SAND2019-15076C

Optimizing a Compressive Imager for Machine Learning Tasks

Brian J. Redman
Sandia National Laboratories
Albuquerque, USA
bjredma@sandia.gov

Daniel Calzada
Sandia National Laboratories
Albuquerque, USA
dacalza@sandia.gov

Jamie Wingo
Sandia National Laboratories
Albuquerque, USA
jwingo@sandia.gov

Tu-Thach Quach
Sandia National Laboratories
Albuquerque, USA
tong@sandia.gov

Meghan Galiardi
Sandia National Laboratories
Albuquerque, USA
mgaliar@sandia.gov

Amber L. Dagel
Sandia National Laboratories
Albuquerque, USA
aldagel@sandia.gov

Charles F. LaCasse
Sandia National Laboratories
Albuquerque, USA
cflacas@sandia.gov

Gabriel C. Birch
Sandia National Laboratories
Albuquerque, USA
gcbirch@sandia.gov

***Abstract*—Images are often not the optimal data form to perform machine learning tasks such as scene classification. Compressive classification can reduce the size, weight, and power of a system by selecting the minimum information while maximizing classification accuracy.**

In this work we present designs and simulations of prism arrays which realize sensing matrices using a monolithic element. The sensing matrix is optimized using a neural network architecture to maximize classification accuracy of the MNIST dataset while considering the blurring caused by the size of each prism. Simulated optical hardware performance for a range of prism sizes are reported.

***Index Terms*—compressive sensing, machine learning, optics**

I. Introduction

Imaging systems are traditionally designed such that single points in object space map as closely as possible to single points in image space. These devices attempt to create an image of the real-world. High image quality can improve the ability for human observers to make conclusions about the data; however, images are often not the most efficient data form to perform machine learning tasks such as scene classification.

Compressive classification systems are designed to only measure the information from the scene that improves classification accuracy. Each detector samples multiple regions of the scene, and the sampling can be optimized to minimize the number of detectors required. This break from the one to one constraint of imaging enables compressing the data before it is recorded. It has been shown that high classification accuracy is achievable by applying the classifier directly to compressed data [1], [2]. Reducing the number of detectors reduces the power required by the optical system and the bandwidth required to process or transmit the data. Removing the constraints of imaging also reduces the requirements to minimize aberrations which can reduce the size and of the optical components.

In our previous work we showed that a compressive sensing system could be implemented as a monolithic architecture using an array of prisms and discrete detector elements [3], [4]. This architectures differs from previous compressive sensing approaches [5], [6] because the measurements of all the detectors are made concurrently, but the prism array is the only optical element which reduces the size weight and power of the system. The prism array is created from optimized compression matrices where each nonzero element of the matrix is realized as an individual prism. Each prism directs one input angle onto a detector with multiple prisms contributing to each detector as shown in Fig. 1. Each prism is also independent of the other prism which lends the architecture to a completely computer generated optical design process which minimizes the design time from an optical engineer and thereby reduces the cost.

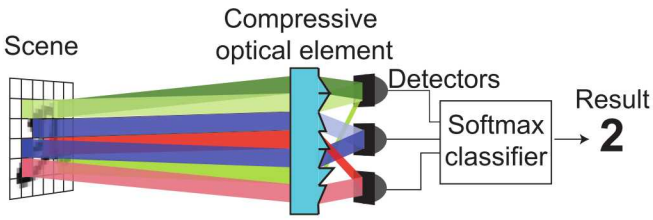

Fig. 1. Compressive classification samples multiple scene locations onto each detector element enabling classification with fewer measurements than would be required for imaging. A prism array enables the multiple sampling using a monolithic optical element to reduce the size weight and power of the optical system. The compressed data is classified instead of reconstruct images to minimize the required data.

The previous work indicated that the prism array blurred the sensing matrix. The geometry of the prisms set the acceptance angle which was larger than a pixel in the scene, and meant that each detector sampled larger regions of the scene than the compression matrix it was designed to realize. This work presents an optimization process that includes the effects of blurring. An additional layer was added to the optimization process which convolves the sensing matrix with a blurring function representing the prism width. This layer allows the optimization to exploit blurring to decrease the sparsity without increasing the number of prisms in a physically realized

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & 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-NA0003525.

version of the sensing matrix. The total number of prisms required to realize the sensing matrix is reduced by increasing the size of the prisms and the decrease in the classification accuracy is minimal.

II. Prism Array

The prism array architecture directly realizes a constrained sensing matrix as monolithic optical component located a fixed distance in front of discrete detectors. An example prism array that would make a measurement of the MNIST data set using 9 detector elements is shown in Fig. 2. A neutral density filter on each prism weights the contribution from each input angle. The detectors are separated so incident light off the designed input angle is not measured. Each prism, however, has a range of angles on either side of the designed input angle which will be directed on to the detector.

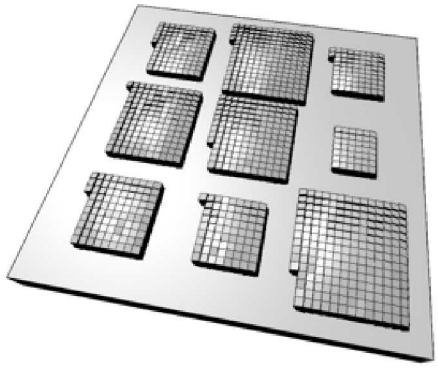

Fig. 2. An example prism array. Each prism realized one non-zero entry from the compression matrix. Each cluster of prisms contributes to a separate detector enabling each detector to sample multiple scene locations.

This range of acceptance angles causes blurring, as shown in Fig. 3. In traditional optical design situations, blurring is an often-undesirable feature that is driven down through iterative optimizations of the optical device. From the perspective of the sensing matrix, blurring increase sensitivity to angles adjacent to the designed input angle, but with a decreased weighting. Therefore, blurring can be thought of as localized relaxation of strict sparsity constraints on a sensing matrix that can be physically realized without additional optical hardware. The angular extent of the blurring is determined by the geometry of the prism array.

The geometry in this work is set to match previous works [3] as shown in Fig. 4. Each detector is 100 μm $\times$ 100 μm which is much larger than the pixels of a consumer camera. These large detectors are possible because the detectors are discrete. The gap between the detectors is require so that the light from outside the designed input angle of each prism will not be recorded. An imaging system would use a field stop to remove the unwanted light, but the prism array is not an imaging component. Because it is not an imaging component, separation between the prism array and the detectors is not set by a focal length. Instead the prism array is optimized for a set distance. A large separation decreases the cross talk between detectors, but also increases the size of the optical system. For this work the prism array-detector separation is set to 9 mm which matches previous simulations of the system. The field of view of 5° also matches previous work and corresponds to a reasonable field of view for remote sensing applications. The width of the prism was variable to show effects of blurring.

III. Simulations

A. Blurring Kernel

In previous work we created three dimensional models for the prism arrays and simulated each detector sensitivity as a function of input angle. Rays from a source were trace were traced for each of the 28×28 possible input angle. The total flux on each detector was recorded for each input angle to build up a system response matrix for the prism array. This matrix was used to compress images and simulate the performance of the optical system.

Non-sequential ray tracing provides an accurate simulation of the optical system but is slow and computationally intensive which makes it infeasible for use in the a neural network optimization process that requires many iterations. Blurring is the primary deviation away from the ideal sensing matrix caused by the prism array, so an analytic model for the blurring is required to include the properties of the prism array in the sensing matrix optimization. The detector sensitivity for a given angle is determined by the the overlap between the beam from a prism and the detector.

As a first order approximation the blurring was calculated for a prism directly above the detector as shown in Fig. 3 (a). The acceptance angle of the prisms can be modeled as a convolution between the effective angular size of the detector and beam as determined by the width of the prism. Both the prism and detector are rectangular shapes, and the convolution between two rectangular function results in a trapezoidal function as shown in Fig. 3 (b). The burring kernel goes to zero at the angle where the edge of the beam from the prism completely misses the detector as shown by the red beam in Fig. 3 (a). The flat part of the trapezoid starts at the angle where the smaller of the detector or the prism beam is completely encompassed inside the other shape. The overlapping area is independent of angle when the beam is entirely encompassed inside the detector (when the detector is larger than the prism), or the beam overfills the detector (when the prism is larger than the detector).

A two dimensional blurring kernel was computed for each prism size. The prism size, prism-array-detector separation, and the detector size was used to compute the angles for each part of the blurring kernel The profile is a piece wise defined function, so to avoid sampling errors, the blurring kernel was computed at ten times the resolution of the MNIST images and integrated to downsample. The blurring kernel created a close approximation of the blurring simulated using non-sequential ray tracing as shown in Fig. 5

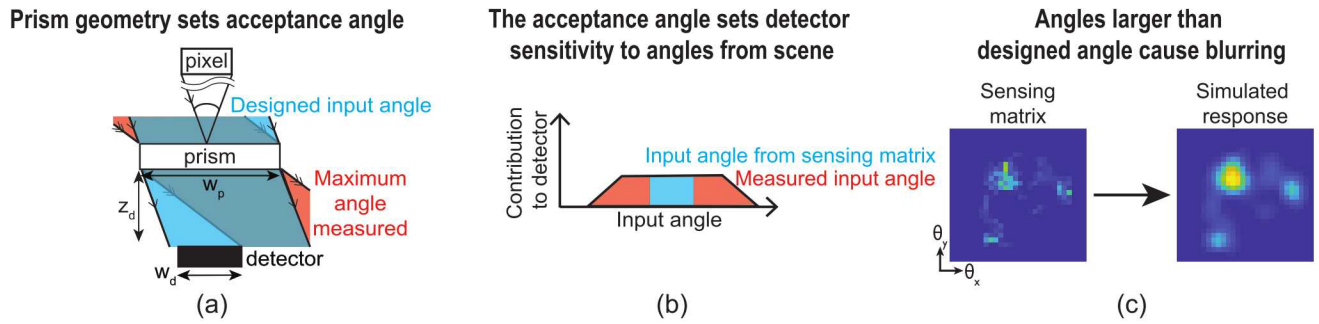

Fig. 3. (a) The separation between the prism array and the detector along with the width of the prism relative to the width of the detector sets which input angles are measured. Ideally each prism would only accept the angle that is subtended by a pixel in the scene. (b) The angular sensitivity can be modeled by how much of the beam from the prism is on the detector for each input angle. (c) If the sensing matrix is optimized assuming point sampling then the prisms accepting larger input angles will blur the system response matrix.

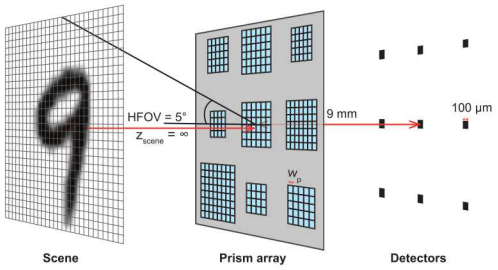

Fig. 4. The width of the prisms, the separation between the prism array and the detector, and the size of the detector affected the blurring of the sensing matrix. For this work, the separation between the detector and the prism array was fixed as was the size of the detector. The 28×28 pixel MNIST images were set to subtend a 5° half field of view.

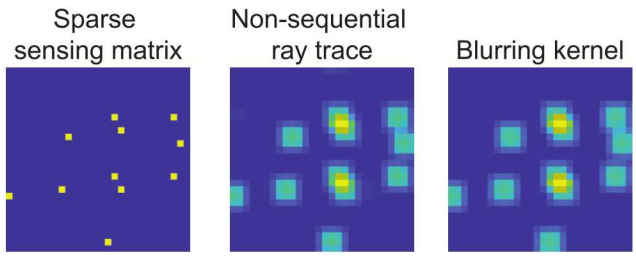

Fig. 5. Sensing matrices blurred with the blurring kernel showed close agreement to the system response matrices simulated using non-sequential ray tracing.

B. Neural Network Optimization

The prism array realized a compression matrix; however, the prism array cannot realize a general compression matrix. To allow the compression operation to be realized with optical hardware, it must be a linear operation. Thus, we create a transformation function C such that

$$C_n : \mathbb{R}^{784} \ni \mathbf{x} \mapsto \mathrm{B}(A)^\top \mathbf{x} \in \mathbb{R}^n, A \in \mathbb{R}^{784 \times n}, \quad (1)$$

where $\mathbf{x}$ is the flattened image being transformed, n is the dimensionality of the compressed space (the number of detectors), and the function $\mathrm{B}(A)$ is the blurring operator for the compression matrix A. To optimize this transformation for digit classification, we must optimize A jointly with another differentiable model whose output is the digit class. We have chosen a feed-forward neural network for this, allowing the whole process to be modeled in TensorFlow [7] as a differentiable feed-forward neural network consisting of three fully-connected layers. The first layer is the regularized and blurred compression transformation. The final layer classified the data into digit classes form 0 to 9 using a Softmax classifier.

To impose our soft sparsity constraints on the compression transformation, we apply L1 regularization to A in Equation 1 and constrain the entries of A such that $A_{ij} \geq 0$. L1 regularization is well-known to encourage sparsity in linear models and neural networks [8], [9]. While this increases the sparsity in A, the optimization phase exposes the undesirable property that the weights A_{ij} can become indefinitely small while the subsequent layer compensates by multiplying its weights by the same factor. The end result is A containing many small but non-zero weights. To dissuade this undesirable behavior, we add a small amount of Gaussian noise to the compressed representation,

$$\begin{aligned} C'_n : \mathbb{R}^{784} \ni \mathbf{x} \mapsto \mathrm{B}(A)^\top \mathbf{x} + \delta \in \mathbb{R}^n, A \in \mathbb{R}^{784 \times n}, \\ \delta \sim \mathcal{N}(\mathbf{0}, \sigma^2 I_n), \sigma = 0.02, \end{aligned} \quad (2)$$

where I_n is the $n \times n$ identity matrix and $\mathcal{N}(\mu, \Sigma)$ is the Gaussian distribution. Thus, if the weights in A become too small, the noise becomes the dominant factor in C'_n and little information is passed to the rest of the network. As an added benefit, the noise acts as regularization that can be turned off at test time [10]–[12].

To further encourage sparsity, once our model has converged, we approximate the compression transformation with a LASSO regression [9]. Depending on the regularization constant chosen, this further reduces the number of non-zero weights by a factor of 4-10x while generally achieving an $R^2 > 0.95$ relative to the original transformation. After this step, we fine-tune the rest of the model while freezing the sparse compression transformation.

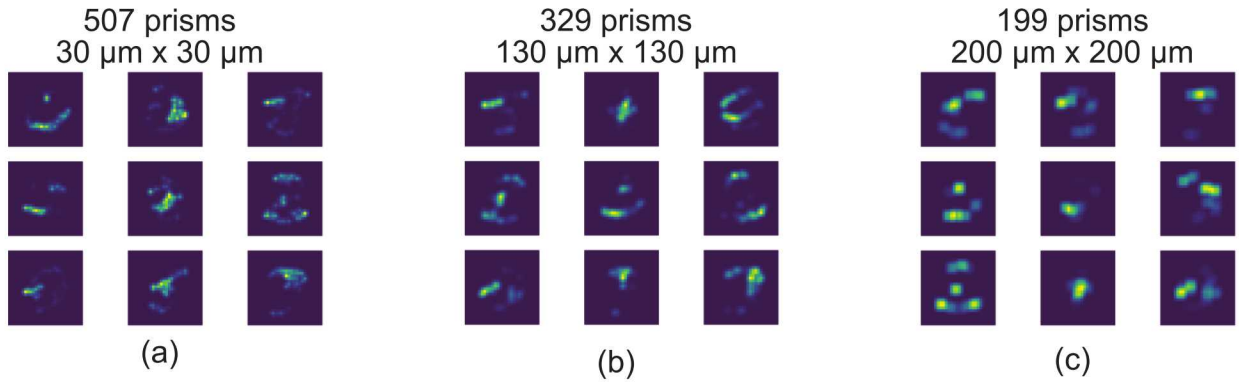

Fig. 6. (a) A sensing matrix describes the sensitivity of each detector to locations in the scene. The brighter a region the more contribution to that detector. This diagram is showing an example with nine detectors. k is the detector number. (b) Non-sequential ray tracing was used to simulate the sensitivity of each detector for a prism array 30 μm prisms. The small prisms resulted in minimal blurring. (c) The same prism array but with larger prisms is more blurred because the Larger prisms accept larger input angles.

IV. Results

The neural network optimization was used to create sensing matrices to classify the MNIST dataset using between 2 and 9 detectors. Each matrix was optimized both to maximize classification accuracy and for sparsity. The LASSO tuning parameter that sets the sparsity pressure was set to 0.01 to balance between the loss of classification accuracy from fewer measurements and the increased number of prisms from a less sparse sensing matrix. The optimization was run 17 times for each number of detectors to get a mean and variance for classification accuracy and number of nonzero entries in the compression matrix. Each neural network was trained on the 60,000 training images in the MNIST dataset. The number of nonzero entries was of particular interest because the prism array realizes each entry as a separate prism; therefore, sparse sensing matrices are smaller and less complex.

Fig. 6 shows example nine detector sensing matrices for a range of prisms sizes. Many of the same regions are sampled in the sensing matrices optimized for different size prisms. For example, the top right detector for the 30 μm prisms (Fig. 6 (a)) and the center detector for the 130 μm prisms (Fig. 6 (b)) both sample the center and an arc along the lower region of the MNIST images. The prism array with 130 μm prisms requires fewer prisms to realize the same sampling because each of the larger prisms sample a larger angular extent of the scene. Although, if the prism size becomes too large, the sampling regions significantly change because the blurring is larger than the regions that sensing matrix is optimizing to sample as demonstrate by the 200 μm prisms shown in Fig. 6 (c).

The performance of the system is determined by how well it can classify the MNIST dataset. A Softmax classifier was used to classify the 10,000 test images of the MNIST dataset compressed by each of the 17 sensing matrices created for each number of detectors from 2 to 9. The mean and standard deviation of the classification accuracy was recorded for each of the sensing matrices sets at each number of detectors and is shown in Fig. 7 (a). Three prism sizes are shown demonstrating that the the neural network optimization could achieve similar classification accuracies across a wide range of prism sizes. Ac cross all prism sizes and number of detectors, the performance of the optimized sensing matrices far exceeds the classification accuracy of the the same test images compressed using random Gaussian sampling.

The difference between the prism sizes become clear when comparing the number of prisms required to achieve the classification accuracy. Fig. 7 (b) shows the mean and standard deviation for the number of prisms required to implement each sensing matrix. The number of prisms required increases less than linearly as the number of detectors increases because each detector samples less of the scene. The larger prisms greatly reduce the number of prisms required to achieve the high classification accuracy.

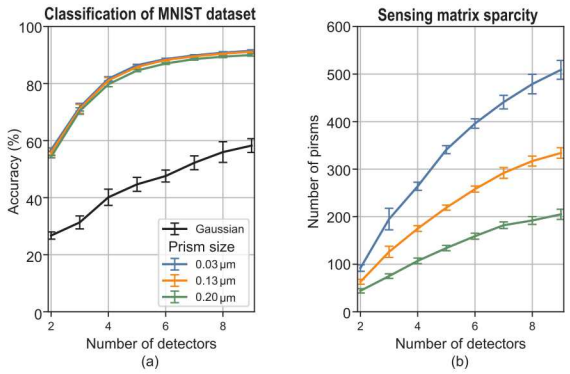

Fig. 7. (a) The classification of the the MNIST images from data compressed with the optimized sensing matrices had higher classification accuracy than images compressed with random Gaussian sampling. The decrease in classification accuracy due to increasing prism size was a small effect. The error bars are the standard deviation from 17 sensing matrices. (b) The number of non-zero entries in the sensing matrix sets the number of prisms required to realize the optical component. The larger prism required fewer prism to achieve similar classification accuracy. The random Gaussian sampling was not optimized for sparsity, so it has $784 *$ (Number of detectors) non-zero elements.

V. Conclusion

In this work we demonstrated a compressive sensing architecture using a prism array optimized based on the geometrical properties of the optical component. The prism array realized a compression matrix with a single optical element before discrete detectors. One challenge of the prism array architecture is that the size of the prisms sets the acceptance angle. Prisms effectively blur the sensing matrix. Reducing the size of the prisms reduces the blurring, but also decreases throughput of the optical system. Instead, the blurring was leveraged to decrease the number of prisms required to implement the sensing matrix.

The effects of this blurring were included in the neural network optimization of the sensing matrix by convolving the sensing matrix with a blurring kernel. An analytical model for the blurring kernel was created an verified against the results from non-sequential ray tracing. The number of prism was greatly reduced by increasing the size of the prisms and the decrease in classification accuracy was minimal.

Acknowledgement

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & 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-NA0003525.

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 United States Government.

References

[1] R. Calderbank, S. Jafarpour, and R. Schapire, "Compressed learning: Universal sparse dimensionality reduction and learning in the measurement domain," *preprint*, 2009.

[2] R. Timofte and L. Van Gool, "Sparse representation based projections," *Proceedings of the 22nd British machine vision conference-BMVC*, pp. 61–1, 2011.

[3] B. J. Redman, G. C. Birch, C. F. LaCasse, A. L. Dagel, T.-T. Quach, and M. Galiardi, "Design and evaluation of task-specific compressive optical systems," in *Defense and Commercial Sensing*, vol. 10990. SPIE, 2019. [Online]. Available: https://doi.org/10.1117/12.2520187

[4] G. C. Birch, T.-T. Quach, C. F. L. Meghan Galiardi, and A. L. Dagel, "Optical systems for task-specific compressive classification," in *Optics and Photonics for Information Processing XII*, vol. 10751. International Society for Optics and Photonics, 2018. [Online]. Available: https://doi.org/10.1117/12.2321331

[5] J. M. Duarte-Carvajalino and G. Sapiro, "Learning to sense sparse signals: Simultaneous sensing matrix and sparsifying dictionary optimization," *IEEE Transactions on Image Processing*, vol. 18, no. 7, pp. 1395–1408, 2009.

[6] A. Ashok, P. K. Baheti, and M. A. Neifeld, "Compressive imaging system design using task-specific information," *Applied optics*, vol. 47, no. 25, pp. 4457–4471, 2008.

[7] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: https://www.tensorflow.org/

[8] Y. Shi, J. Huang, Y. Jiao, and Q. Yang, "A semismooth newton algorithm for high-dimensional nonconvex sparse learning," *IEEE transactions on neural networks and learning systems*, 2019.

[9] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 58, no. 1, pp. 267–288, 1996.

[10] G. An, "The effects of adding noise during backpropagation training on a generalization performance," *Neural computation*, vol. 8, no. 3, pp. 643–674, 1996.

[11] Y. Bengio, "Deep learning of representations: Looking forward," in *International Conference on Statistical Language and Speech Processing*. Springer, 2013, pp. 1–37.

[12] C. M. Bishop, "Training with noise is equivalent to tikhonov regularization," *Neural computation*, vol. 7, no. 1, pp. 108–116, 1995.