

LiSens — A Scalable Architecture for Video Compressive Sensing

Jian Wang,[†] Mohit Gupta,[‡] and Aswin C. Sankaranarayanan[†]

[†] ECE Department, Carnegie Mellon University, Pittsburgh, PA

[‡] CS Department, Columbia University, New York, NY

Abstract

The measurement rate of cameras that take spatially multiplexed measurements by using spatial light modulators (SLM) is often limited by the switching speed of the SLMs. This is especially true for single-pixel cameras where the photodetector operates at a rate that is many orders-of-magnitude greater than the SLM. We study the factors that determine the measurement rate for such spatial multiplexing cameras (SMC) and show that increasing the number of pixels in the device improves the measurement rate, but there is an optimum number of pixels (typically, few thousands) beyond which the measurement rate does not increase. This motivates the design of LiSens, a novel imaging architecture, that replaces the photodetector in the single-pixel camera with a 1D linear array or a line-sensor. We illustrate the optical architecture underlying LiSens, build a prototype, and demonstrate results of a range of indoor and outdoor scenes. LiSens delivers on the promise of SMCs: imaging at a megapixel resolution, at video rate, using an inexpensive low-resolution sensor.

1. Introduction

Many computational cameras achieve novel capabilities by using spatial light modulators (SLMs). Examples include light field cameras [16, 24], lensless cameras [9], and compressive cameras [4, 7, 13, 20]. The measurement rate of these cameras is often limited by the switching speed of the SLM used. This is especially true for the compressive single-pixel camera (SPC) [4] where a photodetector is super-resolved by a digital micromirror device (DMD). The SPC inherits both the spatial resolution of the DMD as well as its operating rate. On one hand, this allows capturing images at a high resolution in spite of the sensor being a single pixel. This is especially beneficial while imaging in non-visible wavebands, where the combination of a single photodetector

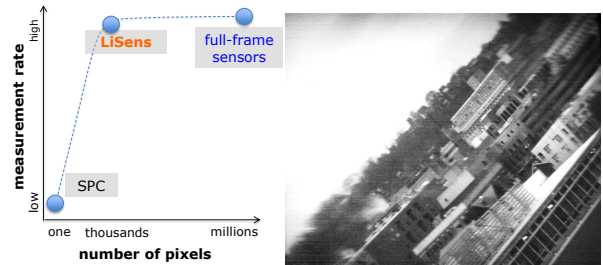


Figure 1. **Motivation.** Varying the resolution of the sensor in a multiplexing camera provides interesting tradeoffs. A single-pixel camera is inexpensive but has limited measurement rate. Full-frame sensors deliver high measurement rates but often at a steep price especially in non-visible wavebands. The proposed camera, LiSens, uses a 1024 pixel line-sensor and achieves a measurement rate in MHz. (right) A 1024×768 pixel image of an outdoor scene using our prototype.

coupled with a high-resolution DMD often provides an inexpensive alternative to full-frame sensors.¹ On the other hand, the SPC also inherits the operating rate of the DMD which, for commercially-available units, is in tens of kHz. At this measurement rate, acquiring images and videos at the spatial resolution of the DMD is feasible only at a very low temporal resolution.

It is instructional to compare the two widely differing models of sensing (see Figure 1) in terms of their cost, determined by the number of pixels in the camera, and measurement rate, defined as the number of measurements that the device obtains in unit-time. Full-frame 2D sensors that rely on Nyquist sampling are capable of achieving measurement rates in tens of MHz, but are prohibitively expensive in many non-visible wavebands. The SPC, which uses multiplexing as opposed to sampling, has low measurement rates, but is inexpensive for exotic wavebands. However, these imaging models are only the two extremes of a continuous camera design

¹Megapixel sensors in short-wave infrared, typically constructed using InGaAs, cost more than USD 100k.

space with varying number of pixels. Our main observation is that the design that achieves the maximum measurement rate at low cost lies between the two extremes.

Based on this observation, we propose a novel architecture for video compressive sensing (CS) that is capable of delivering measurement rates in MHz. The proposed camera, which we call *Line-SENSOR-based compressive camera (LiSens)*, uses a 1D array of pixels to observe a scene via a DMD. Each pixel on the sensor is optically mapped to a row of micromirrors on the DMD and hence, the camera obtains a coded line-integrals of the image formed on the DMD. The high measurement rate of LiSens enables CS of images and videos at high spatial resolutions and reasonable frame-rates. We provide the optical schematic for building such a camera and illustrate its novel capabilities with a lab prototype.

2. Prior work

Multiplexed imaging. This paper relies heavily on the concept of optical multiplexing. Suppose that we seek to sense a signal $\mathbf{x} \in \mathbb{R}^N$. A conventional camera samples this signal directly and has an image formation model given by

$$\mathbf{y} = \mathbf{x} + \mathbf{e},$$

where \mathbf{e} is the measurement noise. Multiplexing expands the scope of the camera by allowing it to obtain any arbitrary linear transformation of the signal:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{e}.$$

Assuming that $\mathbf{A} \in \mathbb{R}^{N \times N}$ is invertible, the signal can be estimated as $\hat{\mathbf{x}} = \mathbf{A}^{-1}\mathbf{y}$. It is well known that, when the entries of $\mathbf{A} = [a_{ij}]$ are restricted to $|a_{ij}| \leq 1$, using a Hadamard measurement matrix is optimal provided the noise is signal independent [6]. Note that in the absence of any priors, we need at least N measurements to recover an N -dimensional signal.

Compressive sensing (CS). CS aims to sense a signal $\mathbf{x} \in \mathbb{R}^N$ from an under-determined linear system, i.e., measurement of the form

$$\mathbf{y} = \Phi\mathbf{x} + \mathbf{e},$$

where $\Phi \in \mathbb{R}^{M \times N}$ with $M < N$. For an arbitrary signal in \mathbb{R}^N this is impossible since the map $\Phi : \mathbb{R}^N \mapsto \mathbb{R}^M$ is many-to-one and non-invertible. CS handles this by restricting the signal \mathbf{x} to belong to a distinguished class; for example, sparse signals in a transform basis. The main results of CS state that when the measurement matrix Φ has a special structure and \mathbf{x} is K -sparse in a transform basis, then we can robustly recover \mathbf{x} , provided $M = O(K \log(N/K))$ [1].

Video models for CS. There have been many models proposed to handle CS of videos. Wakin et al. [26] proposed the use of 3D wavelets as a sparsifying transform for videos. This model is further refined in Park and Wakin [18] where a lifting scheme is used to tune a wavelet to better approximate the temporal variations. Dictionary-based models were used in Hitomi et al. [7] for temporal resolution of videos; it is shown that a highly overcomplete dictionary provides high quality reconstructions. Inspired by the use of motion-flow models in video compression, Reddy et al. [20] and subsequently, Sankaranarayanan et al. [21] employed an iterative strategy where optical flow derived from an initial estimate is used to further constrain the video recovery problem. This provides a significant improvement in reconstruction quality, albeit at a high computational cost. Gaussian mixture models were used in [27] for video CS; a hallmark of this approach is that the mixture model is learned directly from the compressive measurements providing a model that is tailored to the specifics of the scene being reconstructed.

Compressive imaging hardware prototypes. The original prototype for the SPC used a DMD as the programmable SLM [4]. In [9], a variant of the SPC is proposed where an LCD panel is used in place of the DMD; the use of a transmissive light modulator enables a lensless architecture. Sen and Darabi [23] use a camera-projector system to construct an SPC exploiting a concept called dual photography [22]; the hallmark of this system is its use of active illumination.

There have been many multi-pixel extensions to the SPC. The simplest approach [11] is to map the DMD to a low-resolution sensor array, as opposed to a single photodetector, such that each pixel on the sensor observes a non-overlapping “patch” or a block of micromirrors on the DMD.² SMCs based on this design have been proposed for sensing in mid-wave infrared camera [15] and short-wave infrared [3]. Measurement matrix designs for such block-based compressive imaging architectures are presented in Kerviche et al. [12] and Ke and Lam [10].

While we focus on spatial multiplexing, many architectures have been proposed for multiplexing other attributes of a scene including temporal [5, 7, 8, 14, 20], angular [16, 24], and spectral [13, 25]. All of these architectures use a high-resolution sensor and sacrifice this spatial resolution partially to obtain higher resolution in

²The interested reader is referred to earlier work by Nayar et al. [17] where a scene is observed by a sensor via a DMD to enable programmable imaging. However, this device was not used for spatial multiplexing or compressive sensing.

time, spectrum and/or angle. In contrast, the imaging architecture proposed in this paper takes a device with a high temporal resolution and sacrifices the temporal resolution in part to obtain higher spatial resolution.

3. Measurement rate of a camera

The measurement rate of a camera is given by the product of its spatial resolution, in pixels, and the temporal resolution, in frames per second. A spatial multiplexing camera (SMC) captures one or more coded linear measurements of a scene via a spatial light modulator (e.g., DMD). Multiple measurements are taken sequentially by changing the code displayed on the DMD. Thus, the measurement rate of an SMC is limited by the resolution and the frame-rate of the sensor, as well as the resolution and switching speed of the DMD.³ The operating rate of the DMD is determined by the rate at which the micromirrors can be switched from one code into another. Denoting this rate as R_{DMD} Hz, we note that the frame rate of the SMC cannot be greater than R_{DMD} since the specific linear mapping from the scene to sensor is determined by the micromirror code. The frame rate of a sensor is also limited by its readout speed — typically, determined by the operating rate of the analog-to-digital converter (ADC) in the readout circuit. Suppose that we have an ADC that can perform readout at a rate of R_{ADC} Hz. Then, given a frame with F pixels, the ADC limits the sensor to R_{ADC}/F frames per second. Hence, the measurement rate of the SMC is given as

$$F \times \min\left(R_{DMD}, \frac{R_{ADC}}{F}\right) \text{ meas. per sec.}$$

Operating speeds of commercially-available DMDs are 2-3 orders of magnitude lower than that of ADCs in sensors. As a consequence, for small values of F , the measurement rate of an SMC is dominated by the operating speed of the DMD (see Figure 2). Once $F \geq R_{ADC}/R_{DMD}$, then the bottleneck shifts from the DMD to the ADC. Further, the smallest sensor resolution (in terms of number of pixels) for which the measurement rate is maximized is

$$(\text{minimum number of pixels}) \quad F_{\min} = R_{ADC}/R_{DMD}.$$

In essence, at $F = F_{\min}$ we can obtain the measurement rate of a full-frame sensor but with a device with potentially a fraction of the number of pixels. This is

³Compressive cameras can further increase the spatial or temporal resolution by using scene priors.




Figure 2. **Measurement rate and sensor resolution.** There are two regimes of operations: for $F < F_{\min}$, the bottleneck is the DMD and for $F > F_{\min}$, the bottleneck is the sensor readout. At $F = F_{\min}$ lies a design that obtains the highest measurement rate with the least sensor resolution.

invaluable for sensing in many wavebands, for example short-wave infrared.

As a case study, consider an SMC with a DMD operating rate $R_{DMD} = 10^4$ Hz and an ADC with an operating rate $R_{ADC} = 10^7$ Hz. Then, for a sensor with $F_{\min} = 10^3$ pixels, we can obtain 10^7 measurements per second. An SPC, in comparison, would only provide 10^4 measurements per second. This significant increase in measurement rate motivates the multi-pixel SMC design that we propose in the next section.

4. LiSens

The optical setup of LiSens is illustrated in Figure 3. The DMD splits the optical axis into two axes — one each for $\pm 11^\circ$ orientations of the micromirrors. Along one of these axes, the 2D image of the scene formed on the DMD plane is mapped onto the 1D line-sensor using a relay lens and a cylindrical lens. The image captured by the line-sensor can be represented as a 1D integral of the 2D image (along rows or columns) formed on the DMD plane. This is achieved by aligning the axis of the cylindrical lens with that of the line-sensor and placing it so that it optically mirrors the aperture plane of the relay lens onto the sensor plane (see Figure 4).

Suppose that the physical dimensions of the DMD and the line-sensor are (w_D, h_D) and (w_L, h_L) , respectively, with $h_L \ll w_L$. The relay lens is selected to produce a magnification of w_L/w_D so that the DMD maps to the line-sensor along its width; this determines the ratio of the focal lengths used in the relay lens. The exact value of the focal lengths is optimized to minimize




Figure 3. **Schematic of the LiSens imaging architecture.**




Figure 4. **The DMD-to-sensor mapping.** Shown are the ray diagrams for the mapping of the DMD to line-sensor along two orthogonal directions — (top) along and (bottom) orthogonal to the axis of cylindrical lens. The cylindrical lens focuses the aperture plane of the relay lens onto the sensor plane. Hence, the sensor plane is focused on the DMD along one-axis (top) but defocused along the other (bottom).

vignetting and needs to be determined along with the objective lens. If $w_D = w_L$, we can use a 1:1 relay lens with a focal length f_r and the distance between the relay lens and the line-sensor is $2f_r$. The parameters and the placement of the cylindrical lens are determined so that the aperture of the relay lens is magnified (or shrunk) within the line-sensor. This corresponds to a magnification of h_L/d_r , where d_r is the diameter of the relay lens. Let f_c be the focal length of the cylindrical lens, and suppose that it is placed at a distance of u_c from the line-sensor and v_c from the relay lens. Then,

$$u_c + v_c = 2f_r, \quad \frac{1}{u_c} + \frac{1}{v_c} = \frac{1}{f_c}, \quad \frac{u_c}{v_c} = \frac{h_L}{d_r}$$

Hence,

$$f_c = 2f_r \frac{h_L}{d_r + h_L} \frac{d_r}{d_r + h_L} \approx 2f_r \frac{h_L}{d_r}. \quad (1)$$

In practice, we observe that, for a marginal loss of light




Figure 5. **Field-of-view enhancement by the cylindrical lens.** Shown are images acquired with our prototype.




Figure 6. **Benefits of frame transfer.** The DMD was operated with $400\mu s$ per pattern. The sensor had a readout of $350\mu s$ per frame. (a) Simultaneous readout allows for a $400\mu s$ exposure. (b) Sequential readout cut the exposure time down to $50\mu s$, resulting in a noisy image. Shown are images acquired with our prototype.

throughput, we could obtain flexibility in both the choice of the cylindrical lens (focal length) as well as its positioning. Figure 5 shows the improvement in the vertical field-of-view when the cylindrical lens is introduced.

4.1. Practical advantages

Before we proceed, it is worth pondering on alternative and potentially simpler designs for multi-pixel SMCs. An intuitive multi-pixel extension of the SPC would be using a low-resolution 2D sensor array [3, 15]. Having a 2D sensor array provides a simpler mapping from the DMD that can be achieved using just relay lenses. Yet, there are important considerations that make the LiSens camera a powerful alternative to using 2D sensors.

Specifically, it is simpler and inexpensive to obtain frame transfer⁴ on a line sensor without requiring complex circuitry or loss of light due to reduced fill-factor.

⁴Frame transfer is a technology used to enable simultaneous exposure and readout in a sensor. In a traditional camera, during the readout process, the sensor is not exposed; this reduces the duty cycle of the sensor and leads to inefficiencies in light collection. With frame transfer, there is a separate array of storage pixels beside the photosensitive array. After exposure, charge from the pixels are immediately transferred to the storage array allowing the sensor to be exposed again.

Figure 6 demonstrates the benefits of frame transfer on our lab prototype. The 1D profile of the sensor also provides the possibility of having a per-pixel ADC; this provides a dramatic increase in the readout rate of the sensor. Finally, we benefit from the fact that line-sensor have long been manufactured for spectroscopy which requires very precise, low-noise sensors with high dynamic range and broad spectral response; these properties make them highly desirable for our application.

4.2. Imaging model

Suppose the DMD has a resolution of $N \times N$ micromirrors, and the line-sensor has N pixels such that each pixel maps to a row of micromirrors on the DMD. At time instant t , let $\mathbf{X}_t \in \mathbb{R}^{N \times N}$ be the scene image formed on the DMD, and let $\Phi_t \in \mathbb{R}^{N \times N}$ be the binary pattern displayed on the DMD. Then, the measurement obtained at the line-sensor, $\mathbf{y}_t \in \mathbb{R}^N$, is given as

$$\mathbf{y}_t = (\mathbf{X}_t \circ \Phi_t) \mathbf{1} + \mathbf{e}_t,$$

where \circ denotes the entry-wise product, $\mathbf{1}$ is the vector of ones and \mathbf{e}_t is the measurement noise.

For simplicity, we use DMD patterns where every column is the same, i.e., patterns of the form

$$\Phi_t = \mathbf{1} \phi_t^T,$$

where $\phi_t \in \mathbb{R}^N$ is a binary vector. This alleviates the need for extensive calibration. The sensor measurements are then given as

$$\mathbf{y}_t = \mathbf{X}_t \phi_t + \mathbf{e}_t.$$

For the experiments in the paper, we use rows of a column-permuted Hadamard matrix for ϕ_t .

4.3. Recovery

Sensing images. When the scene is static, i.e., $\mathbf{X}_t = \mathbf{X}$ for $t \in \{1, \dots, T\}$, we can write the imaging model as

$$\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_T] = \mathbf{X} \Phi + E,$$

where $\Phi = [\phi_1, \dots, \phi_T]$ is the measurement matrix and $E = [\mathbf{e}_1, \dots, \mathbf{e}_T]$ is the measurement noise.

If $T \geq N$ and Φ were well-conditioned, say for example, a Hadamard matrix, then we can obtain an estimate of the scene image as

$$\hat{\mathbf{X}} = \mathbf{Y} \Phi^\dagger. \quad (2)$$

In practice, it is unreasonable to expect a scene to be static over a large duration and hence, we can expect

$T < N$. To regularize the inverse problem, we enforce sparsity in the gradients of the recovered image using a minimum total-variation prior [2]. This leads to the following optimization problem.

$$\min_{\mathbf{X}} TV(\mathbf{X}), \quad \text{s.t.} \quad \|\mathbf{Y} - \mathbf{X} \Phi\|_F \leq \epsilon, \quad (3)$$

where $TV(\mathbf{X})$ is the total-variation norm that captures strength of the image gradients. We used the isotropic TV-norm defined as

$$TV(\mathbf{X}) = \sum_{i=1}^N \sum_{j=1}^N \sqrt{G_x^2(i, j) + G_y^2(i, j)},$$

where $G_x, G_y \in \mathbb{R}^{N \times N}$ are the spatial gradients of \mathbf{X} .

Sensing videos. We enforce sparse spatio-temporal gradients when recovering time-varying scenes. Specifically, suppose that we obtain measurements $\{\mathbf{y}_t, 1 \leq t \leq T\}$ for a time-varying scene $\{\mathbf{X}_t, 1 \leq t \leq T\}$. It is often an overkill to recover an image for every frame of the sensor since the problem becomes computationally overwhelming. Instead, we decide on a target frame-rate (say 10 frames per second) and group together successive measurements so as to obtain the desired frame rate. Let $\{\tilde{\mathbf{X}}_k, k = 1, \dots, Q\}$ be the frames associated with the scene at the desired frame-rate. Then, the imaging model reduces to

$$t = k \frac{T}{Q}, \quad [\mathbf{y}_t, \dots, \mathbf{y}_{t+\frac{T}{Q}}] = \tilde{\mathbf{X}}_k [\phi_t, \dots, \phi_{t+\frac{T}{Q}}] + E_k$$

Grouping together measurements and associating them to a frame of a video also reduces the number of parameters to be estimated. Interested readers are referred to [19] for a frequency domain analysis on judicious selection of frame-rate for the SPC. Finally, similar to (3), we solve for the video by minimizing a spatio-temporal TV-norm constrained by the measurements.

5. Hardware prototype

Our proof-of-concept prototype was built using a Nikkor 50mm F/1.8 objective lens, a DLP7000 DMD, and a Hamamatsu S11156-2048-01 line-sensor. The DMD had a spatial resolution of 1024×768 with a micromirror pitch of $13.8 \mu\text{m}$ and a maximum operating speed of 20 kHz. The line sensor had 2048 pixels with a pixel size of $14 \mu\text{m} \times 1 \text{mm}$; we were only able to use 1024 pixels due to our inability to find a cylindrical lens that was sufficiently long to span the entire line-sensor. Due to the pixel pitch and micromirror pitch being nearly the same, we used a 1:1 relay lens and oriented the line sensor so that the line sensor was aligned

to the width of the DMD. Hence, each pixel on the line sensor summed up 768 micromirrors on the DMD.

Our line sensor provided access to simultaneous exposure and readout using frame transfer. We operated the prototype with an exposure of approximately $500\mu s$. Due to buffer limitations in the readout circuit (a Hamamatsu C11165-01 driver), we were able to obtain 100 frames at $500\mu s$ per frame, followed by a cool-down time of $60ms$. As a consequence, the device provided 900 frames per second and hence, a measurement rate of 9×10^5 measurements per second.

For comparison and validation, we built an SPC on the off-axis of the DMD (see Figure 7). The measurement rate of the SPC was 20 kHz, the speed of the DMD.

Figure 7 shows both the modified schematic of the camera as well as the optical layout. The subtle differences to the schematic shown in Figure 3 stem from a couple of practical constraints. Recall that the DMD reflects the optical axis by 22° . The flange distance of objective lenses (F-mount and C-mount) is insufficient to provide clearance for the cone of light reflected from the DMD. To alleviate this, we optically mirror the image plane of the objective lens using a 1:1 relay lens, thereby providing ample space for the reflected cone of light. We used a 100mm:100mm achromatic doublets for the relay lens, thereby providing an $f_r = 50mm$ and a cylindrical lens with a focal length of $12.7mm$.⁵

Alignment. Misalignment of the optical components, in addition to introducing blur, results in a DMD column mapping to multiple pixels on the line detector. In our prototype, we minimized both effects by manually adjusting each component while observing and quantifying sharpness of the images in real-time. While this produces acceptable results, a calibration procedure is indeed required especially to resolve the image especially at the boundaries of the field of view.

Vignetting. The use of relay lenses to extend the optical axis causes significant vignetting. To alleviate this, we introduced a field lens at the DMD which dramatically reduces vignetting. A second field lens at the image plane of the objective lens further reduces vignetting but with an increase in spatial blur. The results in this paper were produced with a single field lens at the DMD.

6. Experiments

We use two metrics to characterize the operating scenarios for our experiments:

⁵We used relay lenses with diameter $d_r = 25mm$. The value suggested in (1) is $f_c = 4mm$. However, to avoid custom designed optics, we went with a larger focal length.

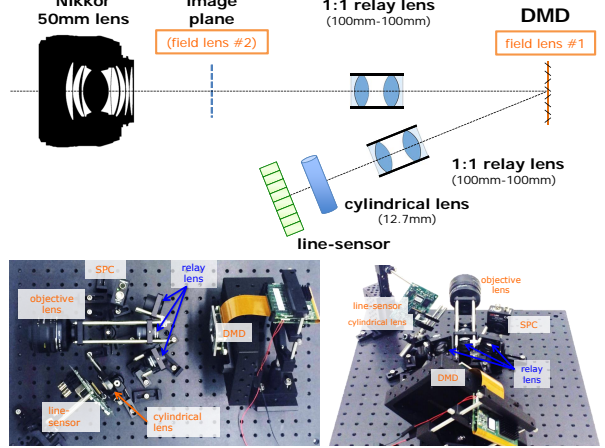


Figure 7. **Prototype of the LiSens camera.**

(i) *The under-sampling* given as the ratio of the dimensionality of the acquired image to the number of measurements, i.e., if we seek to acquire an $N_1 \times N_2$ image with M measurements, then the under-sampling is $N_1 N_2 / M$. When the under-sampling is 1, the system is invertible and we use (2) to recover the image. For values greater than 1, we use the TV prior and solve (3).

(ii) *Temporal resolution / capture duration per frame.* Recall, from Section 4.3, that we pool together measurements and associate them with a single recovered frame. This determines the capture duration per frame and its reciprocal, the temporal resolution. The smaller the capture duration, the fewer the measurements that we acquire and hence, the greater the under-sampling.

These two metrics are linked by the measurement rate of the camera and the resolution at which we operate the DMD. For a capture duration of τ seconds, the under-sampling is given as

$$(N_1 N_2) / (\tau \times \text{measurement rate})$$

where the DMD resolution is $N_1 \times N_2$. Recall, that the measurement rate of our LiSens prototype is nearly 1 MHz while that of the SPC is 20 kHz.

Finally, we map the ± 1 entries of the Hadamard matrix to 0/1 by mapping the '-1's to '0's; in post-processing, the measurements corresponding to the ± 1 matrix are obtained simply by subtracting the mean scene intensity. For time-varying scenes, we repeat the all-ones measurement in the Hadamard matrix once every 100 measurements to track the mean intensity value.

Gallery. We present the images recovered by the LiSens camera on two indoor static scenes for various capture durations in Figure 8. We also quantize the loss




Figure 8. **LiSens reconstructions for a few scenes under varying capture durations / under-sampling.** The reconstructions suffer little loss in quality until a capture time of $110ms$; this corresponds to an under-sampling of $8\times$. Inset on each figure is the reconstruction SNR, in dB, using the $1\times$ reconstruction as the ground truth.

in performance with increased under-sampling using the reconstruction SNR. Specifically, given a ground truth image \mathbf{x}_{gr} and a reconstructed image $\hat{\mathbf{x}}$, the reconstruction SNR in dB is given as

$$-20 \log_{10} \left(\frac{\|\mathbf{x}_{gr} - \hat{\mathbf{x}}\|_2}{\|\mathbf{x}_{gr}\|_2} \right).$$

We used the Nyquist rate reconstruction, i.e., under-sampling of $1\times$, as the ground truth image. We observe that there is little loss in performance until a capture duration of $220ms$ (under-sampling of $4\times$) and with a small drop in performance at $110ms$ ($8\times$). As we will see later, using a spatio-temporal prior improves the reconstruction quality even at an under-sampling of $8\times$.

Comparison to SPC. We compare the performance of LiSens and an SPC built on the off-axis of the DMD. We image a static scene using both cameras and compare the reconstructed image for varying capture durations (see Figure 9). Recall that the measurement rate of the SPC is limited to 20 kHz while the LiSens has nearly 1 MHz. This $50\times$ gap enables the LiSens camera to sense images at a resolution of 1024×768 at 5 – 10 frames per second. As we will see next, using simple spatio-temporal priors like a 3D TV norm will enable higher-quality reconstructions at an increased temporal resolution. We note that the photodetector in the SPC had a lower dynamic range than the line-sensor. For a fairer comparison, we took multiple runs with the SPC and averaged the measurements to artificially boost its dynamic range.

Figure 9 also indicates that LiSens operating at $8\times$ under-sampling produces significantly higher quality results than the SPC operating at $7.4\times$ under-sampling. This is a consequence of the sparsity of images scaling sub-linearly with image resolution. That is, if we looked at sparsity of images in a wavelet basis as a function of image resolution, the ratio of sparsity level to image resolution decreases monotonically. This can also be attributed to images having a lot of energy in low-freq components. Hence, if we attempt to recover a scene at 128×128 and 1024×1024 , each at say $8\times$ under-sampling, then the quality would be significantly better for the high-resolution reconstruction.

Dynamic scenes. Next, we sensed a time-varying scene — a bunny on a turntable spinning at one revolution every 15 seconds. We fix a capture duration of $110ms$ per frame or equivalently a temporal-resolution of 9 frames per second, and recover the frames of the video under a 3D total variation prior. Due to high levels of measurement noise, the recovered videos were subject to a $3 \times 3 \times 3$ median-filter. As seen in Figure 10, the recovered video preserves fine spatial detail. This can be attributed to the use of video priors, which capture inter-frame redundancies, as well as the high measurement rate enabled to our design.

This validates the potential of the LiSens architecture to acquire scenes at high spatial resolution (1024×768) and with a reasonable temporal resolution.

Finally, Figure 11 shows frames of a video of an outdoor scene reconstructed with the same parameters as

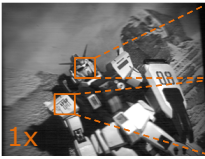



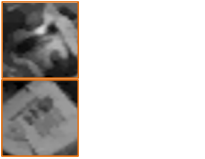



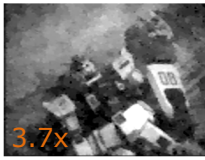

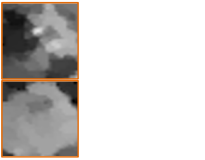


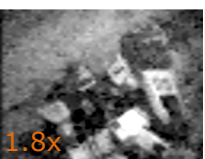
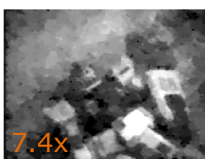
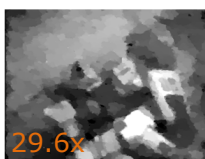


	LiSenS		single pixel camera					
	768x1024		64x64		128x128		256x256	
capture duration	0.88s							
	0.22s							
	0.11s							

Figure 9. **Performance of LiSenS and SPC.** We look at recovered images for various capture durations and, in the case of the SPC, various spatial resolution as well. The remarkable gap in the performance of the devices can be attributed to the difference in their measurement rates. Inset in orange are the under-sampling factors at which images were reconstructed.




Figure 10. **Video CS using LiSenS.** We image a bunny on a turntable (left) and recover the video at a capture duration of 110ms per frame (under-sampling of 8 \times) using a 3D TV prior. Shown are a few recovered frames and spatio-temporal slices.

the indoor bunny scene. The data was significantly noisier which contributed to the reduced reconstruction quality. The use of stronger priors such as motion flow models as in [20, 21] could help in such scenarios.

7. Discussion

We present a multiplexing camera that is capable of delivering very high measurement rates, comparable to that of a full-frame sensor, but with a number of pixels that is a small fraction. At its core, the proposed camera moves the bottleneck in the frame rate of the device from the spatial light modulator to the readout of the sensor.

The ensuing boost in measurement rate enables sensing of scenes at video rate and at the full resolution of the light modulator. Finally, while our prototype was built for visible wavebands, the underlying principles transfer *mutatis mutandis* to other wavebands of interest.

8. Acknowledgement

We thank Zhuo Hui and Yang Gao for the help with the data collection. J. W. and A. C. S. were supported in part by the NSF grant CCF-1117939.



Figure 11. Sensing outdoors with the LiSens prototype.

References

- [1] R. G. Baraniuk. Compressive sensing. *IEEE Signal Process. mag.*, 24(4), 2007.
- [2] A. Chambolle. An algorithm for total variation minimization and applications. *J. Mathematical Imaging and Vision*, 20(1-2):89–97, 2004.
- [3] H. Chen, S. Asif, A. C. Sankaranarayanan, and A. Veeraraghavan. FPA-CS: Focal plane array-based compressive imaging in short-wave infrared. In *CVPR*, 2015.
- [4] M. F. Duarte, M. A. Davenport, D. Takhar, J. N. Laska, T. Sun, K. F. Kelly, and R. G. Baraniuk. Single-pixel imaging via compressive sampling. *IEEE Signal Process. Mag.*, 25(2):83–91, Mar. 2008.
- [5] L. Gao, J. Liang, C. Li, and L. V. Wang. Single-shot compressed ultrafast photography at one hundred billion frames per second. *Nature*, 516(7529):74–77, 2014.
- [6] M. Harwit and N. J. Sloane. *Hadamard transform optics*. New York: Academic Press, 1979.
- [7] Y. Hitomi, J. Gu, M. Gupta, T. Mitsunaga, and S. K. Nayar. Video from a single coded exposure photograph using a learned over-complete dictionary. In *ICCV*, 2011.
- [8] J. Holloway, A. C. Sankaranarayanan, A. Veeraraghavan, and S. Tambe. Flutter shutter video camera for compressive sensing of videos. In *ICCP*, 2012.
- [9] G. Huang, H. Jiang, K. Matthews, and P. Wilford. Lensless imaging by compressive sensing. In *ICIP*, 2013.
- [10] J. Ke and E. Y. Lam. Object reconstruction in block-based compressive imaging. *Optics express*, 20(20):22102–22117, 2012.
- [11] K. Kelly, R. Baraniuk, L. McMackin, R. Bridge, S. Chatterjee, and T. Weston. Decreasing image acquisition time for compressive imaging devices, Oct. 14 2014. US Patent 8,860,835.
- [12] R. Kerviche, N. Zhu, and A. Ashok. Information-optimal scalable compressive imaging system. In *Computational Optical Sensing and Imaging*, 2014.
- [13] X. Lin, G. Wetzstein, Y. Liu, and Q. Dai. Dual-coded compressive hyperspectral imaging. *Optics letters*, 39(7):2044–2047, 2014.
- [14] P. Llull, X. Liao, X. Yuan, J. Yang, D. Kittle, L. Carin, G. Sapiro, and D. Brady. Compressive sensing for video using a passive coding element. In *Computational Optical Sensing and Imaging*, 2013.
- [15] A. Mahalanobis, R. Shilling, R. Murphy, and R. Muise. Recent results of medium wave infrared compressive sensing. *App. optics*, 53(34):8060–8070, 2014.
- [16] K. Marwah, G. Wetzstein, Y. Bando, and R. Raskar. Compressive light field photography using overcomplete dictionaries and optimized projections. *ACM Trans. Graphics (TOG)*, 32(4):46, 2013.
- [17] S. K. Nayar, V. Branzoi, and T. E. Boulton. Programmable imaging using a digital micromirror array. In *CVPR*, 2004.
- [18] J. Y. Park and M. B. Wakin. A multiscale framework for compressive sensing of video. In *Pict. Coding Symp.*, 2009.
- [19] J. Y. Park and M. B. Wakin. Multiscale algorithm for reconstructing videos from streaming compressive measurements. *J. Electronic Imaging*, 22(2), 2013.
- [20] D. Reddy, A. Veeraraghavan, and R. Chellappa. P2C2: Programmable pixel compressive camera for high speed imaging. In *CVPR*, June 2011.
- [21] A. C. Sankaranarayanan, C. Studer, and R. G. Baraniuk. CS-MUVI: Video compressive sensing for spatial-multiplexing cameras. In *ICCP*, 2012.
- [22] P. Sen, B. Chen, G. Garg, S. R. Marschner, M. Horowitz, M. Levoy, and H. Lensch. Dual photography. *ACM Trans. Graphics (TOG)*, 24(3):745–755, 2005.
- [23] P. Sen and S. Darabi. Compressive dual photography. *Computer Graphics Forum*, 28(2):609–618, 2009.
- [24] S. Tambe, A. Veeraraghavan, and A. Agrawal. Towards motion aware light field video for dynamic scenes. In *ICCV*, 2013.
- [25] A. Wagadarikar, R. John, R. Willett, and D. Brady. Single disperser design for coded aperture snapshot spectral imaging. *App. Optics*, 47(10):44–51, 2008.
- [26] M. B. Wakin, J. N. Laska, M. F. Duarte, D. Baron, S. Sarvotham, D. Takhar, K. F. Kelly, and R. G. Baraniuk. Compressive imaging for video representation and coding. In *Pict. Coding Symp.*, 2006.
- [27] J. Yang, X. Yuan, X. Liao, P. Llull, D. J. Brady, G. Sapiro, and L. Carin. Video compressive sensing using gaussian mixture models. *IEEE Trans. Image Proc.*, 23(11):4863–4878, 2014.