study of computational imaging in read-noise versus shot-noise limited settings.

5 Lasers for Live Energy-Efficient Imaging

Scanning-based laser projectors operate on a different principle than mask-based projectors. These projectors use a MEMS mirror to steer a laser beam and raster-scan across the scene, sending all their energy to just one pixel at a time. We consider these projectors at the more granular timescale of a single *scanline*, where they act as impulse projectors whose "impulses" are scanlines (Section 3.1). By aligning the projector's scanlines with the rows of a rolling-shutter camera—whose ability to mask individual rows electronically provides a sequence of "impulse" masks—we obtain a system with very high energy efficiency for many transport probing tasks.

Our prototype is portable and consists of nothing more than an off-the-shelf camera and off-the-shelf projector whose only modifications are synchronization electronics—no passive or active optical components are required beyond a simple color filter. This avoids the need for sophisticated hardware devices and extra optics, all of which introduce aberrations, limit working volume, reduce light throughput, and adversely impact size, weight and cost.

Laser projectors for "impulse" illumination We use a Microvision ShowWX+ pico-projector (resolution: 848×480, frame rate: 60Hz, light power: 5 lumens per color channel) for light projection. Fortunately, even though the projector's raster-scan path is fixed in hardware, it implements a globally-optimal sequence of illumination codes for many of the probing tasks we consider (Section 3.1).

Rolling-shutter cameras for "impulse" masking For outdoor experiments we use an IDS Imaging UI-3250CP-M-GL (1600×1200) monochrome CMOS rolling-shutter camera, fitted with a 532nm laser line cut filter (centered on the projector's green laser wavelength) with a passband of 10nm. For indoor experiments we use an IDS Imaging UI-3240CP-C-HQ (1280×1024) color CMOS camera. In both cases, we use the cameras' rolling shutter to implement sensor masking [Muller 2012].

At any point in time, the rolling shutter can expose either a single row or a band of rows. The rolling shutter is triggered by the VSYNC signal generated by the projector. Figure 13 shows our cameras' timing diagram. We use t_p to denote the timeslice for which the projector dwells on a single scanline. The speed at which

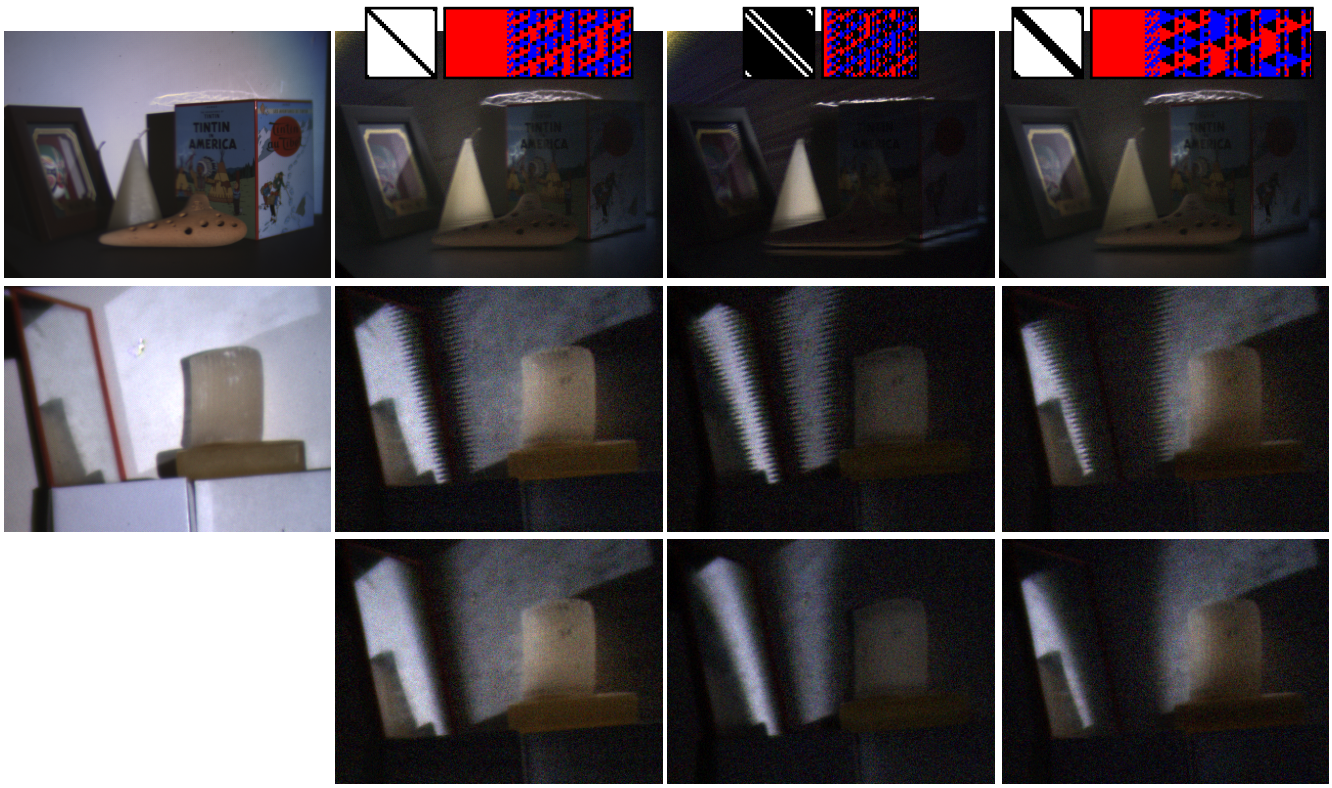

(a) conventional photos (b) robust non-epipolar indirect (c) short-range indirect (d) long-range indirect

Figure 12: *New transport-probing abilities enabled by homogeneous factorization of high-rank matrices. Probing matrix* $\mathbf{\Pi}$ *and mask matrix* $\mathbf{M}$ *(after dithering) are shown above each column. (a) We show results for two scenes: (Top) A scene containing, from left to right, a Chinese mask within a display case; a translucent candle; an ocarina; and a tin box. (Middle and bottom) A scene with a mirror and a translucent candle. (b) Making non-epipolar probing robust to geometric misalignments by building robustness into matrix* $\mathbf{\Pi}$ *itself. Here we set the elements of* $\mathbf{\Pi}$ *to zero within a small band of the diagonal to ensure no direct light "leaks" into the photo, and compute a code sequence of length* $K = 352$*. (c) Short-range indirect imaging captures interreflections, subsurface scattering, and a bright vertical band of mirror reflections. We used* $K = 176$*. (d) Long-range indirect photos captured using a code sequence of length* $K = 416$*. Since subsurface scattering is a short- to mid-range transport phenomenon, candles appear darker compared to (c). In the bottom row we used* $4\times$ *longer sequences to improve image quality. Note that all the imaging tasks above are considerably "harder" than epipolar-only and non-epipolar probing, requiring less energy-efficient codes: exposure time was* $T = 1/3sec$ *for all images, about* $10\times$ *longer than in Figures 9 and 10.*

the rolling shutter progresses down the rows of the image (t_c) is determined by the pixel clock frequency; we choose this frequency and the focal length of camera lens so that the downward velocity of exposed camera rows matches the downward velocity of the projector's scanline. Increasing the row exposure (t_e) increases the thickness of the band of camera rows exposed for each projector scanline. Changing the delay (t_o) between the VSYNC signal from the projector and the trigger signal passed to the camera changes the offset between the illuminated row on the projector and the imaged row(s) on the camera.

5.1 Epipolar and Non-Epipolar Imaging

To probe with an epipolar matrix $\mathbf{\Pi}$ (Figure 5), we position the camera and projector so that they emulate the rectified stereo configuration shown in Figure 3 as precisely as possible. Low-distortion camera lenses are critical for this alignment; we use the Lensagon CVM0411ND varifocal lens in all our experiments.

To capture live epipolar video, we project a plain white pattern continuously and choose sensor masks so that for each projector scanline, only the corresponding row of sensor pixels is exposed. This has the effect of transmitting energy from a projector row to a sensor row on the same epipolar plane, while blocking everything else. We do this by matching the camera's row exposure time (t_e) and row readout time (t_c) to the time the projector dwells on a single scanline (t_p).

By changing the timing parameters on the camera so that the masks are inverted, the same setup can be used for probing with non-epipolar $\mathbf{\Pi}$ (Figure 5). This is done by setting t_e to be t_p less than the projector cycle time and adding t_p to camera trigger offset t_o. Thus, at any point in time, every sensor row is exposed except the one corresponding to the current projector scanline. This blocks all energy transmission between projector rows and sensor rows on the same epipolar plane while leaving all other light unblocked.

Ideally, we would be able to configure the rolling shutter so that only the rows of sensor pixels illuminated by the projector at any timestep would be exposed (Figure 13). In practice, the projector we use generates distorted scanlines that are not absolutely straight. Additionally, we observe synchronization jitter and small perturbations in the trajectory of the projector's laser during each exposure cycle. This means that the region in the camera image corresponding to a projector scanline is constrained to lie inside a narrow band in the image, not along a single row of pixels. To accommodate these bands, during epipolar imaging we thicken the region of un-

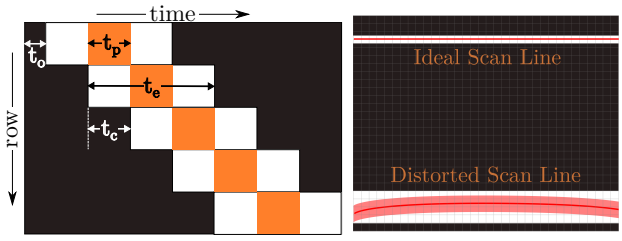

(a) projector and camera timing (b) effects of distortion

Figure 13: *(a) How to combine masking with a rolling shutter camera and raster scanning with a laser projector. At each timestep (of duration t_p), the projector illuminates a single scanline (orange). The camera rolling shutter exposes one or more rows of pixels (white). The masks are defined by three controllable parameters: the exposure time t_e; the time it takes the rolling shutter to read a row of pixels t_c; and the offset between the projector sync output and the camera trigger input t_o. (b) Ideally, a single projector scanline corresponds to a single camera row (top). In practice, because of distortion and jitter, each scanline corresponds to a band of sensor pixels (bottom).*

blocked pixels in each mask by increasing the pixel exposure time t_e and adjusting the trigger offset t_o. As a result, some short-range indirect light may leak into the epipolar image.

Figure 14 shows frames from live epipolar and non-epipolar video captured with our prototype, for scenes with significant global light transport. This prototype performs non-epipolar imaging at least four times more efficiently than O'Toole *et al.*'s [2014] because DMD-based codes must block 75% of the energy from the projector's light source. Moreover, it captures epipolar video at full video rates and in a single readout, whereas energy-efficient codes for DMD-based epipolar imaging require two readouts and a computational subtraction step (Section 4), and output video at half the video rate [O'Toole et al. 2014].

5.2 Epipolar Structured Light

Combining epipolar imaging with the projection of structured-light patterns is equivalent to probing with the epipolar structured light matrix $\mathbf{\Pi}$ in Figure 5. By probing with $\mathbf{\Pi}$ instead of projecting structured light patterns conventionally, we gain two benefits: robustness to global light transport and robustness to ambient illumination.

Robustness to global light transport Global transport effects like interreflections and scattering can cause severe, systematic errors when using structured light to reconstruct scenes with concavities and objects made of optically challenging materials. Epipolar structured light blocks the majority of this indirect light, resulting in more accurate reconstructions of difficult scenes where global transport effects are pronounced. The reconstructions obtained in Figure 15 show how effective epipolar structured light is in the face of interreflections and subsurface scattering. Figure 19 shows how our method improves reconstruction quality in the presence of volumetric scattering.

Structured light in bright sunlight Active light sources are typically orders of magnitude weaker than daylight. Since conventional cameras have limited dynamic range, the signal from the source is overwhelmed by ambient illumination when regular imaging is used. This makes active illumination methods, such as structured light, hard to use outdoors. Even though in theory our prototype minimizes pixel exposure only for the task of epipolar probing (Section 3.1), in practice the exposure time of individual scanlines is sufficiently small that the energy received from ambient light is small compared to the energy transmitted from the projector—even under challenging illumination conditions and even when the projector light source is weak. Refer to Figures 15, 16, and 17 for some example reconstructions performed under challenging ambient illumination conditions. All reconstructions were computed using ten high-frequency gray code patterns.

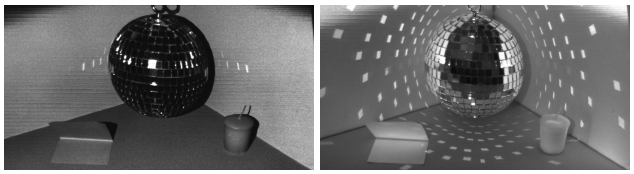

(a) epipolar imaging (b) non-epipolar imaging

Figure 14: *Separating components of illumination: we can capture live epipolar images in one shot, and much more efficiently than prior work. Note that most interreflections appear in the non-epipolar component: reflections from the ball appear in the epipolar image only when they agree with the epipolar geometry. These images were captured by exposing epipolar stripes wider than a pixel; as a result, some short-range subsurface scattering appears in the epipolar image instead of the non-epipolar image. Vignetting artifacts at the corners are due to projector distortion.*

It should be noted that our method is complementary to other methods for suppressing ambient light like placing a narrow-band spectral filter on the camera that matches the output wavelength of the active light source. Also note that, unlike the ambient light suppression done in time-of-flight cameras (which is electronic and subject to shot noise), epipolar structured light blocks ambient illumination before it reaches the sensor.

Analysis of range and power The working range of an active illumination system is limited by read noise and shot noise. The contribution of ambient light to the image is essentially independent of distance, while light from the active illumination source has an inverse square fall off. When the ambient light reaching the camera is small compared to the light from the active source, range is limited by read noise. As the relative contribution of ambient light to the image grows, the light from the source is lost in the shot noise of the ambient component.

Weakening the effect of ambient light by coding or optical filtering can increase the working range of an active illumination system at a given power level. Let k be the factor by which an imaging system can effectively weaken the influence of ambient light. Consider these four cases: (1) using no coding or optical filtering ($k = 1$); (2) using a narrow band light source and a filter on the camera ($k = 20$); (3) an idealized (distortion free) system that uses our epipolar codes with a 480 row projector in conjunction with a filter ($k = 480 \times 20 \approx 10,000$); and (4) a system like ours where the exposed pixel bands have been expanded to be 6 pixels thick to accommodate distortion ($k = 1,600$).

As shown in Figure 18, systems with a higher ability to block ambient light show significantly slower degradation in working range as ambient light levels increase.[2] To attain a working range of $3m$ with little or no blocking of ambient light (low k values) requires a very powerful light source, while systems like ours with a high k value are much more energy efficient.

[2]We define working range as the maximum distance at which a binary projected pattern can be reliably decoded (see the supplementary materials).

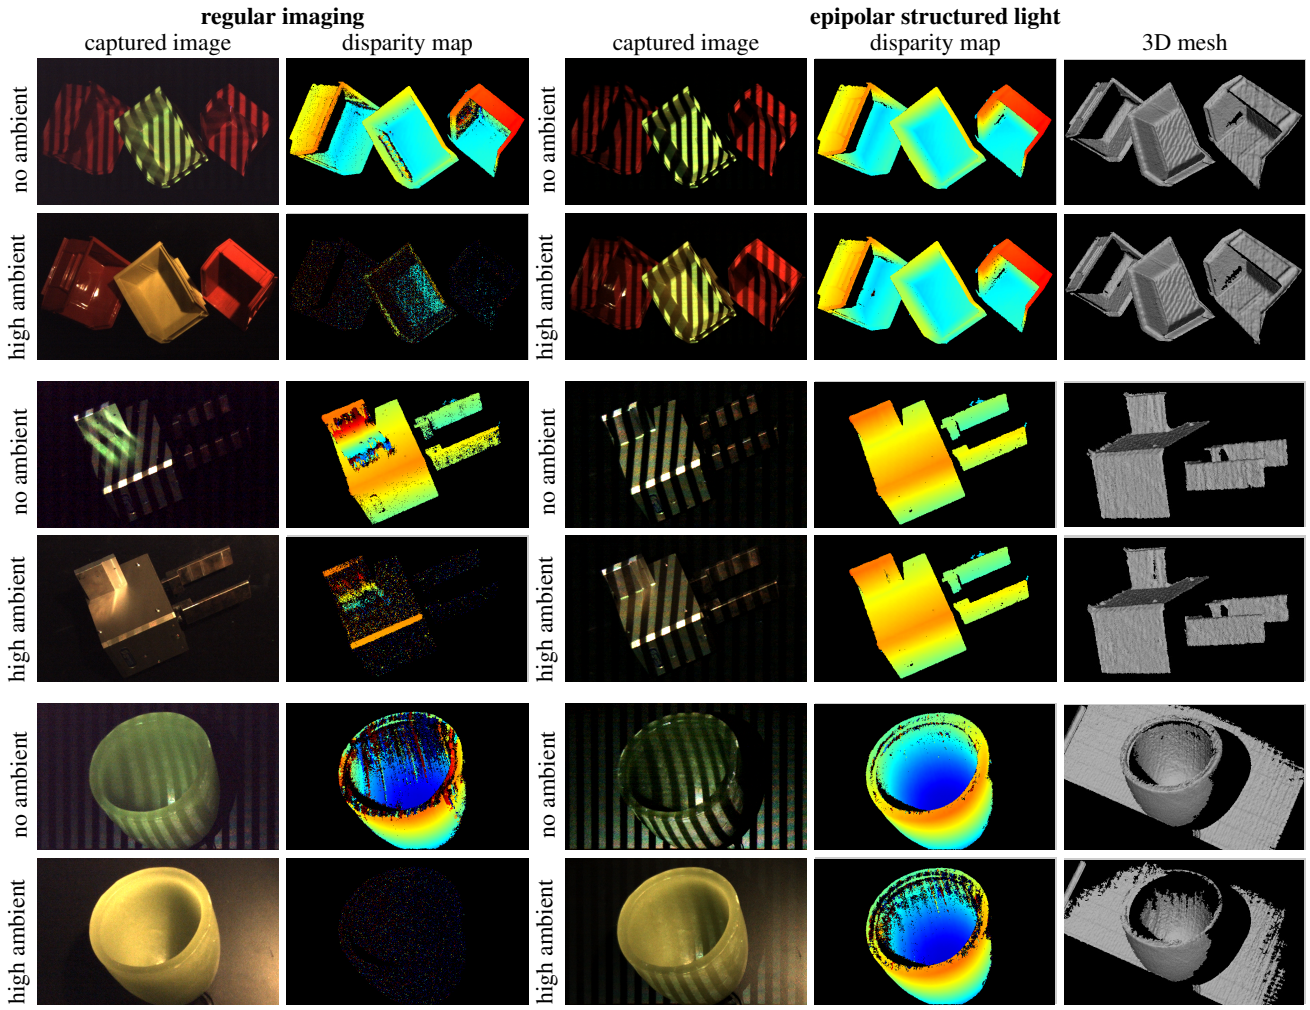

Figure 15: *Imaging and reconstructing challenging objects with strong global illumination effects under ambient lighting. Structured light with regular imaging is prone to errors due to global light transport effects (notice the effects of interreflections on the plastic bins and industrial gripper and subsurface scattering in the green wax bowl). Additionally, since active illumination sources are often orders of magnitude weaker than ambient light, structured light patterns are easily overwhelmed by ambient light. Epipolar structured light blocks a significant fraction of both. In these examples an ambient light level of 10 kilolux causes structured light with regular imaging to break down. In contrast, epipolar structured light still works well. We used the same acquisition time for both methods in each scene.*

5.3 Disparity Gating

When the camera and projector are rotated about their optical axis by 90° so that projector scanlines and camera rows are perpendicular to the epipolar planes, the correspondence between scanlines and rows becomes dependent on scene depth. By changing the trigger offset t_o we can attenuate energy transmission on the basis of scene depth. This corresponds to probing with the disparity gating matrix shown in Figure 5. By capturing a sequence of images where each image "unblocks" direct surface reflections from narrow range of disparity values, we can recover the shape of a scene. In Figure 20 we demonstrate how disparity gating could be used to recover depth through participating media.

5.4 Live Dual Videography

Dual photography [Sen et al. 2005; Koppal and Narasimhan 2015] is a technique that makes it possible for a projector and a camera in general position to "exchange" their viewpoints. For instance, one can synthesize a photo of exactly what the scene would look like from the projector's point of view even though there is no camera at that location. Creating such an image generally requires computing the full transport matrix between the projector and the camera; the technique in [Sen et al. 2005] required thousands of actual images of the scene to be captured, and hours of processing.

Here we use our laser-based prototype to demonstrate that while capturing the full transport matrix of a scene is still very hard, capturing an approximation of the *epipolar* image from the projector's viewpoint is not only easy, it can be done in real time with no computational processing whatsoever (Figure 21).

We place the projector and camera in a disparity-gating configuration (Section 5.3) and adjust timings so that the projector's scanline and the camera's "active" row are identical. This corresponds to disparity-gating of the plane at infinity and is described by the dual videography probing matrix in Figure 5. In this configuration, a textured plane at infinity would form the same image on both the projector plane and the sensor plane. Nearby objects, however, would form different images because of depth variations and because of indirect light transport. We then place a line diffuser in front of the

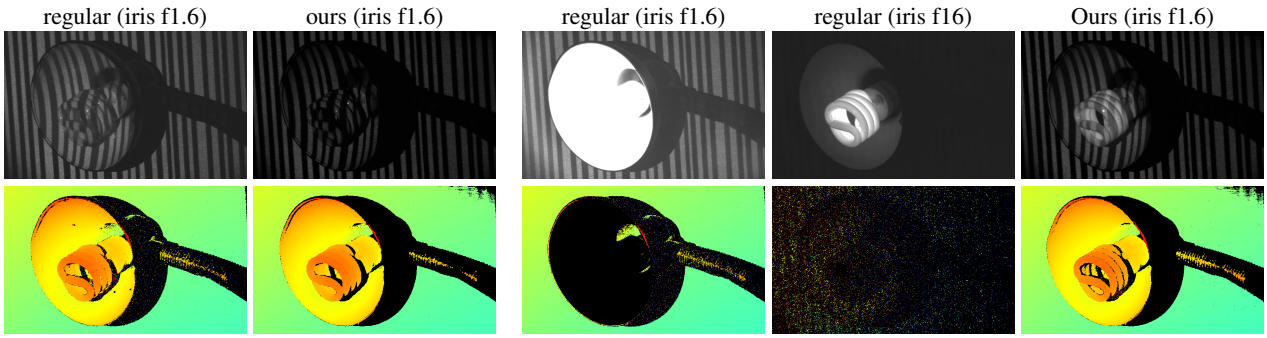

(a) lamp off (b) lamp on

Figure 16: *Imaging and scanning a 1600 Lumen lamp with a 5 Lumen projector: When the lamp is off (a), the pattern projected by the projector is visible with both regular imaging and our method. By projecting a series of structured light patterns the lamp can be reconstructued. When the bulb is turned on (b), regular imaging breaks down. With a large aperture the image is saturated, and with a small aperture the projected pattern is not visible to the camera on the bulb or the shade. Our rolling shutter-based implementation exposes sensor pixels only while they can receive light from the projector. As a result, most of the light from the bulb is blocked, the pattern is clearly visible even on the bulb, and the bulb can be reconstructed even though it is on.*

(a) regular (high f-number) (b) ours (low f-number) (c) disparity map (d) 3D mesh

Figure 17: *Active illumination with a 5 Lumen laser projector in bright sunlight (80 kilolux): With regular imaging (a), the active illumination patterns are overwhelmed by sunlight and are not visible despite using a wavelength filter. Our rolling shutter based implementation blocks a large fraction of the ambient light. This allows the projected pattern to be seen by the camera (b) and makes 3D structured light reconstruction possible (c,d).*

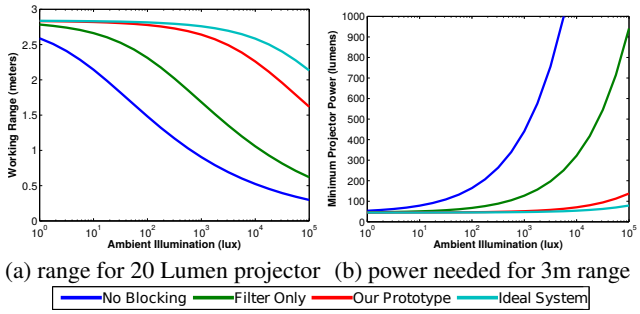

(a) range for 20 Lumen projector (b) power needed for 3m range

Figure 18: *Range analysis. (a) Maximum working range of a system with a 20 Lumen projector with a 45° horizontal field of view running at $60Hz$ as a function of ambient light level for different ambient blocking schemes (see supplementary material for complete list of simulation parameters). (b) The projector power required to attain a $3m$ range. By effectively blocking a large fraction of the ambient light, redistributive projectors with suitably-designed sensor masks can provide energy-efficient operation even under bright ambient illumination conditions.*

camera, with its axis of diffusion oriented horizontally, *i.e.*, along the epipolar planes. This smears along the epipolar planes the image of every scene point they contain, making it invariant to depth variations in the scene. As a result, each pixel in the camera's active column captures direct surface reflections from the scene point illuminated by the projector—no matter what its actual depth. As the projector scanline and active row sweep across the scene, an image is formed at the camera from the projector's viewpoint. Key to this method's success is again our prototype's energy efficiency: even though the diffuser scatters light a great deal, enough of it reaches the camera to create a clear image that is relatively free of noise.

6 Conclusion

We presented a novel framework for modeling and maximizing the energy efficiency of an imaging system under controlled masking and lighting. We introduced the notion of *redistribution ratio* that models the degree to which a light source's energy can be used for imaging rather than be wasted due to blocking/attenuation. We prove that optimal energy efficiency requires impulse illumination, where all the energy of the source can be concentrated in just one projector pixel that can be scanned rapidly. For projectors that cannot do this, we provide an algorithm that computes energy-efficient codes by solving a homogeneous matrix factorization problem through local optimization.

We also combine a laser projector and a rolling-shutter camera to demonstrate several novel capabilities, made possible only because our prototype can make near-optimal use of light source energy. That said, much remains to be understood on the problem of maximizing energy efficiency—especially for projectors that cannot use the energy of their source very efficiently (no closed form solution exists in that case). Last but not least, as dynamic vision sensors become more practical [Matsuda et al. 2015], and as redistributive laser systems and electronic masks become more flexible, it should be possible to build systems that push the energy efficiency frontier even further.

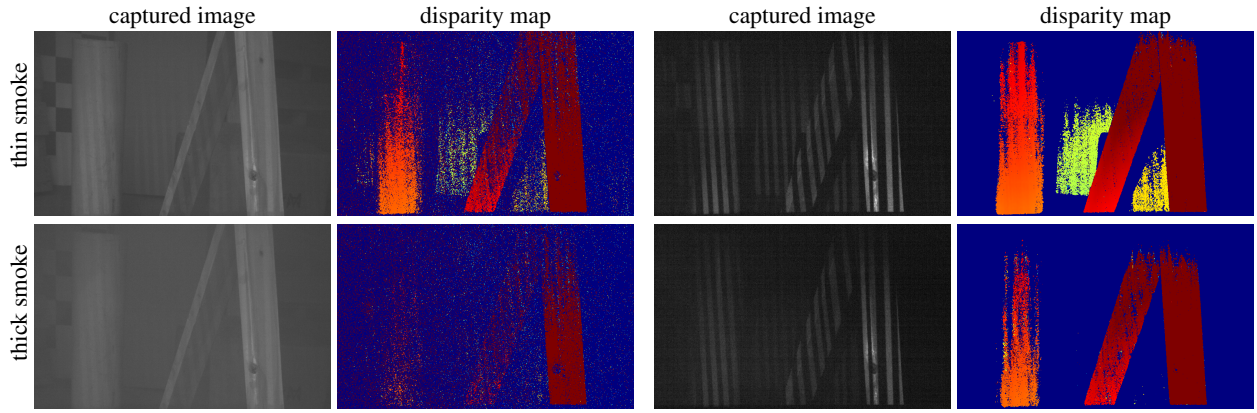

Figure 19: *Epipolar structured light scanning in participating media. We project structured light patterns and reconstruct shape in the presence of volumetric scattering due to smoke. With regular imaging (a), the smoke degrades the contrast of the projected patterns, reducing the range over which shape can be recovered. Epipolar structured light (b) blocks a lot of the scattered light, preserving image contrast and allowing a larger range of depths to be recovered at a given smoke concentration.*

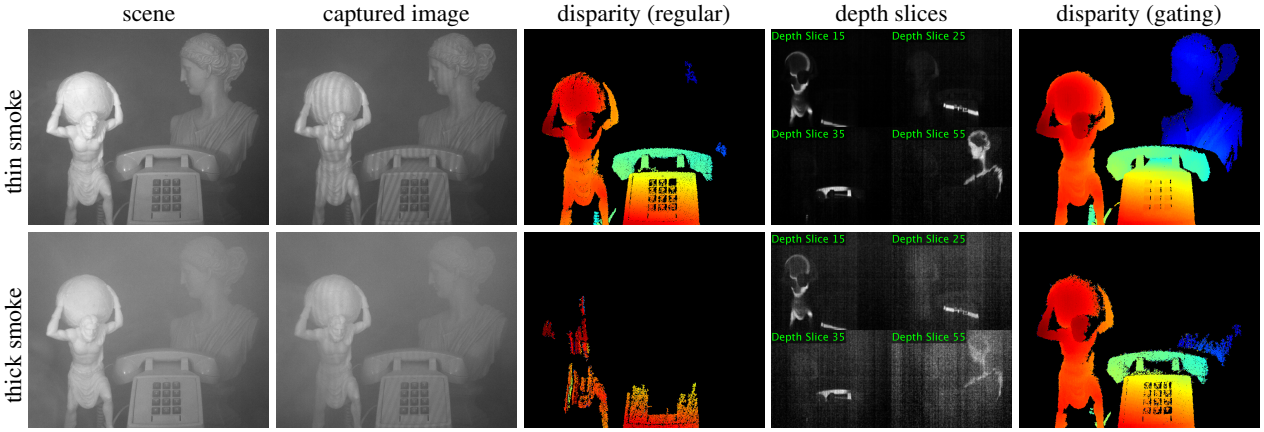

Figure 20: *Disparity gating in participating media. We compare structured light with regular imaging to disparity gating. Disparity gating is able to reconstruct objects at a further distance through the smoke than regular structured light. For structured light, we use 20 patterns and average three 16.67ms exposures for each pattern. For disparity gating we divide the range of disparities in the scene into 60 slices and capture each slice with a 16.67ms exposure.*

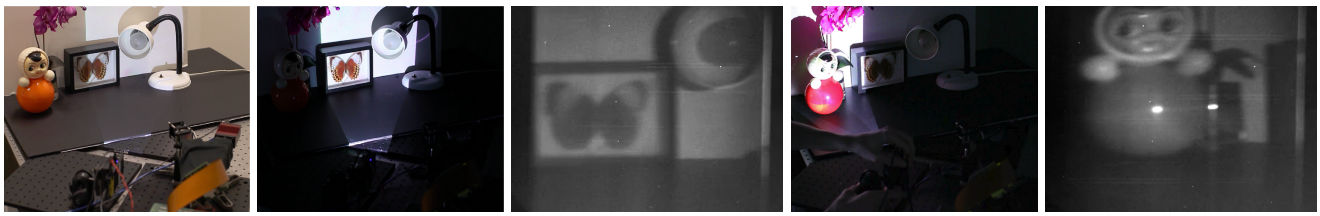

Figure 21: *Live dual videography. We demonstrate how combining an impulse projector with a masked camera can be used to capture live video of a scene from the point of view of the projector. As the projector pans from right (b) to left (d), the apparent viewpoint of the video frame captured by the static camera changes from (c) to (e).*

Acknowledgements

O'Toole and Kutulakos gratefully acknowledge the support of the QEII-GSST program, and the Natural Sciences and Engineering Research Council of Canada under the RGPIN, RTI, and SGP programs. Achar and Narasimhan received support in parts by US NSF Grants 1317749 and 0964562, and by the US Army Research Laboratory (ARL) Collaborative Technology Alliance Program, Cooperative Agreement W911NF-10-2-0016. The authors would like to thank Christoph Mertz and John Kozar for their help with the projector hardware and Chao Liu and Minh Vo for their help with experiments.

References

BACH, F., MAIRAL, J., AND PONCE, J. 2008. Convex Sparse Matrix Factorizations. *Preprint arXiv: 0812.1869*.

COSSAIRT, O., GUPTA, M., AND NAYAR, S. 2012. When Does Computational Imaging Improve Performance? *IEEE TIP 22*, 2, 447–458.

DAMBERG, G., AND HEIDRICH, W. 2015. Efficient freeform lens optimization for computational caustic displays. *Opt. Express 23*, 8 (Apr), 10224–10232.

DAMBERG, G., BALLESTAD, A., KOZAK, E., KUMARAN, R., AND MINOR, J. 2014. Efficient, High Brightness, High Dynamic Range Projection. In *ACM SIGGRAPH 2014 Emerging Technologies*, 18:1–18:1.

DECKER, JR, J. A., AND HARWIT, M. 1969. Experimental operation of a Hadamard spectrometer. *Appl. Opt 8*, 12, 2552–2554.

GU, J., NAYAR, S., GRINSPUN, E., BELHUMEUR, P. N., AND RAMAMOORTHI, R. 2013. Compressive Structured Light for Recovering Inhomogeneous Participating Media. *IEEE T-PAMI 35*, 3, 1–14.

GUPTA, M., YIN, Q., AND NAYAR, S. 2013. Structured Light in Sunlight. In *Proc. IEEE ICCV*, 545–552.

HAEFFELE, B., YOUNG, E., AND VIDAL, R. 2014. Structured Low-Rank Matrix Factorization: Optimality, Algorithm, and Applications to Image Processing. In *Proc. ICML*, 2007–2015.

HARWIT, M., AND SLOANE, N. J. A. 1979. Hadamard Transform Optics. In *Hadamard Transform Optics*. Academic Press.

HITOMI, Y., GU, J., GUPTA, M., MITSUNAGA, T., AND NAYAR, S. K. 2011. Video from a single coded exposure photograph using a learned over-complete dictionary. In *Proc. IEEE ICCV*, 287–294.

HOSKINSON, R., STOEBER, B., HEIDRICH, W., AND FELS, S. 2010. Light reallocation for high contrast projection using an analog micromirror array. *ACM SIGGRAPH Asia 29*, 6 (Dec.), 165:1–165:10.

IBBETT, R. N., ASPINALL, D., AND GRAINGER, J. F. 1968. Real-Time Multiplexing of Dispersed Spectra in Any Wavelength Region. *Appl Optics 7*, 6, 1089–1093.

KOPPAL, S., AND NARASIMHAN, S. 2015. Beyond perspective dual photography with illumination masks. *IEEE Transactions on Image Processing 24*, 7 (July), 2083–2097.

LANMAN, D., HIRSCH, M., KIM, Y., AND RASKAR, R. 2010. Content-adaptive parallax barriers: Optimizing dual-layer 3d displays using low-rank light field factorization. *ACM SIGGRAPH Asia 29*, 6 (Dec.), 163:1–163:10.

MATSUDA, N., COSSAIRT, O., AND GUPTA, M. 2015. MC3D: Motion Contrast 3D Scanning. In *Proc. IEEE ICCP*, 147–156.

MERTZ, C., KOPPAL, S. J., SIA, S., AND NARASIMHAN, S. G. 2012. A low-power structured light sensor for outdoor scene reconstruction and dominant material identification. In *Proc. IEEE PROCAMS*, 15–22.

MITRA, K., COSSAIRT, O., AND VEERARAGHAVAN, A. 2014. Can we beat Hadamard multiplexing? Data driven design and analysis for computational imaging systems. In *Proc. IEEE ICCP*, 1–9.

MITRA, K., COSSAIRT, O. S., AND VEERARAGHAVAN, A. 2014. A Framework for Analysis of Computational Imaging Systems: Role of Signal Prior, Sensor Noise and Multiplexing. *IEEE T-PAMI 36*, 10, 1909–1921.

MULLER, M., 2012. Confocal imaging device using spatially modulated illumination with electronic rolling shutter detection, Aug. 7. US Patent 8,237,835.

NAYAR, S. K., BRANZOI, V., AND BOULT, T. 2004. Programmable imaging using a digital micromirror array. In *Proc. CVPR*, 436–443.

O'TOOLE, M., RASKAR, R., AND KUTULAKOS, K. N. 2012. Primal-dual coding to probe light transport. *ACM SIGGRAPH 31*, 4 (July), 39:1–39:11.

O'TOOLE, M., MATHER, J., AND KUTULAKOS, K. N. 2014. 3D Shape and Indirect Appearance by Structured Light Transport. In *Proc. CVPR*, 3246–3253.

PARIKH, N., AND BOYD, S. 2014. Proximal Algorithms. *Foundations and Trends in Optimization 1*, 3.

RASKAR, R., AGRAWAL, A., AND TUMBLIN, J. 2006. Coded exposure photography: Motion deblurring using fluttered shutter. *ACM SIGGRAPH 25*, 3 (July), 795–804.

SCHECHNER, Y. Y., NAYAR, S. K., AND BELHUMEUR, P. N. 2007. Multiplexing for optimal lighting. *IEEE T-PAMI 29*, 8, 1339–1354.

SEN, P., CHEN, B., GARG, G., MARSCHNER, S. R., HOROWITZ, M., LEVOY, M., AND LENSCH, H. P. A. 2005. Dual photography. *ACM SIGGRAPH 24*, 3 (July), 745–755.

TAKHAR, D., LASKA, J. N., WAKIN, M. B., DUARTE, M. F., SARVOTHAM, D. B. S., KELLY, K. F., AND BARANIUK, R. G. 2006. A New Compressive Imaging Camera Architecture using Optical-Domain Compression. In *Proc. Computational Imaging IV*, 43–52.

VEERARAGHAVAN, A., REDDY, D., AND RASKAR, R. 2011. Coded strobing photography: compressive sensing of high speed periodic videos. *IEEE T-PAMI 33*, 4, 671–686.

WETZSTEIN, G., LANMAN, D., HIRSCH, M., AND RASKAR, R. 2012. Tensor displays: Compressive light field synthesis using multilayer displays with directional backlighting. *ACM SIGGRAPH 31*, 4 (July), 80:1–80:11.

XU, Y., AND YIN, W. 2013. A block coordinate descent method for regularized multiconvex optimization with applications to nonnegative tensor factorization and completion. *SIAM J. Imaging Sci. 6*, 3, 1758–1789.