

Optical Engineering

OpticalEngineering.SPIEDigitalLibrary.org

Performance evaluation of two optical architectures for task-specific compressive classification

Brian J. Redman
Amber L. Dagele
Meghan A. Galiardi
Charles F. LaCasse
Tu-Thach Quach
Gabriel C. Birch

SPIE.

Brian J. Redman, Amber L. Dagele, Meghan A. Galiardi, Charles F. LaCasse, Tu-Thach Quach, Gabriel C. Birch, "Performance evaluation of two optical architectures for task-specific compressive classification," *Opt. Eng.* **59**(5), 051404 (2020), doi: 10.1117/1.OE.59.5.051404

Performance evaluation of two optical architectures for task-specific compressive classification

Brian J. Redman,^{a,b,*} Amber L. Dagel,^a Meghan A. Galiardi,^a
Charles F. LaCasse,^a Tu-Thach Quach,^a and Gabriel C. Birch^a

^aSandia National Laboratories, Albuquerque, New Mexico, United States

^bUniversity of Arizona, Tucson, Arizona, United States

Abstract. Many optical systems are used for specific tasks such as classification. Of these systems, the majority are designed to maximize image quality for human observers. However, machine learning classification algorithms do not require the same data representation used by humans. We investigate the compressive optical systems optimized for a specific machine sensing task. Two compressive optical architectures are examined: an array of prisms and neutral density filters where each prism and neutral density filter pair realizes one datum from an optimized compressive sensing matrix, and another architecture using conventional optics to image the aperture onto the detector, a prism array to divide the aperture, and a pixelated attenuation mask in the intermediate image plane. We discuss the design, simulation, and trade-offs of these systems built for compressed classification of the Modified National Institute of Standards and Technology dataset. Both architectures achieve classification accuracies within 3% of the optimized sensing matrix for compression ranging from 98.85% to 99.87%. The performance of the systems with 98.85% compression were between an $F/2$ and $F/4$ imaging system in the presence of noise. © The Authors. Published by SPIE under a Creative Commons Attribution 4.0 Unported License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.OE.59.5.051404](https://doi.org/10.1117/1.OE.59.5.051404)]

Keywords: compressive sensing; compressive classification; task-specific; MNIST.

Paper 191357SS received Sep. 26, 2019; accepted for publication Dec. 18, 2019; published online Jan. 14, 2020.

1 Introduction

Classification of images is an active area of research for fields such as self-driving cars,^{1,2} facial recognition,³ medical imaging,^{4,5} and remote sensing.^{6,7} High-resolution data optimized for human observers is commonly passed into a machine learning algorithm that processes the data and returns a classification decision. Many of the images will never be seen by a person because the classification is the desired product and the volume of the data is large. High resolution images increase the volume of the data and can be unnecessary for the machine learning algorithms which reduce the dimensionality of the data as part of the processing. Compressive sensing implements some of this compression in the optical hardware, thereby reducing the size, weight, and power of the measurement system and reducing the bandwidth required for transmission of the data. In our past work, it was shown that optimizing a sensing matrix using a neural network to maximize classification accuracy instead of minimizing information loss enables better performance of a classification task than traditional compressed sensing methods.⁸

In this work we demonstrate two optical architectures to realize optimized compressive sensing matrices and expand upon past work.⁹ The first architecture achieves compressive measurements through the use of an array of prisms and neutral density filters in a nonimaging design. Each prism and filter pair enables the realization of one nonzero element within an optimized sensing matrix. The second architecture utilizes a more conventional approach, with a less complex prism array dividing the aperture into channels which are imaged onto an intermediate image plane. The sensing matrix weighting in the second architecture was achieved using a digital micromirror device (DMD) in the intermediate image plane. These two architectures

*Address all correspondence to Brian J. Redman, E-mail: bjredma@sandia.gov

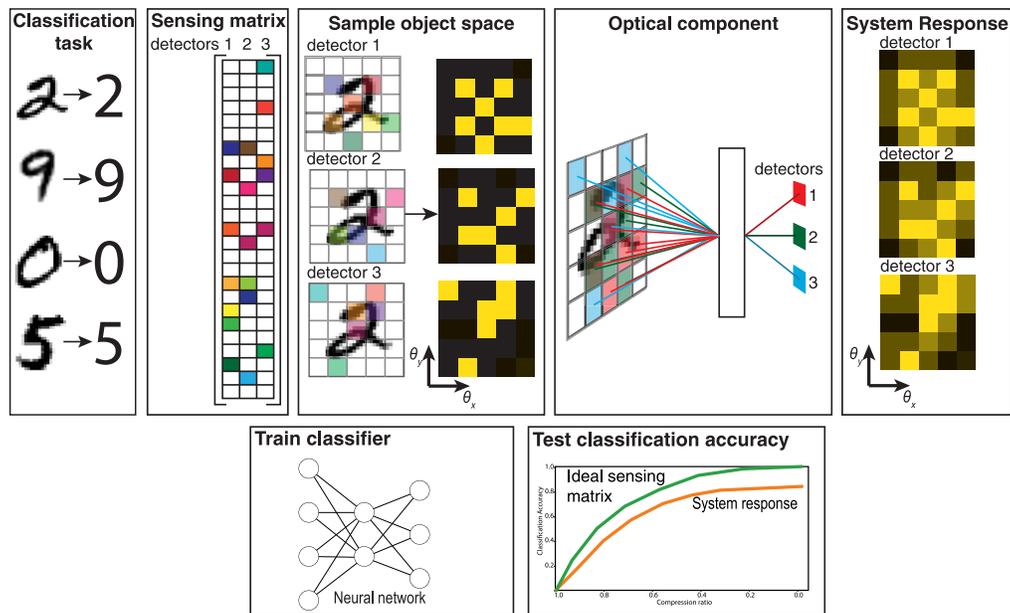


Fig. 1 For task-specific compressive classification, a sensing matrix is optimized to minimize the dimensionality of the measurement while maximizing the classification accuracy. If it is constrained to be non-negative, the sensing matrix becomes a weighted (transmission) mapping of the rows (input angles) to the columns (detectors). An optical component can be created to realize the weighted mapping. Then simulation or measurements of the optical components are used to create system response matrices that represent each detector's sensitivity to input angles. The performance of the system is determined by how well a machine learning classifier can classify scenes compressed using the system response matrix.

represent different approaches to realizing the same optimized sensing matrix. The positive and negative attributes of these approaches will be discussed throughout the work.

Figure 1 shows a high-level overview of realizing a sensing matrix as an optical component. The classification task chosen for this paper was classifying the handwritten digits of the Modified National Institute of Standards and Technology (MNIST) dataset. The generation of task-specific compressive sensing matrices was established previously⁸ and is discussed briefly in Sec. 2. Section 2 discusses the physical parameters of the system and sampling object space. Sections 2.1 and 2.2 discuss the creation of the optical components. Section 3.1 discusses the optical simulation of the system response. Section 3.2 discusses the radiometric properties of the compressive sensing systems as compared to imaging systems. Section 3.3 combines the optical simulations and radiometric model with a noise model to compare the compressive sensing systems to imaging systems. Section 4 presents the performance of the holistic optical device and algorithmic classifiers. The positive and negative attributes of each sensing architecture are discussed in Sec. 5. Finally, the conclusions are discussed in Sec. 6.

1.1 Background

Research in compressive sensing has shown that images can be reconstructed from datasets sampled below the Shannon–Nyquist sampling limit.¹⁰ The compression is achieved by transforming the data into a sparse basis, most commonly using the discrete cosine transform, and sampling with random Gaussian compression matrix, which is computationally easy to create and minimizes the loss of information for natural scenes.¹¹ Optimized sparsifying transforms and compression matrices have been explored to minimize the loss caused by the sensing matrix.^{12–15} Being able to reconstruct images from a compressed measurement dataset has been demonstrated in simulation of multiple architectures¹⁶ and experiments such as the Rice University single-pixel compressive imager.^{17,18}

It has been shown that if the compression follows the restricted isometry property (RIP), which is a common optimization metric for compressive imaging systems, then directly

classifying compressed data does not reduce classification accuracy.¹⁹ Even for compression ratios large enough to break the RIP, classifiers have been developed to classify the data. Convolutional neural network classifiers have been applied to MNIST images compressed using random Gaussian projections. The classification accuracies were greater or equal to 58.94% for compression ratios less than or equal to 98.98%.²⁰ Classification of compressed data has been shown to maintain high classification accuracies across a wide range of datasets and compression techniques.²¹

Removing the requirement to reconstruct an image separates compressive classification from compressive imaging. The requirements of the sensing matrix also change for compressive classification. Compression that maximizes task-specific information and removes information not relevant to the task can improve the performance of detection or localization.²² This concept has been demonstrated for systems using sequential measurements on a single detector and a changing spatial light modulator or parallel measurements using a lenslet array and fixed masks.²³ The improved performance is because the dimensionality of the space is the number of measurements, and adding dimensions that do not differentiate the classes increases the computational cost²⁴ and can reduce classification performance.²⁵ The compression matrix needs to be optimized to maximize the classification accuracy, but for compressive sensing, the sensing matrix also has to be optimized to be realizable with optical hardware. Negative values are not directly realizable because the irradiance on a detector is inherently positive. In addition, the number of elements in the matrix sets the complexity of the optical design. We developed sensing matrices optimized using neural networks to maximize classification accuracy of the MNIST dataset and constrained to be non-negative and sparse.⁸

In this work we present a comparison between a DMD architecture and a prism array architecture that directly realize a compressed, task-specific sensing matrix. The behavior of both architectures is simulated using nonsequential ray tracing to consider factors such as stray light. The ray trace simulation is combined with radiometric models of the systems to allow for the comparisons of noise between the compressive classification systems and traditional imaging. High classification accuracies are demonstrated for the simulated systems at extremely high compression ratios.

2 Optical Designs

We developed two optical architectures to realize optimized sensing matrices, as established by Birch et al.⁸ for the MNIST task. These optical designs highly compress the data from the 784 pixels in the images to between one and nine measurements. The small number of detector elements reduces the constraints on physical placement, detector colocation, and pixel size.

Both the images of the MNIST dataset and the sensing matrix are mathematical constructs which have to be translated to physical parameters to create an optical system. The sensing matrix can be considered as a mapping of brightness in object space to detector values in image space. Each row of the sensing matrix is one object space location, and the column determines the detector. For a non-negative sensing matrix, the whole matrix can be constrained to be between zero and one by dividing by the maximum value in the matrix. Weights of this normalized sensing matrix correspond to the transmission from each input angle. There are 28 by 28 input angles because the images in the MNIST dataset are 28 by 28 pixels.

The sensing matrix was created using a neural network optimization. The first layer of the neural network was the compression matrix, enabling the neural network to learn the form of the compression matrix concurrently with learning a classifier to maximize the classification accuracy. The compression matrix was constrained to be non-negative so that the weights could be realized using transmission of an optical component. Basis pursuit²⁶ was used to sparsify the sensing matrices to decrease the complexity of the optical components. The final layer of the neural network was a Softmax classifier.

The design problem was inherently underdefined because the images in the MNIST dataset do not have physical properties such as size or radiant exitance. To make the problem tractable, we defined the source and detector geometry. We set the distance to the scene as infinite; therefore, the light from each pixel was a collimated source or a plane wave. With the object at infinity, the size was defined in angular space. We assumed a half field of view (HFOV) of

5 deg, corresponding to a fairly narrow field of view system. The size of the detectors was set to be 100 by 100 μm , which is much larger than the pixels of a consumer camera. The small number of separated detectors allowed for the larger size.

2.1 Prism Array Architecture

The sensing matrix maps values in object space to measurements in image space. Lenses cannot directly realize this mapping because lenses have a one-to-one mapping from input angle to output location, while each column of the sensing matrix has multiple separated nonzero values. Mapping multiple input angles to a single detector encourages the use of an array of elements. A prism is an element that maps an input angle to an output angle. If the position of the prism and detector are known, a prism can be used as an element that maps an input angle to an output location. The transmission of the prism corresponds to the weighting of the sensing matrix. In this paper, we discuss a process to design a prism array to realize an arbitrary sensing matrix. Figure 2 shows the physical dimensions assigned to the system to make the underdefined problem of realizing mathematical constructs tenable. Figure 3 shows an overview of the processing workflow to convert a sensing matrix into a physical component.

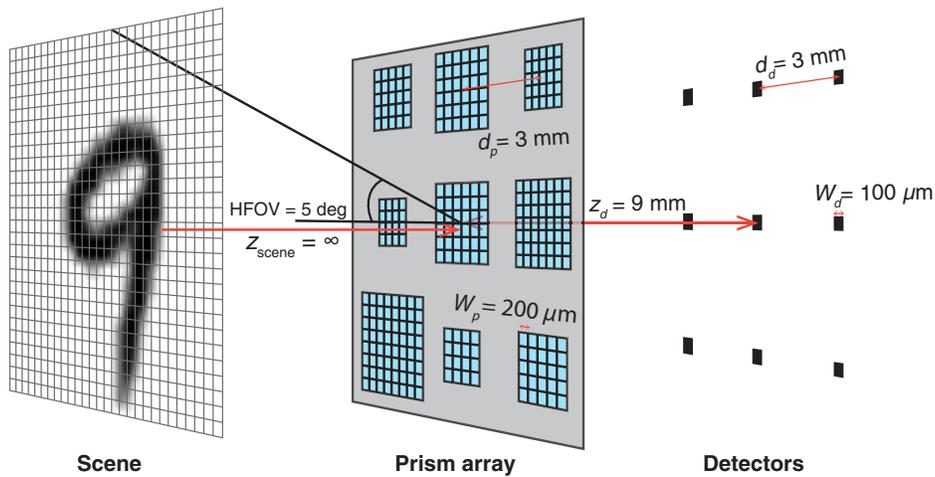


Fig. 2 The dimensions of the prism array. The diagram shows a nine-detector example.

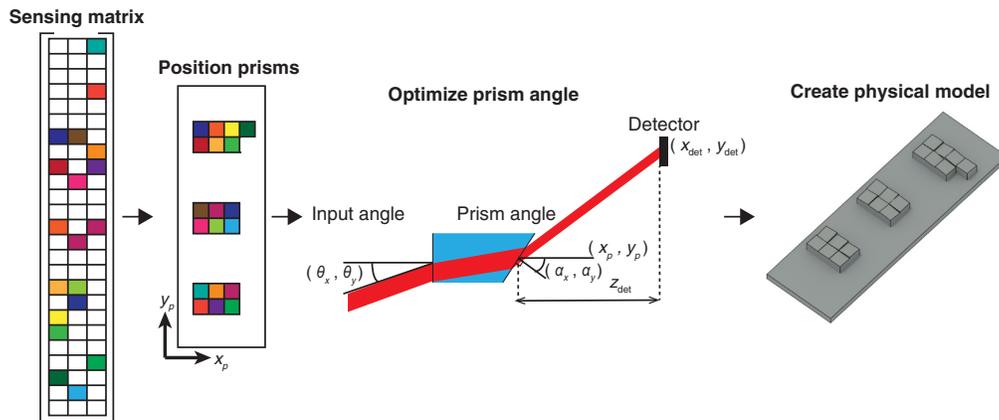


Fig. 3 The prism array was created as a physical realization of the sensing matrix. First, each nonzero entry from the sensing matrix was assigned to a physical location. Then the angles of each prism were optimized to map the input angle, determined by the location in the sensing matrix, to the detector position. A model was created from the prism geometry to allow for simulation of the optical system.

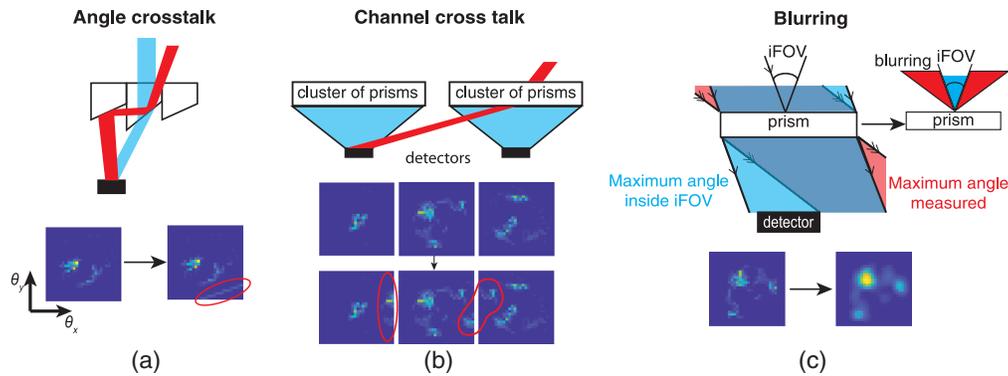


Fig. 4 The considerations when designing a compressive classification system are different from the traditional imaging systems. The error cases come from angles being mapped incorrectly onto the detectors. (a) Angle cross talk was caused by scattering or spurious reflections. (b) Channel cross talk was caused when the prisms from one channel contributed light to the detector of an adjacent channel. (c) Blurring was caused by a detector accepting a larger angular field of view than was designed for.

The positions of the detectors were set before the prism positions because the prisms were clustered around the detector to which they contributed. The clusters of prisms avoided having lines of prisms with similar angles, which could cause angle cross talk, as shown in Fig. 4(a). The separation between the detectors, and the distance between the prism array and the detector determined the stray light between the prism clusters. This channel cross talk, as shown in Fig. 4(b), could be decreased by widening the separation between the detectors; however, a larger separation increased the total size of the component.

Decreasing the distance between the prism array and the detector decreased the required separation of the detectors, but increased the angle between the prisms and the detector. For this work, the separation between the detectors was set to 3 mm. The distance between the prism array and the detector was set to 9 mm.

With the detector positions set, each nonzero element in the sensing matrix was assigned a position on a grid centered on the corresponding detector. The grid spacing was determined by the size of the prisms. The size of the prism determined if the detector was underfilled or overfilled over the range of angles accepted by each prism. For this design, the prism size was set to 200 by 200 μm which overfills the detector for the designed field of view of each prism. The large prisms also increased the power on the detector at the expense of blurring the system response, as shown in Fig. 4(c). The height of the prisms was set so that the lowest point of the prisms was touching the substrate.

The position of the prisms relative to the detector was used by a sequential ray trace program to optimize the angle of each prism. The refractive index of the prism material was assumed to be 1.5 for all wavelengths. The physical parameters of the prism array were used to create a physical model in a nonsequential ray tracing program where each prism was modeled as a separate rectangular solid. The weighting of the sensing matrix was implemented with coatings on the tilted surface of each prism. Transmissions were set to the weight of the normalized sensing matrix, and the reflectance was set to a uniform 6%. The coating was an approximation of an absorptive neutral density filters where the reflection would be due to the glass filter interface, and the attenuation would be due to absorption inside the filter. It was not feasible to implement the transmissions with floating point precision, so the transmissions were uniformly binned into 128 values between 0 and 1.

2.2 Digital Micromirror Device Architecture

The prism array architecture requires many elements with small feature sizes and sharp edges which is not feasible to fabricate using traditional methods. An alternative architecture using a DMD and less complex prism array enabled the use of commercial off-the-shelf optical elements and a simple custom optical component. The DMD architecture presented in this work takes

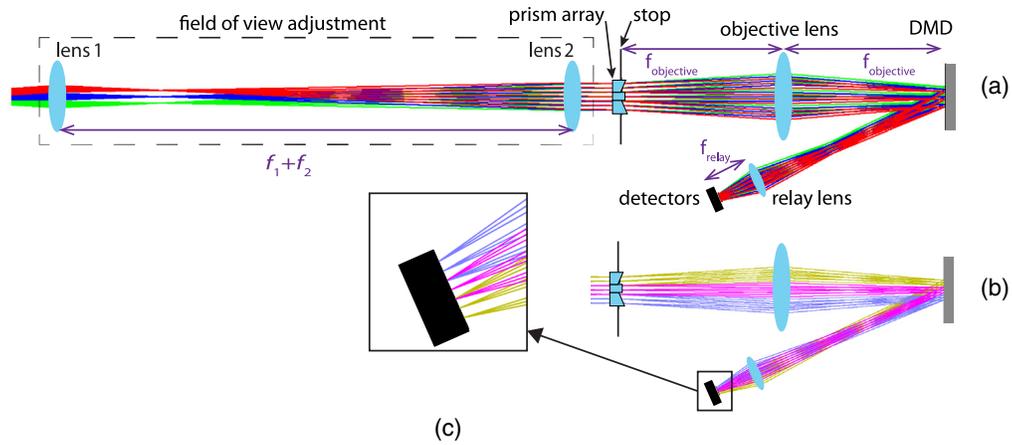


Fig. 5 (a) The DMD architecture uses a prism array to divide the stop. The object is imaged onto the DMD and the stop is imaged onto the detector. (b) Both the channels and the fields are separated at the DMD. (c) The fields are combined to be overlapping at the detector, but the channels are separated.

measurements simultaneously with parallel channels as opposed to the sequential measurements made by other architectures such as the Rice University single-pixel compressive sensor.²⁷

To make parallel measurements required a channel for each detector which was spatially resolved at the DMD plane, but uniform irradiance at the detector plane. The spatially resolved and separated channels allowed for the sensing matrix weighting to be realized by varying the duty cycle of the DMD pixels. The angular information had to be recombined into uniform irradiance at the detector to avoid preferential sensitivity to some angles.

We achieved spatially resolved channels on the DMD plane using an objective lens to image the scene onto the DMD. The uniform irradiance on the detector plane was achieved by using a relay lens to image the aperture onto the detector, as shown in Fig. 5. The aperture was imaged onto the detector because the irradiance in the aperture is approximately uniform. The stop was located at the front focal point of the objective lens making the system telecentric,²⁸ allowing the distance between the DMD and the relay lens to be changed without any change of magnification.

The separate channels were formed using a prism array to divide the aperture. The angle of the prism set the separation between the channels in the intermediate image plane. Only one prism was required for each detector, and the size of the prisms was larger than the prism array architecture, allowing them to be fabricated as separate components using polishing techniques.

The parameters were optimized using paraxial optics as a proof of concept for the system, but each component was chosen to be possible to implement with commercial off-the-shelf components. The first component set was the DMD, because the size of the DMD sets the requirements of the remaining elements in the system. As a starting point, the design was created around the Texas Instruments® DLP® LightCrafter™ 6500, which is a commercially available component with a large active area of 14.52 by 8.16 mm. There are 1920 by 1080 micromirrors across the active area. Therefore, the resolution of the DMD is much higher than the 28 by 28 pixels required for each channel. The channels on the DMD were arranged in a 3 by 3 grid, so the maximum size of each channel at the DMD plane was 2.72 by 2.72 mm. The light incident on the objective lens was collimated and imaged onto the DMD one focal length behind the objective lens. Therefore, the width of each channel at the DMD was

$$w_{\text{channel}} = 2 \tan(\text{HFOV})f_{\text{objective}}, \quad (1)$$

where HFOV is the half field of view and $f_{\text{objective}}$ is the focal length of the objective lens. Using the 5-deg HFOV design constraint required the objective focal length to be shorter than 13.0 mm to avoid overfilling the DMD. The short focal length was not practical because the lens needed to

be far enough away from the DMD that the reflection would not clip on the mounting hardware for the lens. In addition, the size of the aperture imaged onto the detector plane was set by the ratio of the objective focal length to the relay focal length, and it was not reasonable to make the focal length of the relay lens significantly shorter than 13 mm.

A Keplerian telescope was added to the front of the system to maintain the 5-deg HFOV in object space and decrease the HFOV received by the objective lens. An HFOV magnification of 4 was chosen, so the HFOV received by the objective lens was 1.25 deg. With the reduced HFOV, the objective focal length needed to be less than 52.1 mm. Setting the focal length to 50 mm allowed for many off-the-shelf solutions and reduced the size of the channels at the DMD plane. The small unused region around each channel reduced the likelihood of channel cross talk. With the focal length of the objective set, the angle of the prism was optimized to -5.85 deg for a 2.58-mm separation between the channels at the DMD. The separation smaller than 2.72 mm increased the unused pixels around the outside edge.

The focal length of the relay lens was set to 5 mm to give a 10 times magnification of the detector area. The magnification increased the effective detector area at the aperture and therefore increased the throughput. The focal length of the relay lens had to be balanced against the speed of the lens. The entrance pupil of the lens needed to be larger than the beam diameter, and the focal length needed to be longer than the entrance pupil diameter for additional off-the-shelf options. In addition, increasing the effective area of the detector required a larger prism size or the channels would underfill the detector. Increasing the prism size also increased the beam diameter.

For this work, the prisms were set to 4 by 4 mm which required a relay lens entrance pupil diameter of 11.7 mm, requiring an $F/0.43$ lens which is not feasible. However, if the size of the prisms was reduced to 1 by 1 mm (the effective size of the detector at the prism plane), the beam diameter at the relay lens was 3.59 mm requiring an $F/1.39$ lens, which was a commercially off-the-shelf option.

The beam for each channel was narrowest at the DMD. After that, each beam expanded but the centers of the beams crossed. The centers of the beams converging creates a point where the total beam diameter was narrowest. The location and diameter of this beam waist were determined by a combination of the field of view, prism angle, and objective focal length. In this design, the relay lens was located farther from the DMD than the minimum beam diameter because the minimum beam diameter was close to the objective lens, as shown in Fig. 5(b). The close proximity of the two lenses would not allow for mounting hardware.

The relay lens imaged the aperture onto the detectors. The detectors for this design were located in a 3 by 3 grid at the rear focal length of the relay lens with no separation between the detectors. For the measurements using fewer than nine detectors, the unused sections of the DMD would be set to no transmission. Therefore, no hardware changes were required for any of the measurements presented in this work.

3 Analysis

A ray trace simulation was used to analyze errors in the system which changed the system response matrix. The system response matrix compressed the images of the MNIST data, and the performance of the system was measured from the accuracy with which a classifier could determine the digits. A radiometric model was used to determine the power transmitted to the detector for both the prism array and DMD task-specific architecture. The system response was combined with the radiometric model and noise was added to the detectors to determine how the performance of the compressive sensing systems would compare to traditional imaging systems.

3.1 Nonsequential Ray Tracing

It was not feasible to analytically calculate the scattering and stray light across many element designs presented in this work. To evaluate the performance of the holistic compressive architectures performing an MNIST classification task, a nonsequential ray trace was performed using Zemax OpticStudio® (ZOS).

A source rectangle was added before the first surface of each design. The source created 10^6 randomly positioned collimated rays simulating the input from one object space location. The position and size of the source were set to fill the aperture of each system. To build up the system response matrix, flux on each detector was recorded for each input angle. For the MNIST dataset, there are 28 by 28 input angles, requiring 784 total ray traces.

An automatic process to build up the system response was created. The input angle was scanned using Python scripts controlling ZOS through the application programming interface. Flux on each detector for each angle was saved into a matrix. A system response matrix was created from the detector measurements by dividing the matrix by the maximum value, normalizing the matrix between 0 and 1.

Performance of the compressive optical systems was simulated using the system response matrices. Images of the MNIST dataset were compressed by multiplying the system response matrices. A random forest²⁹ ensemble classifier was trained on compressed data from all 60,000 preassigned training images and tested on compressed data from all 10,000 test images. Python's scikit-learn module, version 0.20.2, implementation of random forest classifier was used. The number of trees in the forest was set to 100. Split quality was measured using entropy, and each split had to contain at least two data points. The maximum number of features that each tree considered was $\text{floor}[\ln(\text{number of detectors}) + 1]$, where \ln is the natural logarithm.

3.2 Radiometric Model

Radiometric throughput is used to calculate the signal power on the detector for a given source radiance and enables analysis of the proposed compressive optical systems' expected sensitivity requirements as compared to traditional imaging systems. The flux Φ on a detector is

$$\Phi = GL, \tag{2}$$

where G is the throughput of the detector and L is the radiance of the source.

In a traditional imaging device, the throughput can be defined by the area, A , of a pixel and the projected solid angle, Ω , subtended by the exit pupil,

$$G_{\text{imaging}} = A\Omega = w_d^2\pi \sin^2(\theta), \tag{3}$$

where θ is the half-angle subtended by the exit pupil and w_d is the width of a pixel.³⁰ Assuming that the exit pupil is the same size as the entrance pupil sets the angle as

$$\theta = \tan^{-1}\left(\frac{1}{2F/\#}\right), \tag{4}$$

where $F/\#$ is the f -number of the lens.²⁸ The throughput is approximately constant for all detector elements across the narrow field of view of the system.

The throughput of the system can be extended to a radiometric system response matrix by calculating the sensitivity of each detector from all the input angles. The imaging systems were assumed to be ideal, perfectly mapping each angle from object space onto a single detector. Therefore, the imaging system response matrices were 784 by 784 element identity matrices. The imaging radiometric system response matrix Ξ is given as

$$\Xi_{\text{imaging}} = G_{\text{imaging}}I, \tag{5}$$

where I is an identity matrix with height and width equal to the number of angles in the scene.

Compressive sensing systems do not have a one-to-one mapping between object space and detector values, which makes calculating the throughput less straightforward. Each detector has contributions from multiple nonclustered input angles. The throughput is a summation of the solid angle for all the contributing input angles times the aperture area that limits the size of each beam. The contribution of each beam is also multiplied by the transmission defined by the system response matrix. The generalized throughput of the compressive sensing systems is given as

$$G_{\text{compressive}}[k] = \sum_{m=1}^M A_{\text{apt}}[m, k] \Omega[m] \Theta[m, k], \quad (6)$$

where m is the angle number, k is the detector number, M is the number of pixels in the scene, A_{apt} is the area of the aperture for each input beam, $\Omega[m]$ is the projected solid angle subtended by each pixel in the scene, and $\Theta[m, k]$ is the system response matrix weight for each input angle and detector. This throughput can be calculated at any surface in the optical system; however, these calculations can be complex. To simplify this, we calculate the throughput at the aperture where the irradiance is constant. In both of the compressive classification architectures, this aperture surface is at the prism array.

The aperture area of each beam is the area of detector that is illuminated. The throughput calculation is in the aperture plane, so the detector area is projected to the aperture plane. The aperture area will be the minimum between the area of the prism and the area of the detector when it is projected onto the prism plane,

$$A_{\text{apt}}[m, k] = \min(A_{\text{prism}}, A'_{\text{detector}}), \quad (7)$$

where A_{prism} is the area of the prism and A'_{detector} is the effective area of the detector projected onto the prism plane. Here A_{prism} is always larger than A'_{detector} for the case where the detector is overfilled across the instantaneous field of view (iFOV) as is the case for the compressive sensing systems presented here. Therefore, the area for the throughput calculation is the area of the detector times the magnification. The aperture area becomes independent of the detector because all the detectors are of the same size and also independent of the angle because all angles are contributing to the detectors of the same size.

The projected solid angle in the throughput calculation is the iFOV for each input angle times a cosine projection term. In this work, the iFOV was defined as the solid angle subtended by one of the 28 by 28 pixels in object space. For the DMD architecture, the HFOV was multiplied by the field of view magnification from the telescope. The projected solid angle was approximated as the differential solid angle times the cosine of the center angle of the pixel. The approximation was valid because the iFOV is small and the cosine of the largest input angle is approximately 1. A differential solid angle can be calculated by the differential area that subtends it:³⁰

$$d\omega = \frac{dA}{r^2}, \quad (8)$$

where $d\omega$ is the differential solid angle, dA is the differential area, and r is the distance to the area. From this equation, we set the distance to 1 and assume a rectangle defined by the pixel in angular space,

$$\text{iFOV} \approx d\omega \approx 4 \tan\left(\frac{\theta_{\text{iFOV}}}{2}\right)^2, \quad (9)$$

where θ_{iFOV} is the angle subtended by one of the pixels. Combining the area, solid angle, and weighting from the normalized system response matrix gives the radiometric system response matrix for the compressive sensing systems as

$$\Xi_{\text{compressive}}[m, k] = A'_{\text{detector}} \text{iFOV} \cos[\theta(m)] \Theta[m, k], \quad (10)$$

where $\theta(m)$ is the center angle of each pixel in the scene. Summing the radiometric system response matrix across all input angles gives the throughput of the compressive sensing systems as

$$G_{\text{compressive}}[k] = \sum_{m=1}^{784} \Xi_{\text{compressive}}[m, k] \quad (11)$$

3.3 Noise Analysis

The system response matrices give the ideal performance of the system, and the radiometric model determines how much signal will reach the detectors. We combined the two models to estimate how the performance of the compressive sensing systems compares to the imaging systems in the presence of noise. The system model was

$$\mathbf{V} = \mathbf{L}\Xi\rho + \mathbf{N}, \quad (12)$$

where \mathbf{V} is the voltage measured by the detectors, \mathbf{L} is the radiance of the scene, Ξ is the radiometric system response matrix, ρ is the responsivity of the detectors, and \mathbf{N} is a vector containing the noise for each detector.

The underdefined nature of measuring the MNIST dataset made absolute noise models from source to detector difficult to define. Instead, the relative performance of the architectures was considered for this work. The radiometric system response for each system were normalized by the throughput of the $F/2$ imaging system giving a normalized radiometric sensing matrix as

$$\hat{\Xi} = \frac{\Xi}{G_{F/2}}. \quad (13)$$

The scene term was system-independent, and the responsivity of the detectors was assumed to be the same for all the systems. This term was multiplied by the throughput of the $F/2$ imaging system to remove the units. The unitless object vector $\hat{\mathbf{x}}$ became

$$\hat{\mathbf{x}} = G_{F/2}\rho\mathbf{L}. \quad (14)$$

The absolute amplitude of the object vector was not needed for a relative comparison between the systems, so the images of the MNIST dataset were normalized to have values between 0 and 1.

The noise was assumed to have a Gaussian probability distribution with zero mean and variance σ . Signal-independent noise ignores the Shot noise of the system and assumes that the noise comes from thermal processes in the electronics. The complete relative system model was

$$\hat{\mathbf{y}} = \hat{\mathbf{x}}\hat{\Xi} + N(0, \sigma), \quad (15)$$

where $\hat{\mathbf{y}}$ are the measured values passed to a random forest classifier and $N(0, \sigma)$ are the Gaussian distributed noise values with zero mean and standard deviation σ .

Signal-to-noise ratio (SNR) is a good metric to compare the performance of the systems; however, the SNR of each system depends on the signal at the detector which is dependent on the throughput of the system. We defined a peak SNR (pSNR) relative to the $F/2$ imaging system. The peak value of the scene is always 1, so the pSNR of the $F/2$ imaging system is

$$\text{pSNR}_{F/2} = \frac{1}{\sigma}. \quad (16)$$

4 Results and Discussion

The system response matrices of each system were created using the ray trace simulation procedures described in Sec. 3.1. Example system response matrices for the case of nine detectors are shown in Fig. 6. For a perfect optical system, the system response matrices would exactly match the sensing matrix that they were designed from (i.e., the ideal sensing matrix shown in Fig. 6(a) would be perfectly reproduced using the optical hardware). However, this was not the case. The prism array architecture had a blurred system response matrix compared to the sensing matrix as shown in Fig. 6(b). This blurring was caused by the prism accepting angles larger than the iFOV. The DMD optical system closely reproduced the original system response matrix. However, there are rows and columns where the response was zero (e.g., a column of data

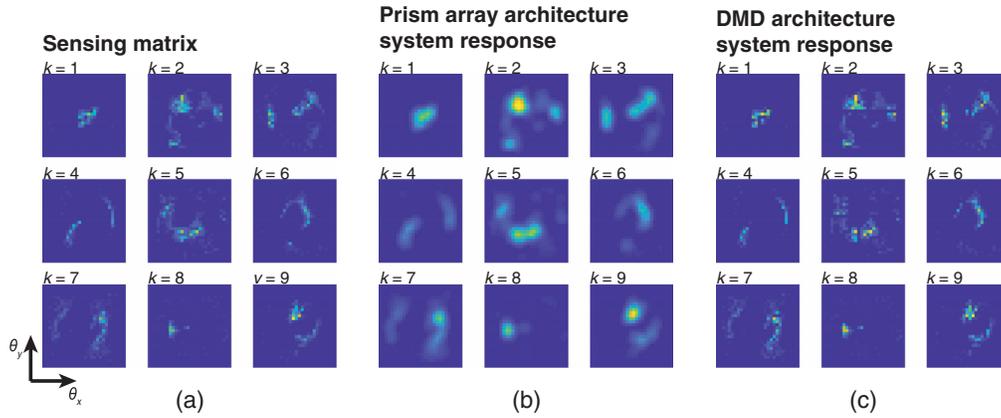


Fig. 6 Side-by-side comparisons between the system response matrices for the nine detector configuration for the (a) ideal system response matrix, for the (b) prism array architecture, and for the (c) DMD architecture.

is lost in Fig. 5(c), $k = 5$ detector response). This zero response results from regions where input angles focused directly on the edge of a DMD mirror in the simulation. The error would likely be removed for a real system with an iFOV instead of perfectly collimated light. However, the error indicates that scattering will be a possible problem.

The system response matrices from these ray trace simulations were used to compress the 60,000 training images from the MNIST dataset. A random forest classifier²⁹ was trained on this compressed data. The trained classifier was then used to classify a compressed test dataset consisting of 10,000 images, and the classification accuracy was recorded. The accuracy of the compressive sensing systems was compared to the classification accuracy of the ideal sensing matrix.

Figure 7 shows the classification accuracy for task-specific classification systems designed with one to nine detectors. Error bars are set by the standard deviation from 10 training classification cycles. Both optical systems had similar performance to the ideal sensing matrix, converging to over 90% accuracy using nine detectors [Fig. 7(a)], and less than 3% difference from the ideal sensing matrix across all configurations [Fig. 7(b)]. The performance difference between the ideal sensing matrix and system response matrices [Fig. 7(b)] showed the interesting result that the prism array had better performance than the ideal sensing matrix when a five-, eight-, or nine-detector configuration was used. We expect that the improved performance was

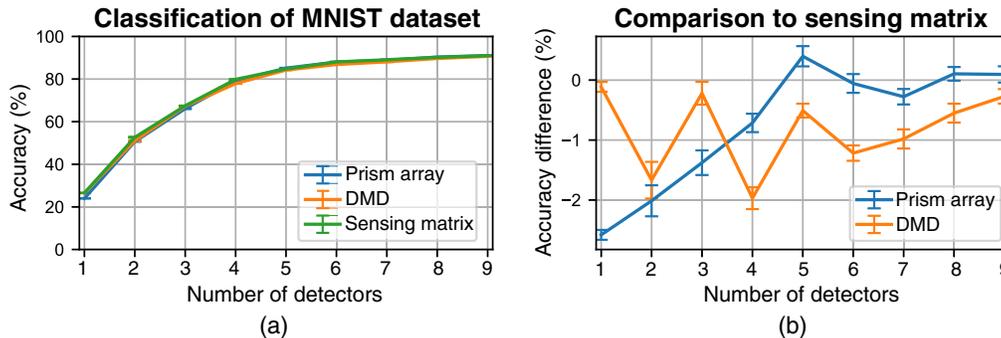


Fig. 7 (a) The classification accuracy of the prism array architecture and the DMD architecture had similar performance to the ideal sensing matrix as the number of detectors was increased. The accuracy of both architectures was close to the classification accuracy of the sensing matrix. The difference of the classification accuracy for each architecture relative to the sensing matrix was plotted separately to show detail and is provided in (b). (b) The DMD architecture was slightly worse than the ideal sensing matrix for all the number of detectors. The prism array architecture had performance exceeding the ideal sensing matrix for five, eight, and nine detectors because the blurring relaxed the sparsity.

due to the blurring relaxing the sparsity constraint of the optimization. The prism array had significant blurring as shown in Fig. 6(b). The blurring locally released some of the sparsity constraint, allowing for improved performance without increasing the number of prism elements. This indicates that classical design techniques may not be the best method for optimizing the compressive sensing systems. The below ideal performance of the DMD optical design was likely due to the dead rows and columns caused by an input angle being focused onto the edge of a mirror in the DMD, and the system performance would likely be improved by simulating the full iFOV instead of a collimated source.

The radiometric models set a comparison between the compressive sensing systems and imaging systems. Table 1 shows the calculated throughput for both ideal imaging systems and the nine-detector case for both compressive sensing architectures. For both cases, the source was assumed to fill the field of view of the systems. The compressive sensing systems had higher throughput than an imaging system with an $F/2$ lens and $5 \mu\text{m}$ active area of pixels which would be a high-throughput consumer camera.

Adding noise to the detectors of the systems decreased the classification accuracy, as shown in Fig. 8. The $F/2$ imaging system had the highest classification accuracy for SNR greater than 2.3 which was expected because the compressive sensing systems were using 98.85% fewer detector elements than the imaging systems. The high throughput of the compressive sensing systems resulted in both the prism array and DMD systems with more than five detectors having higher classification accuracies than the $F/4$ imaging system when the SNR was less than 8. The DMD architecture was insensitive to noise due to the high throughput of the architecture.

Table 1 The throughput of the two imaging systems versus the throughput of the compressive sensing systems. The imaging systems are assumed to be ideal and calculated for the on-axis pixel. The throughput is unique for each detector of the compressive sensing systems. The values shown are for the nine-detector configuration of each system. The order of the detectors for the compressive classification systems correspond to the system response matrices shown in Fig. 6.

System	Throughput ($\mu\text{m}^2\text{sr}$)								
Imaging systems									
For each pixel									
$F/2$ lens $5 \mu\text{m}$ pixel	4.62								
$F/4$ lens $5 \mu\text{m}$ pixel	1.21								
Compressive systems									
	Detectors								
Prism array	8.94	22.65	17.01	7.22	20.36	8.86	14.41	5.60	12.91
DMD	20.67	55.23	42.89	15.56	46.89	22.75	38.62	15.46	32.76

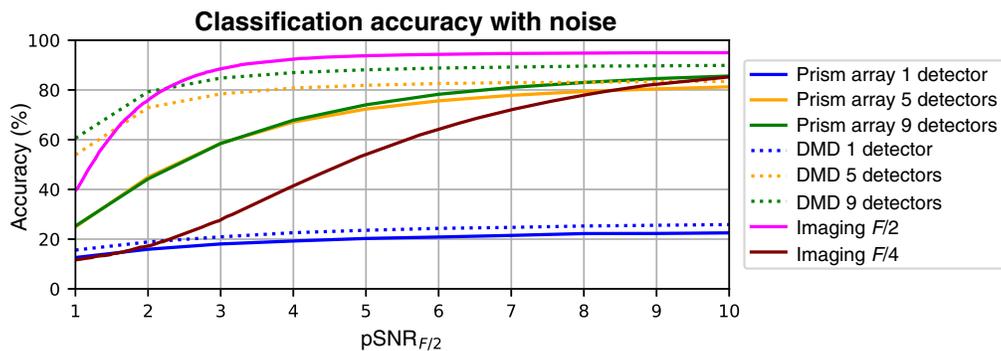


Fig. 8 The noise was random Gaussian added to the detector measurements. The standard deviation was set so that the bright regions of the MNIST images imaged by the $F/2$ system had the displayed SNR.

5 Comparison

The two compressive sensing architectures described in this paper are both realizations of the same sensing matrices; however, the architectures used to realize these sensing matrices have different strengths and weaknesses. The prism array architecture is a monolith element that requires specialized fabrication. The DMD architecture was much larger and required more power, but the optics could be realized using off-the-shelf components and an easily fabricated prism array.

The construction of the two architectures have different challenges. The fabrication of the prism array requires specialized tools such as additive manufacturing. However, assembly of the system is trivial after the part is fabricated. The prism array has to be aligned to the detector elements. The overfill of the detectors allows for some misalignment. The only custom components required for the DMD architecture are the nine prisms used to split the channels. All the other components can be realized with commercial off-the-shelf components, and the sensing matrix can be changed without changing the hardware. Assembling the system is more involved because it requires aligning seven elements.

The form factor of the two systems is also a consideration. The prism array is designed to mount 9 mm in front of the detectors. The only external components needed would be to control the field of view. The DMD architecture is much larger because it requires the vertical space for the optics to adjust the field of view and to image the scene onto the DMD and the horizontal space for the relay lens and the detectors. Optimization can reduce the footprint, but the system will not reduce to the size of the prism array architecture.

From a throughput standpoint, the prism array makes more efficient use of the aperture area, but there is no magnification and so the aperture area is limited by the prism area or the detector area, whichever is smaller. The DMD architecture can be optimized to increase the magnification, thereby increasing the effective area of detector at the aperture plane. The higher throughput of the DMD architecture makes it more resistant to noise.

For errors in the sensing matrix, the DMD is capable of more faithfully reproducing the sensing matrix. The objective lens has a one-to-one mapping of input angles, so no angle cross talk was seen or is expected. If the spacing between the channels on the DMD is too narrow, then the overlapping channels will cause significant channel cross talk. The prism array architecture blurs the sensing matrix. However, the analysis of system performance indicated that the blurring may increase classification accuracy. The increased performance is likely due to the blurring relaxing the sparsity constraint.

The high throughput and adaptability of the DMD architecture made it ideal for a generalized system. The small form factor without any moving parts of the prism array made it ideal for tasks that require low size, weight, and power.

6 Conclusion

Machine learning algorithms have become a mainstay of image analysis, and with the increased use of machine learning, large datasets of labeled data have become commonly available. Compressive sensing systems can make use of these datasets to minimize the data being collected and release the constraints of imaging. This research has shown a workflow for creating optical components to realize a compressive sensing matrix using the basic example of classifying the MNIST dataset. The workflow directly scales to more complicated sensing matrices, only limited by the print area for the prism array and the DMD size and resolution for the DMD architecture. This workflow enables the rapid creation of compressive sensing systems for arbitrary datasets.

We have presented two optical architectures for the creation of task-specific compressive imagers. The first architecture was a monolithic part that used a prism array to directly implement the sensing matrix. The second architecture used a simple prism array, conventional optics, and a DMD to implement the sensing matrix.

The sensing matrix of each architecture was simulated using the nonsequential ray tracing for configurations of detectors ranging from one to nine. The prism array was shown to blur the sensing matrix, where the DMD architecture was shown to reproduce the sensing matrix with

much greater fidelity. However, the classification performance of the two systems was shown to be similar despite the blurring.

The radiometric throughput of both systems was found to be greater than an $F/2$ lens imaging onto $5\text{-}\mu\text{m}$ pixels. The noise analysis of the systems showed that the $F/2$ imaging system had better performance than the compressive classification systems when the SNR was >2 . However, the high throughput of the compressive sensing systems resulted in higher classification accuracy than an $F/4$ imaging system when the SNR was <8 . These results indicate that the compressive classification systems could have similar performance, or better in low SNR condition, to imaging systems while requiring 98.85% fewer detector elements.

The discrete detectors of the demonstrated architectures open the door to using unique detectors for each channel increasing the spectral bandwidth or combining extremely high speed measurements for temporal Fourier analysis along with spatial measurements for localization.

The prism array opens up the possibility of a solid-state optical component with the space between the detectors and the prisms filled with an optical epoxy, making the entire system extremely compact and vibration resistant. The DMD architecture enables a reconfigurable compressive sensing system where the sensing matrix could be dynamically changed to perform different task-specific functions.

Acknowledgments

We would like to thank John van der Laan for his constructive comments. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the U.S. government. The work was supported by the Laboratory Directed Research and Development Program at Sandia National Laboratories, a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

References

1. D. Cireşan et al., "Multi-column deep neural network for traffic sign classification," *Neural Networks* **32**, 333–338 (2012).
2. B. Huval et al., "An empirical evaluation of deep learning on highway driving," CoRR (2015).
3. G. Levi and T. Hassner, "Age and gender classification using convolutional neural networks," in *IEEE Conf. Comput. Vision and Pattern Recognit. (CVPR) Workshops*, pp. 34–42 (2015).
4. R. Vardasca, L. Vaz, and J. Mendes, *Classification and Decision Making of Medical Infrared Thermal Images*, pp. 79–104, Springer International Publishing, Cham (2018).
5. A. Kumar et al., "An ensemble of fine-tuned convolutional neural networks for medical image classification," *IEEE J. Biomed. Health Inf.* **21**, 31–40 (2017).
6. A. Romero, C. Gatta, and G. Camps-Valls, "Unsupervised deep feature extraction for remote sensing image classification," *IEEE Trans. Geosci. Remote Sens.* **54**, 1349–1362 (2016).
7. E. Maggiori et al., "Convolutional neural networks for large-scale remote-sensing image classification," *IEEE Trans. Geosci. Remote Sens.* **55**, 645–657 (2017).
8. G. C. Birch et al., "Optical systems for task-specific compressive classification," *Proc. SPIE* **10751**, 1075108 (2018).
9. B. J. Redman et al., "Design and evaluation of task-specific compressive optical systems," *Proc. SPIE* **10990**, 109900H (2019).
10. E. J. Candès et al., "Compressive sampling," *Proc. Int. Congr. Math.* **3**, 1433–1452 (2006).
11. R. G. Baraniuk, "Compressive sensing," *IEEE Signal Process. Mag.* **24**(4), 118–121 (2007).
12. M. Elad, "Optimized projections for compressed sensing," *IEEE Trans. Signal Process.* **55**(12), 5695–5702 (2007).

13. J. Xu, Y. Pi, and Z. Cao, "Optimized projection matrix for compressive sensing," *EURASIP J. Adv. Signal Process.* **2010**(1), 560349 (2010).
14. G. Li et al., "On projection matrix optimization for compressive sensing systems," *IEEE Trans. Signal Process.* **61**(11), 2887–2898 (2013).
15. S. Xu, R. C. de Lamare, and H. V. Poor, "Distributed compressed estimation based on compressive sensing," *IEEE Signal Process. Lett.* **22**(9), 1311–1315 (2015).
16. M. A. Neifeld and J. Ke, "Optical architectures for compressive imaging," *Appl. Opt.* **46**, 5293–5303 (2007).
17. M. A. Davenport et al., "The smashed filter for compressive classification and target recognition," *Proc. SPIE* **6498**, 64980H (2007).
18. J. M. Duarte-Carvajalino and G. Sapiro, "Learning to sense sparse signals: simultaneous sensing matrix and sparsifying dictionary optimization," *IEEE Trans. Image Process.* **18**(7), 1395–1408 (2009).
19. R. Calderbank, S. Jafarpour, and R. Schapire, "Compressed learning: universal sparse dimensionality reduction and learning in the measurement domain," <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.481.8129&rep=rep1&type=pdf> (2009).
20. S. Lohit, K. Kulkarni, and P. Turaga, "Direct inference on compressive measurements using convolutional neural networks," in *IEEE Int. Conf. Image Process.*, pp. 1913–1917 (2016).
21. R. Timofte and L. Van Gool, "Sparse representation based projections," in *Proc. 22nd Br. Mach. Vision Conf.*, p. 61-1 (2011).
22. M. A. Neifeld, A. Ashok, and P. K. Baheti, "Task-specific information for imaging system analysis," *J. Opt. Soc. Am. A* **24**, B25–B41 (2007).
23. A. Ashok, P. K. Baheti, and M. A. Neifeld, "Compressive imaging system design using task-specific information," *Appl. Opt.* **47**(25), 4457–4471 (2008).
24. R. Bellman, *Dynamic Programming*, Computer Science Series, Dover Publications, New York (2003).
25. G. Hughes, "On the mean accuracy of statistical pattern recognizers," *IEEE Trans. Inf. Theory* **14**, 55–63 (1968).
26. S. Chen and D. Donoho, "Basis pursuit," in *Proc. 1994 28th Asilomar Conf. Signals, Syst. and Comput.*, vol. 1, pp. 41–44 (1994).
27. M. F. Duarte et al., "Single-pixel imaging via compressive sampling," *IEEE Signal Process. Mag.* **25**, 83–91 (2008).
28. J. E. Greivenkamp, *Field Guide to Geometrical Optics*, Vol. **1**, SPIE Press, Bellingham, Washington (2004).
29. L. Breiman, "Random forests," *Mach. Learn.* **45**, 5–32 (2001).
30. B. G. Grant, *Field Guide to Radiometry*, SPIE Press, Bellingham, Washington (2011).

Brian J. Redman received his PhD from the University of Arizona College of Optical Sciences in November 2019. His interests include modeling nontraditional optical systems, long-wavelength infrared imaging, polarization, and channeled imaging systems.

Amber L. Dage is a principal member of the technical staff at Sandia National Laboratories. Her work at Sandia encompasses nonconventional optics and imaging, including x-ray phase contrast imaging (2018 R&D 100 award), lensless imaging, three-dimensional imaging from structured illumination, and the development of a novel x-ray tomography system. She was recognized as a 2018 SPIE DCS rising researcher. Amber earned her MS (2008) and PhD (2011) in optical sciences from the University of Arizona.

Meghan A. Galiardi received her PhD in mathematics from the University of Illinois and is currently a researcher at Sandia National Laboratories. She has experience in mathematical modeling, algorithm design, and numerical analysis and has worked on a variety of machine learning projects, including neural networks and cyber applications.

Charles F. LaCasse received his PhD from the College of Optical Sciences in 2012, and his dissertation was titled "Modulated Polarimetry." He has worked at Sandia National Laboratories since 2013 and has worked on a variety of optical systems and signal exploitation in remote

sensing applications. These research thrusts include hyperspectral sensing, novel sensor design, machine learning/neural networks for hyperspectral imagery exploitation, and optical system performance modeling.

Tu-Thach Quach is a researcher in the Threat Intelligence Center at Sandia National Laboratories working on intelligent algorithms for information security and remote sensing problems. His research interests include the areas of statistical machine learning, information hiding, multimedia security and forensics, and computer vision. He received his PhD in computer engineering from the University of New Mexico.

Gabriel C. Birch received his PhD in optical sciences from the University of Arizona, College of Optical Sciences in 2012. He is a scientist at Sandia National Laboratories. His current research interests include computational imaging systems, optical design for compressive imaging systems, and nontraditional imaging systems.