Optical Engineering

OpticalEngineering.SPIEDigitalLibrary.org

High-resolution imaging using a translating coded aperture

Abhijit Mahalanobis
Richard Shilling
Robert Muise
Mark Neifeld

SPIE.

Abhijit Mahalanobis, Richard Shilling, Robert Muise, Mark Neifeld, "High-resolution imaging using a translating coded aperture," *Opt. Eng.* **56**(8), 084106 (2017), doi: 10.1117/1.OE.56.8.084106.

High-resolution imaging using a translating coded aperture

Abhijit Mahalanobis,^{a,*} Richard Shilling,^a Robert Muise,^a and Mark Neifeld^b

^aLockheed Martin, Missiles and Fire Control, Orlando, Florida, United States

^bUniversity of Arizona, Department of Electrical and Computer Engineering, Tucson, Arizona, United States

Abstract. It is well known that a translating mask can optically encode low-resolution measurements from which higher resolution images can be computationally reconstructed. We experimentally demonstrate that this principle can be used to achieve substantial increase in image resolution compared to the size of the focal plane array (FPA). Specifically, we describe a scalable architecture with a translating mask (also referred to as a coded aperture) that achieves eightfold resolution improvement (or 64:1 increase in the number of pixels compared to the number of focal plane detector elements). The imaging architecture is described in terms of general design parameters (such as field of view and angular resolution, dimensions of the mask, and the detector and FPA sizes), and some of the underlying design trades are discussed. Experiments conducted with different mask patterns and reconstruction algorithms illustrate how these parameters affect the resolution of the reconstructed image. Initial experimental results also demonstrate that the architecture can directly support task-specific information sensing for detection and tracking, and that moving objects can be reconstructed separately from the stationary background using motion priors. © 2017 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: [10.1117/1.OE.56.8.084106](https://doi.org/10.1117/1.OE.56.8.084106)]

Keywords: computational imaging; coded aperture; digital super resolution; task specific imaging.

Paper 170592 received Apr. 22, 2017; accepted for publication Jul. 26, 2017; published online Aug. 22, 2017.

1 Introduction

Traditional imaging systems utilize very large format focal plane arrays (FPAs) to produce high-resolution digital images. High-definition visible band FPAs are relatively inexpensive and abound in commercial cameras. However for other applications, such as infrared imaging, FPAs are still expensive due to inherent limits on yield and fabrication processes. In such cases, digital super-resolution (Here, super-resolution refers to resolution of subpixel image features, and not imaging beyond the diffraction limit of the optical aperture.) techniques have proven useful for enhancing resolution by combining multiple low-resolution images with subpixel shifts. It is well known that simple averaging of frames with subpixel registration can be used to improve resolution by a factor of two or thereabouts, while more powerful algorithms that employ forward models and smoothness constraints can achieve considerably better improvements in image resolution.^{1–6}

Computational imaging (CI) with coded apertures provides an alternate approach for high-resolution imaging using low-resolution measurements⁷ by optically modulating the light field before it is digitally captured, so that high spatial frequencies can be algorithmically recovered from several such encoded measurements. Although it is possible to encode information at the entrance pupil,⁸ many CI architectures employ a spatial light modulator (SLM) in an intermediate image plane to implement the coding strategy. The choice of the SLM is an important part of the CI architecture. Several teams have reported the use of reflective digital micromirror devices (DMDs), which are commercially

available in high-speed large formats. For example, Sun et al.⁹ have suggested the use of a DMD to optically modulate the scene with a series of shifted Hadamard patterns, to encode the scene at high resolution before it is acquired at a lower resolution by the FPA. These coded images can be co-registered on a high-resolution grid to enhance the image resolution. Olivas et al.¹⁰ have evaluated the performance of a DMD-based CI system using multiple basis sets on monochromatic, color, and infrared natural light scenes, and compared the results to a camera with the same number of (conventional pixel) measurements. McMackin et al.¹¹ have described a compressive short-wave infrared (SWIR) imager (900 nm to 1.7 μm) camera with a DMD and a single-element sensor. They showed that the hardware and software combination makes it possible to create images with the resolution of the DMD while employing a substantially lower cost sensor subsystem than would otherwise be required by the use of traditional FPAs. Mahalanobis et al.¹² demonstrated a CI testbed for midwave infrared using a DMD and a programmable FPA with variable pixel sizes. Using this testbed, they experimentally showed that high-resolution images can be successfully reconstructed from compressive measurements made with a much smaller (and potentially cheaper) infrared FPA.

Electronically addressable SLMs are not without their disadvantages. Although versatile, reflective SLMs require folded optical paths that can be challenging to design, particularly for larger fields of view and infrared applications. Transmissive SLMs permit on-axis optical geometries, but are not available in high-speed large formats. Such considerations have led several teams to consider a transparent

*Address all correspondence to: Abhijit Mahalanobis, E-mail: abhijit.mahalanobis@lmco.com

mask with a fixed pattern as an elegant and inexpensive alternative to electronically addressable SLMs that can also scale to arbitrary sizes.

In fact, a moving mask etched with a fixed pattern has been found to produce surprisingly good results as reported by other researchers. Llull et al.¹³ have proposed the mechanical translation of an encoded mask for low power coded aperture compressive temporal imaging. Coding is implemented by a chrome-on-glass binary transmission mask in an intermediate image plane. In contrast with previous approaches, modulation of the image data stream by harmonic oscillation of this mask requires very little operating power. Such masks were also demonstrated for compressive imaging in coded aperture snapshot spectral imagers,¹⁴ which include an intermediate image plane before a spectrally dispersive relay optic. Koller et al.¹⁵ have described a prototype compressive video camera that encodes scene movement using a translated binary photomask in the optical path. The encoded recording can then be used to reconstruct multiple output frames from each captured image, effectively synthesizing high-speed video. Recently, Don et al.¹⁶ have reported an architecture with a rotating coded aperture with an optimum pattern that minimizes the mutual coherence of the sensing matrix. This concept does not require any mechanical translation of the mask and is particularly attractive for applications where the entire platform is rotating about the imaging axis.

Although the concept of imaging with moving coded apertures is well documented in the literature, previous work has not explicitly reported the 64:1 increase in high-resolution image pixels compared to the size of FPA. While we have previously reported initial results showing fourfold increase in resolution,¹⁷ our goal here is to demonstrate that these concepts are scalable, and can be employed to significantly enhance the resolution of images of outdoor scenes containing both stationary and moving objects. Specifically, the contributions of the paper are:

- Experimental results with outdoor scenes that demonstrate the ability to achieve an eightfold increase in image resolution in each dimension. This implies that a 1 Mpixel FPA can be used to obtain 64 Mpixel high-resolution imagery.
- A comparison of results obtained using masks designed with different element sizes for four- and eightfold increases in resolutions. The degradation in resolution compared to the ideal image is modeled as a linear blur, which is shown to be twice as large (i.e., worse) for the fourfold case than it is for eightfold resolution enhancement.
- A comparison of results obtained with two different reconstruction algorithms of different computational complexities. It is shown that the full image can be reconstructed reasonably well using a computationally inexpensive iterative gradient descent algorithm. Alternately, a closed form least-squares algorithm may be used to reconstruct the scene block by block. This approach yields better results, but is computationally more expensive since matrix inversions are required.
- A discussion of the properties of the sensing matrix. The stability of the matrix depends not only on the mask pattern, but also on the displacement of the

mask between measurements. Boundary artifacts between adjoining blocks are less prominent when the matrices are normalized to satisfy the “photon constraint”¹⁸ to account for finite measurement resources (such as integration time or equivalently the number of photons that are available during each measurement).

- Preliminary results showing the detection, tracking, and reconstruction of moving objects using motion priors. This shows how a coded aperture system can sense information for a specific task, and not for image reconstruction alone.

The rest of the paper is organized as follows: Sec. 2 is a review of the architecture and experimental setup for intermediate image plane encoding using a moving mask. We discuss how the mask enables the FPA to sense image data at a higher spatial resolution than it otherwise would based on its pixel pitch. The algorithms used for image reconstruction are described in Sec. 3. For simplicity, we use minimum L2-norm algorithms for image reconstruction, although the architecture allows a powerful sparsity-based technique to be employed. Section 4 presents experimental results that show the reconstruction of images of a scene with a stationary background and moving objects. These experiments demonstrate the ability to achieve the expected improvement in resolution. Finally, Sec. 5 presents our conclusions along with a discussion of the lessons learned and potential future work.

2 Overview of Sensing Architecture and Experimental Setup

The idea of image plane coding with a moving mask is shown in Fig. 1. The objective lens images the scene on the mask that modulates the light with a spatially varying transmittance. The smallest elements of the transmittance function etched on the mask (also referred to in this paper as “mask elements”) must be equal in size to the diffraction limited blur spot formed by the objective lens at the image plane. This ensures that the highest spatial frequencies present in the image will be encoded before the light is collected at the FPA. The relay lens images the mask onto an FPA with demagnification so that each detector receives light from a group of high-resolution mask elements. In the absence of the mask, the FPA would capture an under-sampled image of the scene, with resolution limited by the detector pitch. While the diameter of the objective lens determines the limiting angular resolution of the system, the relay optics determines the entrance pupil diameter and the overall light collection ability of the sensor.

Fig. 1 The high-resolution image formed by the objective lens is encoded by the mask, and reimaged with demagnification on the FPA.

The fundamental parameters governing the architecture design are as follows: let us assume that the field of view (FOV) of the system and the angular resolution required of the optics (θ) are specified. The number of high-resolution pixels required in each dimension is $N_p = \frac{\text{FOV}}{\theta}$. Assuming diffraction limited optics, the diameter of the objective lens is $D = \frac{1.22\lambda}{\theta}$, where λ is the wavelength of light. The diameter of the blur spot formed at the mask is $s = 2.44\lambda F_\#$, where $F_\#$ is the f -number of the objective lens. Therefore, the size of the mask is

$$S_M = s * N_p = 2.44\lambda F_\# \frac{\text{FOV}}{\theta}. \quad (1)$$

Now assume that we are given an FPA of size $S_F = dN_d$, where d is the detector size, and N_d is the number of detectors in the array. Therefore, the relay optics has to image the mask on the FPA with a demagnification factor of

$$M = \frac{S_M}{S_F} = \frac{(\text{FOV})2.44\lambda F_\#}{dN_d\theta}. \quad (2)$$

Equation (2) allows us to determine the required demagnification given the FOV, the optical angular resolution, and the size of the FPA. In general, higher demagnification (which implies greater super-resolution gains) can be achieved by reducing the focal length between the relay optics and the FPA. In practice however, lenses with $F/\#s < 1$ are expensive (if not impractical). It is also important to note that for the reimager to capture all the light collected by the objective lens, the $F/\#$ between the objective lens and the mask has to be equal to the $F/\#$ from the mask to the relay lens. Otherwise, only a fraction of the light is collected by the relay optics, equal to the ratio of the two $F/\#s$. Thus, the system design must trade the mask size and the relay imager optics given the constraints (such as FOV, angular resolution, detector pitch, and the size of the FPA) to achieve the desired resolution improvement.

We limit our discussion to binary mask functions that have a fixed pattern of “open” and “closed” mask elements, equivalent to transmittance values of “one” and “zero,” respectively. Figure 2 shows the relation between the mask elements and the individual detectors on the FPA necessary to achieve a fourfold increase in image resolution. In this example, the mask is imaged with 4:1 demagnification, so that a $d \times d$ detector collects light from a $4d \times 4d$ area of the mask. If the mask elements are also of size $d \times d$, the optics images a 4×4 block of such elements on the detector. The example in Fig. 2 shows one of sixteen mask elements open so that the detector collects light only over 1/16th of its surface area. Other regions of the detector can be exposed by opening the corresponding mask elements one at a time. In a sense, this is equivalent to “subsampling” or “microscanning” the detector’s instantaneous field of view (iFOV), and effectively increases the sampling rate in each direction by a factor of four. Although the opening and closing elements require an electronically addressable SLM, the same effect can be achieved by moving the mask with respect to the detector. For instance, the mask can be tiled with a sequence of 4×4 elements with a different one open in each block, and then translated in the x - y plane to expose

Fig. 2 Each detector on the FPA receives light from a block of elements on the mask. Here, the detector collects light from a group of 4×4 elements, of which only one is open.

the detector to different sections of the iFOV. An example of this is shown in Fig. 3, where $d = 25 \mu\text{m}$ so that $100 \mu\text{m} \times 100 \mu\text{m}$ area on the mask is imaged on a $25 \mu\text{m} \times 25 \mu\text{m}$ detector. If the mask moves from left to right (or vice versa) in increments of $25 \mu\text{m}$, the pattern in Fig. 3(a) allows light from only one of the sixteen possible locations within the 4×4 block (i.e., where the mask element is open) to reach the detector. Effectively, this allows the $25 \mu\text{m}$ detector to sample the intermediate image plane with a surface area of $6.25 \mu\text{m}$, i.e., at four times the spatial resolution than if the mask were absent. In principle, 16 such frames of data can be combined with appropriate upsampling and shifts to reconstruct the image with four times the spatial resolution of the FPA.

In CI vernacular, the pattern in Fig. 3(a) is essentially the “unit” code and equivalent to making measurements with an operator proportional to the “identity” matrix (i.e., with the proportionality constant determined by the photon efficiency of the apparatus). Multiplexed measurements can be readily made using patterns with more than one open element in each 4×4 group, as shown in Fig. 3(b). This pattern is a strip of 16 blocks of 4×4 mask elements with half of them open. Stepping the mask in increments of $100 \mu\text{m}$ results in 16 different multiplexed measurements by the detector, which can be used for a full-rank reconstruction of the corresponding 4×4 area of the high-resolution image. In general, multiplexed measurements result in better signal to noise ratio since more photons are collected during the integration period.

It should be noted that the image resolution can be varied by simply changing the size of the mask elements. For instance in Figs. 2 and 3(a), 10-fold increase in resolution can be achieved by simply reducing the size of the individual mask elements to $10 \mu\text{m}$. The 4:1 demagnification provided by the reimager would now map 10×10 mask elements (over the same $100 \mu\text{m} \times 100 \mu\text{m}$ area of the mask) to a single $25 \mu\text{m}$ detector. Of course, the objective lens must be large enough to provide the necessary angular resolution, and the mask would also have to be stepped in increments of $10 \mu\text{m}$, and a hundred such measurements would be required for a full-rank reconstruction of the high-resolution image.

Fig. 3 With only one element open in each 4×4 block, the mask pattern in (a) enables the detector to microscan the iFOV as it moves in increments of $25 \mu\text{m}$. The pattern in (b) was created by placing sixteen blocks of 4×4 apertures in a row so that a complete set of measurements is obtained by moving the mask in steps of $100 \mu\text{m}$.

2.1 Experimental Setup

The optical system shown in Fig. 4 was used to demonstrate both four- and eightfold improvements in image resolution. The objective lens is a 100 mm, $F/1.5$ lens, with a 12 deg circular FOV that is commercially available from StingRay Optics (SR0939-A01). The plots in Fig. 4(b) show the encircled energy as a function of the “blur spot” size at different distances from the center at the focal plane. Although the lens specifications indicate that $>90\%$ of the encircled energy is within a diameter of $25 \mu\text{m}$ anywhere in the FOV, we determined that the optical blur spot can be reduced to $10 \mu\text{m}$ by stopping down the aperture to $F/4$ and attenuating wavelengths $>700 \text{ nm}$ using a spectral filter. The relay imager was designed as a finite conjugate $F/2$ optics that provides a 4:1 demagnification so that a $16 \text{ mm} \times 16 \text{ mm}$

area on the mask forms a $4 \text{ mm} \times 4 \text{ mm}$ image on the FPA. The object to image distance is 367 mm, and the optics is telecentric at the mask, which helps to ensure that no stray light scatters from the mask into the camera. Two different sensors were used for acquiring the data. Earlier experiments used a Sensors Unlimited 640HSX InGaAs camera with $25 \mu\text{m}$ pixels, which provides 640×512 imagery at 30 Hz. This was later replaced by an Allied Vision visible band CMOS sensor with $2.2 \mu\text{m}$ pixels that produces 12-bit monochrome images in 2592×1944 format at 14 Hz. This visible camera has a software selectable “digital binning” mode, which produces larger “superpixels” created by aggregating blocks of native pixels. This feature was useful for accurate calibration between the mask and the FPA, and for digitally varying the effective pixel size.

Fig. 4 The testbed used to demonstrate the underpinnings of the proposed CI architecture is shown with (a) various components of the optical setup and (b) blur spot of the objective lens.

Fig. 5 The 16 mm × 17 mm masks are etched on a 1-mm thick fused silica glass.

Figure 5 shows pictures of the 16 mm × 17 mm masks that are etched on a 1-mm thick fused silica glass substrate. The masks are mounted on a Physik Instrumente M686 scanner. Using this device, the mask can be translated in increments of 0.3 μm at a speed of 100 mm/s. The controller is set up to move the mask along the optical axis for focusing, and horizontally in the intermediate image plane for generating coded measurements. The frame grabber is synced with the stage controller to acquire an image when the movement settles at each commanded location. The scanner has remarkable linearity and is highly repeatable. A useful feature of the scanner is that it reports the actual location where it stops so that this can be compared to the commanded location, with observed difference between the two <0.001 μm.

Two different masks patterns were used for the experiments. Images of these masks as viewed through the reimager and acquired at high resolution with the Allied Vision camera are shown in Fig. 6. The first is a periodic pattern of 25 μm mask elements shown in Fig. 6(a), and the second is a random pattern of 10 μm mask elements with 40% open shown in Fig. 6(b). These will be henceforth referred to as the periodic mask and random mask, respectively. (For the sake of convenience, we will omit the explicit reference to the size of the elements for the two masks with the understanding that the elements on the random mask are 10 μm and those on the periodic mask are 25 μm.) Since the reimager has 1:4 demagnification and the FPA pixels are 2.2 μm, the 25 μm mask elements occupy ~3 × 3 pixels in Fig. 6(a), while the 10 μm mask elements are ~1 × 1 pixels in Fig. 6(b). The random mask will be used to demonstrate an eightfold

increase in resolution, while the periodic mask supports only fourfold resolution improvement.

3 Image Reconstruction Algorithms

The architecture described in Sec. 2 encodes (i.e., optically modulates) high-resolution spatial information with a moving mask before it is digitally captured at much lower resolution on a relatively small FPA. For a binary mask, the weights of the code are either “1” if the element is open or “0” if it is closed. The code changes as the mask moves, and a sequence of measurements is obtained over time. Thus, the moving mask implements the “measurement” matrix or the “forward transform.”

The forward model for the imaging architecture is shown in Fig. 7. Assume that an $M \times N$ block of the intermediate image $\mathbf{x}(m, n)$ is weighted by the mask function $c_k(m, n)$. The subscript k , ($1 \leq k \leq K$) is a discrete time index for the position of the mask as the k 'th measurement is made. The product of the scene and the mask is filtered by the reimager point spread function (PSF) $h(m, n)$ yielding

$$\hat{\mathbf{x}}_k(m, n) = h(m, n) * [\mathbf{x}(m, n) \mathbf{c}_k(m, n)], \quad (3)$$

where $*$ indicates the two-dimensional (2-D) convolution operation. In matrix-vector notation, this can be expressed as

$$\hat{\mathbf{x}}_k = \mathbf{H} \mathbf{C}_k \mathbf{x}, \quad (4)$$

where $\hat{\mathbf{x}}_k$ and $\mathbf{x}$ are $MN \times 1$ dimensional column vectors obtained by lexicographically rearranging the elements of $\hat{\mathbf{x}}_k(m, n)$ and $\mathbf{x}(m, n)$, respectively; $\mathbf{C}_k$ is an $MN \times MN$ diagonal matrix with the elements of $c_k(m, n)$ along the main diagonal; and $\mathbf{H}$ is an $MN \times MN$ block Toeplitz convolution matrix that implements the spatial filtering operation with $h(m, n)$.

The light arriving at the FPA is integrated over the surface area of each detector before it is digitized. This is modeled as a “box” spatial filter $s(m, n)$ followed by a down sampler that produces the FPA output $y_k(m, n)$ of size $P \times Q$. The box size is $B \times B$, where

Fig. 6 Viewed at high resolution through the reimager, (a) shows an image of the mask with a periodic pattern of 25 μm elements [also shown in Fig. 3(b)], and (b) shows an image of the mask with a random pattern of 10 μm elements with 40% open.

Fig. 7 The sensed data are modeled as the elementwise product of the high-resolution intermediate image and the mask, followed by the optical PSF, and the integration and down sampling operation of the detector.

$$B = \frac{M}{P} = \frac{N}{Q}, \quad (5)$$

is the ratio of the dimensions $x(m,n)$ and $y_k(m,n)$. This “integrate and down-sample” process be formulated as a $PQ \times MN$ matrix operator $\mathbf{S}$ by first defining an $MN \times MN$ block Toeplitz matrix to implement the convolution with the box filter, and then retaining one in every B row (and discarding the rest) so that the output of the FPA at time instant k is given by

$$\mathbf{y}_k = \hat{\mathbf{S}} \mathbf{x}_k, \quad (6)$$

or

$$\mathbf{y}_k = \mathbf{S} \mathbf{H} \mathbf{C}_k \mathbf{x} = \mathbf{A}_k \mathbf{x}, \quad (7)$$

where $\mathbf{A}_k = \mathbf{S} \mathbf{H} \mathbf{C}_k$ is the end-to-end $PQ \times MN$ “measurement” matrix at time instant k . The elements of the $PQ \times 1$ column vector $\mathbf{y}_k$ can be readily rearranged into the $P \times Q$ dimensional FPA output $\mathbf{y}_k(m,n)$.

As the mask translates over time, K such measurements are made which can be collectively expressed as

$$\begin{bmatrix} \mathbf{y}_1 \\ \mathbf{y}_2 \\ \vdots \\ \mathbf{y}_K \end{bmatrix} = \begin{bmatrix} \mathbf{A}_1 \\ \mathbf{A}_2 \\ \vdots \\ \mathbf{A}_K \end{bmatrix} \mathbf{x}, \quad (8)$$

or simply as

$$\mathbf{u} = \mathbf{P} \mathbf{x}, \quad (9)$$

where $\mathbf{P} = [\mathbf{A}_1^T \mathbf{A}_2^T \cdots \mathbf{A}_K^T]^T$ is a $KPQ \times MN$ matrix, and $\mathbf{u} = [\mathbf{y}_1^T \mathbf{y}_2^T \cdots \mathbf{y}_K^T]^T$ is a $KPQ \times 1$ vector. It should be noted that to satisfy the photon constraint,¹⁸ $\mathbf{P}$ must be multiplied by a scalar value such that the L1 norm of its columns is less than or equal to unity. Specifically, the measurement matrix is multiplied by a scalar value to normalize the largest column L1 norm to 1.0, which also ensures that the rest of the columns do not violate the photon constraint. This ensures that the measurement matrix is not changed except for a scalar value.

If $KPQ \ll MN$, then $\mathbf{x}$ can be recovered from Eq. (9) using sparsity-based minimum L1-norm techniques developed for compressive sensing.¹⁹ However, for the purpose of our experiments, we assume that the size of the block to be reconstructed is relatively small and that $KPQ \geq MN$ so that a least-squares technique can be used for estimating $\mathbf{x}$.

Specifically, we seek a solution for $\mathbf{x}$ that minimizes the objective function

$$J(\mathbf{x}) = \|\mathbf{u} - \mathbf{P} \mathbf{x}\|^2 + \delta \|\mathbf{x}\|^2, \quad (10)$$

where δ is a “Tikhonov” regularization parameter that emphasizes a minimum L2-norm solution for $\mathbf{x}$. It is well known²⁰ that the optimum solution is given by

$$\mathbf{x} = (\mathbf{P}^T \mathbf{P} + \delta \mathbf{I})^{-1} \mathbf{P}^T \mathbf{u}, \quad (11)$$

with $\mathbf{I}$ representing the identity matrix. It should be noted that if $\mathbf{P}^T \mathbf{P}$ is invertible, then δ can be zero, leading to the generalized inverse solution. If, however, $\mathbf{P}^T \mathbf{P}$ is ill-conditioned, then nonzero values of δ avoid numerical issues with matrix inversion, but results obtained with large values of δ are often not satisfactory. It is therefore desirable to ensure that $\mathbf{P}^T \mathbf{P}$ has as small a condition number as possible, and that the inverse problem in Eq. (8) is well behaved for relatively small values of δ .

Many factors affect the condition number of $\mathbf{P}^T \mathbf{P}$. While some of these factors (e.g., choice of the mask pattern and reimager blur) cannot be easily changed once implemented in hardware, others (such as the amount by which the mask is shifted between measurements, the number of measurements, and the size of the block of data used for reconstruction) can be selected during data collection and processing. We now discuss how these “software-controlled” parameters can be selected to reduce the condition number of $\mathbf{P}^T \mathbf{P}$ for the particular mask patterns described in Sec. 2.

The reconstruction algorithm operates on one $P \times Q$ low-resolution block at a time, to produce an $M \times N$ block of the high-resolution image, where $B = \frac{M}{P} = \frac{N}{Q}$ is the resolution improvement factor. The matrix $\mathbf{P}$ is defined in terms of the block matrices $\mathbf{A}_k$ whose relation to $c_k(m,n)$ is given in Eq. (7). Thus, $\mathbf{P}$ is constructed using several images of the mask observed at different translated positions for all K measurements. However, the distance by which the mask is physically translated (i.e., the “stepping distance”) between two successive measurements can vary, and thereby affect the condition number of $\mathbf{P}^T \mathbf{P}$.

Our goal is to find good choices for the block size. [Here, “block” refers to the section of the FPA data used to reconstruct the corresponding part of the high-resolution image, and not the block (or region) of the high-resolution image formed at the mask that is captured by a single detector

on the FPA.] and the stepping distance for the two masks [i.e., the periodic mask in Fig. 6(a) and the random mask in Fig. 6(b)]. Toward this end, we calculated the condition number of $\mathbf{P}^T\mathbf{P}$ by varying the stepping distance between measurements in multiples of 5μ using two different block sizes. Although a minimum of B^2 is required for a full-rank reconstruction, the number of measurements was chosen to be $K = 2B^2$ to reduce the effects of noise on the trends. Figure 8(a) shows the results for the periodic mask for $B = 4$, and block sizes of $P = Q = 4$ and $P = Q = 8$, which then correspond to $M = N = 16$ and $M = N = 32$, respectively. We observe that the condition number is the smallest for the 4×4 block size when the mask is translated by five steps (or 25μ) between measurements. We also observe that the condition number for the 8×8 block size initially decreases with increasing step size, but then rises again beyond five steps (or 25μ stepping distance) between measurements, possibly due to the periodicity of the mask pattern. In Fig. 8(b), we see that for the random mask and $B = 8$, the smallest condition number is obtained using a block size of $P = 4, Q = 4$ (i.e., $M = 32, N = 32$), and for a translation of three steps (or a stepping distance of 15μ) between measurements. It is also interesting to note that the condition numbers are generally smaller in Fig. 8(b) for the random mask than they are for the periodic mask in Fig. 8(a). This suggests that in general, better results will be obtained using the random mask, and that the periodic mask will yield comparatively poorer results. We also conclude that the best results for an eightfold increase in resolution is expected using the random mask, block sizes of 4×4 , and a stepping distance of 15μ .

3.1 Gradient Descent

In this section, we discuss an iterative gradient descent method as an alternative technique for image reconstruction that avoids matrix inversion, and thereby allows a large image to be reconstructed as a whole without the need for blockwise operations. Recall that the mask moves to a different position as new measurements are made over time. Given a current estimate of the image to be reconstructed $\tilde{\mathbf{x}}_k$ at time instant k , the forward model in Eq. (7) can be used to compute

$$\tilde{\mathbf{y}}_k = \mathbf{S}\mathbf{H}\mathbf{C}_k\tilde{\mathbf{x}}_k, \quad (12)$$

as an “estimate” for actual FPA output $\mathbf{y}_k$. The mean squared error between $\tilde{\mathbf{y}}_k$ and $\mathbf{y}_k$ at the k 'th iteration is

$$|e_k|^2 = |\tilde{\mathbf{y}}_k - \mathbf{y}_k|^2 = |\mathbf{S}\mathbf{H}\mathbf{C}_k\tilde{\mathbf{x}}_k - \mathbf{y}_k|^2. \quad (13)$$

Setting $\mathbf{A}_k = \mathbf{S}\mathbf{H}\mathbf{C}_k$, the squared error can be expressed as

$$|e_k|^2 = \tilde{\mathbf{x}}_k^T \mathbf{A}_k^T \mathbf{A}_k \tilde{\mathbf{x}}_k + \mathbf{y}_k^T \mathbf{y}_k - 2\tilde{\mathbf{x}}_k^T \mathbf{A}_k^T \mathbf{y}_k. \quad (14)$$

The gradient of $|e_k|^2$ with respect to $\tilde{\mathbf{x}}_k$ is

$$\nabla_{\tilde{\mathbf{x}}_k}(|e_k|^2) = 2\mathbf{A}_k^T \mathbf{A}_k \tilde{\mathbf{x}}_k - 2\mathbf{A}_k^T \mathbf{y}_k = 2\mathbf{A}_k^T e_k. \quad (15)$$

Thus, the updated rule for iteratively estimating the high-resolution image is

$$\tilde{\mathbf{x}}_{k+1} = \tilde{\mathbf{x}}_k - \mu \mathbf{A}_k^T e_k, \quad (16)$$

where a factor of 2 has been absorbed in the definition of the step-size parameter μ . Thus, at each iteration, the reconstructed image is updated by a fraction of the error vector transformed by $\mathbf{A}_k^T$.

4 Results of Experiments

We have previously reported the results of fourfold resolution improvement using the SWIR camera with 25μ pixels in Ref. 21. An example of one such early experiment is shown in Fig. 9. The data were collected indoors (in the laboratory) under good ambient lighting conditions. The gain and exposure time were chosen to ensure that the data were not saturated and that the full dynamic range of the sensor was utilized. The gradient descent algorithm described in Sec. 3.1 was applied to a sequence of coded frames such as that shown in Fig. 9(a) to obtain the reconstruction result shown in Fig. 9(b). The details on the face of the clock (i.e., the numbers indicating the hours and the markings at 1 min intervals) are well resolved, and the improvement in resolution is evident. However, the reconstruction of the minute arm is blurry due to its motion during the time the frames are collected. The plot in Fig. 10 shows that

Fig. 8 The condition number of the matrix $\mathbf{P}^T\mathbf{P}$ is shown for two different block sizes as a function of the stepping distance for the (a) periodic mask and (b) the random mask.

Fig. 9 The gradient descent algorithm used low-resolution coded frame such as in (a) to create the high-resolution image in (b) in $k = 25$ iterations.

Fig. 10 The MSE converged to a minimum value after 25 iterations.

the reconstruction error converged to a minimum value with ~ 25 data frames.

Not surprisingly, the quality of the results is highly dependent on how accurately the actual physical position of the mask is represented by the function $c_k(m, n)$ in Eq. (3). For instance, an ideal version of a mask with the microscanning pattern from Fig. 3(a) is shown in Fig. 11(a), while Fig. 11(b) shows its image acquired by the SWIR camera when mounted in the system. It is clear that the mask is slightly rotated due to imperfect mounting. Further, the mask may be laterally offset by a small but unknown amount, and the demagnification provided by the re-imager may not be exact. Based on such observations, we have previously modeled²¹ the relation between the ideal mask pattern and its true position in the system as an “affine transform.” In principle, the transform parameters can be estimated for all shifted positions of the stage on which the mask is mounted by registering the ideal mask to the observed image, and the transformed version of the ideal mask [such as in Fig. 11(c)] can be treated as an estimate for $c_k(m, n)$ for $1 \leq k \leq K$.

Although the initial results in Fig. 9 were promising, satisfactory results could not be obtained at higher resolution using this approach for mask registration. The reason is that an affine transform is not always sufficient for modeling the distortions introduced by the optics, particularly away from the center of the FOV. Further, for very high-resolution reconstructions, it is essential to register the ideal mask pattern to the observed image of the mask with a spatial accuracy equivalent to the size of the reconstructed pixel, which is not possible if the observed image of the mask [such as in Fig. 11(b)] is acquired at the same low resolution as the coded measurements.

Since accurate mask registration with low-resolution images is difficult, an alternative is to directly obtain $c_k(m, n)$ by recording high-resolution images of the mask for all $1 \leq k \leq K$. Examples of such high-resolution mask images collected with the 2.2μ pixel visible band camera are shown in Fig. 6 and used in the experiments discussed in the rest of this section. While this avoids the approximation errors of numerical registration, it also requires two different FPA settings: high resolution for imaging the mask, and low resolution for acquiring the coded image frames. It should be noted that while the pixel size can be digitally controlled in a testbed for experimental purposes, the factory calibration of a sensor equipped with only a low-resolution FPA may require an in-situ beam splitter to externally capture high-resolution mask images.

4.1 Results of Block Reconstruction

An image of an outdoor scene acquired with a low-resolution FPA is shown in Fig. 12(a). The desired high-resolution version in Fig. 12(b) shows the details in the windows, walls, and roof of the building, a bridge, and electrical pole in the foreground, vegetation, and background objects. The well-known structural similarity (SSIM) measure²² between the two images was calculated to be 0.34. The data were collected in sunny daytime conditions, and the camera gain and exposure were adjusted such that the histogram of the scene covered the full dynamic range of the sensor but without saturating the image. An example of a low-resolution image of scene encoded with the random mask is shown in Fig. 13(a). The high-resolution image in Fig. 13(b) was obtained by tiling together 32×32 blocks, each of which was reconstructed using $K = 128$ coded measurements of the corresponding 4×4 block of the low-resolution data. The SSIM measure between this and the ideal image in Fig. 12(b) is 0.94. Zoomed regions of the data and reconstructed image are shown in Figs. 13(c) and 13(d), respectively, to demonstrate the dramatic increase in image resolution. In fact, the resolution and details (such as various features on the building, the electrical pole, background objects, and the vegetation) in Figs. 13(b) and 13(d) compare very well to that of the ideal image in Fig. 12(b). However, some artifacts due to moving objects are also visible just above the bridge, an issue which will be discussed in Sec. 4.2.

For comparative purposes, Fig. 14 shows the results of reconstruction with a fourfold increase in resolution (i.e., $B = 4$) obtained using the periodic mask. The SSIM measure between this and the ideal image is only 0.66. The result was obtained by tiling 16×16 blocks that were reconstructed using the corresponding 4×4 block of

Fig. 11 The ideal mask pattern is shown in (a) and its image as observed through the reimager is shown in (b). The result of transforming the ideal pattern so that it matches the observed image is shown in (c) and is used as the function $c_k(m, n)$ for image reconstruction.

Fig. 12 An outdoor scene (a) as captured at low resolution by the FPA with an SSIM measure of 0.34 and (b) the ideal high-resolution version showing the details to be reconstructed.

the encoded data, with $K = 32$ measurements per block that were acquired in steps of 25μ . Although the result shows an improvement in resolution compared to the original data in Fig. 12(a), the details in the scene are blurrier compared to the reconstruction with eight-fold resolution improvement in Fig. 13. As noted previously, this result is not unexpected based on Fig. 8 which shows that the condition number of $\mathbf{P}^T \mathbf{P}$ is larger using the periodic mask than it is for the random mask.

To further compare the results obtained with the two different masks, we modeled the degradations observed in the reconstructions as the effects of a 2-D “blur function” $p(m, n)$, estimated as a linear operator that minimizes the mean squared error between the reconstruction $x_r(m, n)$ and ideal image of the scene $x_i(m, n)$ given by

$$mse = \sum_{m,n} |p(m, n) * x_i(m, n) - x_r(m, n)|^2.$$

We obtained a least-squares solution for $p(m, n)$ by treating it as a 13×13 kernel represented as a 169×1 dimensional vector, which when multiplied by a matrix representing the convolution with $x_i(m, n)$ yields $x_r(m, n)$ in vector form. The narrower the width of the blur function, the better the reconstruction. It should be noted that this operator is an approximation, and the actual sources of degradation may be neither linear nor spatially invariant. Nevertheless, the results in Fig. 15 provide interesting insight into the resolution improvements achieved using the different masks, as compared to the original low-resolution data.

In this figure, the horizontal axis denotes the cross section and the vertical axis represents the amplitude of the blur function. The solid line represents the cross section of the blur that relates the ideal scene to the result in Fig. 13(b), while the dotted line marked with “x” represents the same with respect to the result in Fig. 14. Recall that these represent eight- and fourfold increases in resolution,

Fig. 13 The image in (a) is an example of data coded with the random mask. The result in (b) was obtained by tiling together 32×32 image blocks that were reconstructed using the corresponding 4×4 block of encoded data. The SSIM measure between this image and the ideal version in Fig. 13(b) is 0.94. Some details of the images in (a) and (b) are shown in (c) and (d), respectively.

Fig. 14 The results obtained with periodic mask have fourfold greater resolution than the original data (SSIM = 0.66) but are comparatively blurry and noisy due to the poorly conditioned measurement matrix.

respectively. It is intuitively satisfying to note that the width of the solid line plot is approximately half that of the dotted line plot, thereby confirming that the random mask indeed yields twice the resolution achieved using the periodic mask.

4.2 Results of Stationary and Moving Object Reconstruction Using Iterative Gradient Descent

A reconstruction obtained using the random mask and the iterative gradient descent algorithm (described in Sec. 3.1) is shown in Fig. 16. Although the overall image appears to be qualitatively comparable to the result in Fig. 13(b), the SSIM measure is 0.86 and the details in the inset are blurry compared to the close-up in Fig. 13(d). Based on such observations, we conclude that the gradient descent algorithm can reconstruct stationary objects relatively well, but with somewhat less clarity than the blockwise approach.

Fig. 15 The cross section of the blur function $p(m, n)$ that relates the reconstructed and ideal high-resolution image is narrower for the result obtained with the random mask. In fact, it is half the width of the blur estimated for the periodic mask, which is consistent with the difference in resolution between the 8× and 4× reconstructions

Fig. 16 Results obtained with gradient descent algorithm exhibit improved resolution (SSIM = 0.86), but the details are somewhat blurry compared to the results in Fig. 13.

The movement of the mask is necessary to make coded measurements for the reconstruction of the stationary objects and static background. For moving objects however, it may not be possible to move the mask fast enough to match their speeds. Therefore, we propose an alternate approach that keeps the mask stationary, and the moving objects are observed through different regions of the mask by virtue of their motion. This concept is shown in Fig. 17 where a system controller either moves the mask to make measurements about static parts in the scene or holds it stationary to collect data for moving objects. We assume a tracker (which is external to the sensor) can provide strong priors for objects under track such as velocity estimates and probabilities of movement in different directions. This information can be used to modify the adaptive gradient descent algorithm to reconstruct moving objects by simply incorporating a “motion model” to propagate the updated term.

Essentially, the motion model is the probability that the object may translate by $(\Delta m, \Delta n)$ pixels between the observation at time instances $k - 1$ and k , i.e.,

$$\begin{aligned} \mathbf{V}(\Delta m, \Delta n) &= \mathbf{P}[x_k(m, n)] \\ &= \delta(m - \Delta m, n - \Delta n) * \mathbf{x}_{k-1}(m, n), \end{aligned} \quad (17)$$

where $\delta(m, n)$ is the delta function. Figure 18 shows an example of a motion model centered at the most likely displacement of the object between frames, which in this case is $\Delta m = 25$, $\Delta n = 75$. The ellipsoid shape of the distribution also indicates that the uncertainty is greater along Δn than Δm . In practice, such models for object motions are readily derived from the track state variables, or from prior knowledge of object velocities based on traffic patterns and surveillance.

To incorporate the motion model in the adaptive gradient descent algorithm, we represent the current estimate of the reconstructed object image as a convolution of the motion model and the previous estimate, i.e.,

$$\hat{\mathbf{x}}_k(m, n) = \mathbf{V}_k(m, n) * \hat{\mathbf{x}}_{k-1}(m, n), \quad (18)$$

which can be expressed in matrix vector notation as

$$\hat{\mathbf{x}}_k = \mathbf{V}_k \hat{\mathbf{x}}_{k-1}, \quad (19)$$

with $\mathbf{V}_k$ defined as a Toeplitz matrix that implements the convolution operation with the motion model.

It is easy to show that the updated equation now becomes

$$\hat{\mathbf{x}}_{k+1} = \hat{\mathbf{x}}_k - \mu \mathbf{V}_k \mathbf{A}_k^T e_k. \quad (20)$$

In other words, the motion model propagates the error term to the expected location of the moving object so that its image is correctly updated as it moves across the scene.

We now present the results of initial experiments that demonstrate the reconstruction of moving objects using a motion model. This experiment was conducted using the SWIR camera which has a convenient “burst” mode that allows it to rapidly acquire images of moving objects. Recall that large pixel size (25μ) limits the accuracy of the mask registration so that only fourfold resolution enhancement is possible using the periodic mask. As before, the encoded measurements were made as the mask was stepped in increments of 5μ . Additionally however, a burst of 16 frames was also collected at each position while the mask was held stationary.

Figure 19(a) shows an example of a low-resolution coded frame with a vehicle on the road, captured in burst mode. We determined the apparent velocity of the vehicles moving from left to right on the bridge to be about 2 pixels between successive frames. Assuming constant linear velocity, the motion model is $V_k(m, n) = \delta(m, n + 2)$ for this simple case of vehicles moving from left to right along the bridge. It should be noted that this does not incorporate any blur uncertainty, and more general motion models for turns or other types of maneuvers need to be based on the tracker’s predicted velocity and uncertainty estimates. The result in Fig. 19(b), where the reconstructed image of a small truck is clearly visible, was obtained using this motion model in Eq. (20) and a burst of 16 frames. It is also interesting

Fig. 17 A system controller moves the mask for reconstructing static background information, or holds it stationary to measure information about moving objects. This is used to reconstruct objects under track that are moving consistently with motion models provided by the tracker.

Fig. 18 The motion model is the probability that the object may translate by $(\Delta m, \Delta n)$ pixels between consecutive observations

to note that the rest of the image is blank and the algorithm did not reconstruct the stationary background or those objects whose movement is inconsistent with the motion model (such as vehicles moving in the opposite direction). This shows the potential for selectively reconstructing moving objects using the motion model as a strong and discriminating task prior. The image in Fig. 19(c) is the simultaneous reconstruction of the background obtained as before using data acquired at different translated positions of the mask.

5 Summary and Future Work

In this paper, we have demonstrated that a translating mask can be successfully used to encode information and

Fig. 19 Sixteen low-resolution coded frames such as in (a) were captured in a burst while the mask was held stationary. The reconstruction of the moving vehicle in (b) was obtained using a linear constant velocity motion model. (c) The stationary background was reconstructed as before using the data acquired at different mask positions.

reconstruct images with considerably greater spatial resolution than the FPA used to acquire the data. Our goal is not only to demonstrate the principle, but also to describe the optomechanical issues that must be addressed for practical applications. In Sec. 3, we described two different reconstruction algorithms. A close form least-squares algorithm was discussed that reconstructs the image block by block. It should be noted that sparsity-based reconstruction techniques may be used instead of using the same measurements. We showed that the stability of the inverse problem depends not only on the mask pattern, but also on the displacement of the mask between measurements, and experimentally determined the best distances for mask movement to improve the condition number of the matrix to be inverted. In Sec. 4, we presented experimental results with outdoor scenes that demonstrated the ability to achieve an eightfold increase in image resolution in each dimension. As noted earlier, this implies that a 1 Mpixel FPA can be used to obtain 64 Mpixel high-resolution imagery, which can potentially alleviate the need for large format FPAs, particularly in application regimes where such devices are expensive. Finally, we presented preliminary results showing the detection, tracking and reconstruction of moving objects using motion models. We showed that while data acquired with a moving mask allow us to reconstruct a stationary scene, the mask can be held stationary to acquire information and reconstruct objects whose movement is consistent with particular motion models. This is an example of task-specific sensing, where information relevant for a particular task is measured and processed as part of the sensing process. In the context of a task-specific imaging and exploitation system, this holds the promise to dramatically reduce the volume of data that needs to be measured in today's traditional imaging systems.

On a comparative note, the proposed computational approach for reconstructing high-resolution images from encoded low-resolution measurements differs from conventional super-resolution techniques in at least one important way. While our method is designed to reconstruct images at the resolution of the finest mask element, it does not explicitly remove aliasing effects in the same way as digital super-resolution algorithms. It is therefore possible that algorithmic techniques can be further applied to the reconstructed image to remove aliasing artifacts, such as evident in Fig. 9(b). Further analysis is also required to quantify the residual noise in the reconstructions. We believe that the fixed-pattern noise and other reconstruction artifacts arise from residual errors in the mask alignment and calibration process, and that this can be further mitigated using a camera with pixel pitch of around 1 μm to record the mask images in Fig. 6. Effects of finite dynamic range and quantization effects could be significant²³ and will need to be carefully addressed in future studies. Finally, the extensions of this approach to infrared wavelengths must account for stray light and other thermal noise sources. Fortunately, the optics can be made telecentric at the mask so that it reflects back into dewar assembly, and therefore the mask does not have to be placed in the cold space. However, future applications to infrared sensing will require detailed analysis of the radiometric losses associated with the proposed architecture.

Acknowledgments

We gratefully acknowledge the support of the Office of Naval Research (ONR) and the MONITOR program under contract number N00014-14-C-0173. We would also like to thank Dr. Ravi Athale and Dr. Andrew Pipino for their constructive feedback and support during the course of this research effort.

References

1. S. Farsiu et al., "Advances and challenges in super-resolution," *Int. J. Imaging Syst. Technol.* **14**, 47–57 (2004).
2. M. Elad and A. Feuer, "Restoration of a single superresolution image from several blurred, noisy, and undersampled measured images," *IEEE Trans. Image Process.* **6**, 1646–1658 (1997).
3. R. C. Hardie et al., "High-resolution image reconstruction from a sequence of rotated and translated frames and its application to an infrared imaging system," *Opt. Eng.* **37**(1), 247–260 (1998).
4. R. C. Hardie and K. J. Barnard, "Fast super-resolution using an adaptive wiener filter with robustness to local motion," *Opt. Express* **20**, 21053–21073 (2012).
5. A. Stern et al., "Enhanced-resolution image restoration from a sequence of low-frequency vibrated images by use of convex projections," *Appl. Opt.* **40**(26), 4706–4715 (2001).
6. G. Carles, J. Downing, and A. R. Harvey, "Super-resolution imaging using a camera array," *Opt. Lett.* **39**, 1889–1892 (2014).
7. Y. Rivenson, A. Stern, and B. Javidi, "Single exposure super-resolution compressive imaging by double phase encoding," *Opt. Express* **18**(14), 15094–15103 (2010).
8. A. Ashok and M. A. Neifeld, "Point spread function engineering for iris recognition imaging system design," *Appl. Opt.* **49**, B26–B39 (2010).
9. M. J. Sun et al., "Improving the signal-to-noise ratio of single-pixel imaging using digital microscanning," *Opt. Express* **24**(10), 10476–10485 (2016).
10. S. J. Olivas et al., "Single pixel compressive imaging of laboratory and natural light scenes," in *Topical Meeting on Imaging and Applied Optics*, CTu1C.2, OSA (2013).
11. L. McMackin et al., "A high-resolution SWIR camera via compressed sensing," *Proc. SPIE* **8353**, 835303 (2012).
12. A. Mahalanobis et al., "Recent results of medium wave infrared compressive sensing," *Appl. Opt.* **53**(34), 8060–8070 (2014).
13. P. Lull et al., "Coded aperture compressive temporal imaging," *Opt. Express* **21**(9), 10526–10545 (2013).
14. M. E. Gehm et al., "Single-shot compressive spectral imaging with a dual-disperser architecture," *Opt. Express* **15**(21), 14013–14027 (2007).
15. R. Koller et al., "High spatio-temporal resolution video with compressed sensing," *Opt. Express* **23**, 15992–16007 (2015).
16. M. L. Don, C. Fu, and G. R. Arce, "Compressive imaging via a rotating coded aperture," *Appl. Opt.* **56**, B142–B153 (2017).
17. A. Mahalanobis et al., "Pixel resolution improvement using a sliding mask (PRISM)," in *Imaging Systems and Applications (ISA)*, Heidelberg, IW5F.4 (2016).
18. H. S. Pal and M. A. Neifeld, "Multispectral principal component imaging," *Opt. Express* **11**(18), 2118–2125 (2003).
19. E. J. Candes and Y. Plan, "Near ideal model selection by ℓ_1 minimization," *Ann. Stat.* **37**(5), 2145–2177 (2009).
20. A. N. Tikhonov and V. Y. Arsenin *Solution of Ill-Posed Problems*, Winston and Sons, Washington (1977).
21. A. Mahalanobis et al., "Pixel resolution improvement using a sliding mask," in *15th Workshop on Information Optics* (2016).
22. Z. Wang et al., "Image quality assessment: from error visibility to structural similarity," *IEEE Trans. Image Process.* **13**(4), 600–612 (2004).
23. A. Stern et al., "Quantization error and dynamic range considerations for compressive imaging systems design," *J. Opt. Soc. Am. A* **30**(6), 1069–1077 (2013).

Abhijit Mahalanobis received his PhD from Carnegie Mellon University in 1987 and currently is a senior fellow of Lockheed Martin Corporation in Orlando. His area of expertise in computational imaging, pattern recognition, and image processing, with an emphasis on the development of new technologies for sensors and machine vision systems. He has over 150 journal and conference publications in this area. Currently, he is a topical editor for *Applied Optics* and has previously served on the OSA board of meetings. He is a fellow of SPIE, OSA, and IEEE.

Richard Shilling received his PhD from Georgia Institute of Technology in 2009. Currently, he is a member of the engineering staff at Lockheed Martin. He has published several conference

and journal papers, and is actively involved in research and development of algorithms for computational sensing and imaging.

Robert Muise currently serves as a senior staff engineer in the Applied Research Department at Lockheed Martin, where he is worked since 1998. He obtained his PhD in mathematics at UCF in 2003. He holds several patents and has served as the principle investigator on several DARPA and ONR programs, as well as internally funded research and development efforts.

Mark Neifeld received his PhD from Caltech in 1991 and has since been a faculty member in the electrical and computer engineering and

optical sciences. He recently was a program manager at DARPA/DSO, where he started programs on quantum information, compressive sensing, and computational imaging. He is a fellow of OSA and SPIE and a senior member of IEEE. He has coauthored more than 120 journal articles and 250 conference papers in the general areas of optical physics and engineering. Further, he was a two-term topical editor, and a three-time guest editor of special issues, for *Applied Optics*. He has served on the organizing committees of numerous conferences and symposia.