CALIBRATION OF HIGH DIMENSIONAL COMPRESSIVE SENSING SYSTEMS: A CASE STUDY IN COMPRESSIVE HYPERSPECTRAL IMAGING

Graduate Students: Phillip Poon[1] and Matthew Dunlop[2]
Professional Staff: Esteban Vera[2] and Dathon Golish[2]
Faculty Advisor: Michael Gehm[1,2]
[1]College of Optical Sciences, [2]Department of Electrical and Computer Engineering, The University of Arizona, Tucson, AZ, USA

ABSTRACT

Compressive Sensing (CS) is a set of techniques that can faithfully acquire a signal from sub-Nyquist measurements, provided the class of signals have certain broadly-applicable properties. Reconstruction (or exploitation) of the signal from these sub-Nyquist measurements requires a forward model—knowledge of how the system maps signals to measurements. In high-dimensional CS systems, determination of this forward model via direct measurement of the system response to the complete set of impulse functions is impractical. In this paper, we will discuss the development of a parameterized forward model for the Adaptive, Feature-Specific Spectral Imaging Classifier (AFSSI-C), an experimental compressive spectral image classifier. This parameterized forward model drastically reduces the number of calibration measurements.

Keywords: Compressive Sensing, hyperspectral imaging, calibration.

1. INTRODUCTION

A hyperspectral imaging system produces images wherein each spatial location (pixel) has a corresponding spectrum [1, 2]. The input spectral density $S_o(x, y, \lambda)$ can be represented by a hyperspectral cube which has one wavelength and two spatial coordinates, as shown pictorially in Fig. 1. Hyperspectral imaging is used in many applications in remote sensing where the spectrum of a spatial region is of interest such as agriculture and geology [3, 4].

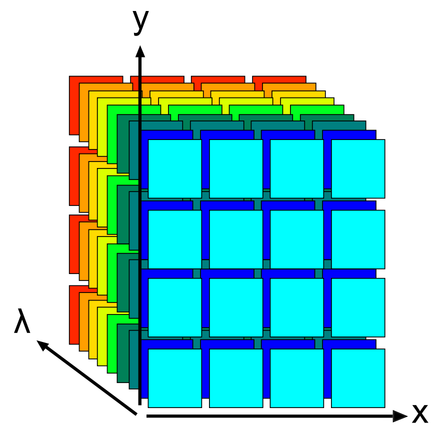

Figure 1. A diagram of a hyperspectral datacube with dimensionality $4 \times 4 \times 9$.

In an optical system with post-measurement additive noise $\mathbf{n}$, we can represent the relation between the signal $\mathbf{f}$ and the measurements $\mathbf{g}$ with

$$\mathbf{g} = \mathbf{Hf} + \mathbf{n}, \tag{1}$$

where $\mathbf{H}$ is the system measurement matrix.

Hyperspectral images have higher dimensionality than conventional RGB images, which generally increases the measurement time. A relatively low resolution hyperspectral cube with $128 \times 128 \times 76$ dimensions has 1.2 million voxels. Measuring the spectrum at each spatial coordinate with a spectrometer (i.e. the whiskbroom technique) would take over 1.6×10^4 measurements. A pushbroom technique simultaneously measures over a row or column of the hyperspectral image, by implementing parallel whiskbroom measurements with multiple sensors [5]. Nonetheless, the pushbroom technique still requires 128 spatial measurements to measure the entire hyperspectral cube. Reducing the number of measurements is important for many applications such as the classification of hazardous materials.

For the purposes of this paper we will refer to a conventional measurement system as any system for which the measurement of the signal, $\mathbf{g}$, resembles the signal itself, $\mathbf{f}$. This implies that the dimensionality or number of measurement samples is equal to the dimensionality of the signal. For example, the whiskbroom technique is conventional because it takes 128×128 spatial measurements for each of the 128×128 spatial locations.

Another name for conventional measurement systems is 'isomorphic', as the sampled measurements intentionally resemble the input signal. For a slit spectrometer, one can read out a linear detector array to obtain the spectrum without any digital signal processing. Similarly, in a conventional camera, the read out of the CCD pixels forms the desired image. Any use of compression algorithms to reduce the amount of data occurs after the signal is sampled.

In a compressive sensing (CS) system the number of measurements is much less than the native dimensionality of the signal [6–8]. This implies that the measurement matrix $\mathbf{H}$ is no longer approximately the identity matrix (hence, the measurement is no longer isomorphic). Accurate measurement in this context is feasible if the underlying signal $\mathbf{f}$ is sparse (most of the elements of a signal must be zero) or compressible (most of the elements are small) in some basis. In a CS measurement system the signal is compressed before it is sampled. A few examples of CS measurement systems include the single pixel camera, compressive MRI and single-shot spectrometers [9–11].

2. THE ADAPTIVE, FEATURE-SPECIFIC SPECTRAL IMAGING CLASSIFIER (AFSSI-C)

Our experimental prototype for compressive hyperspectral imaging is the Adaptive, Feature-Specific Spectral Imaging Classifier (AFSSI-C). The phrase feature-specific measurement is used interchangeably with compressive measurement [12].

The AFSSI-C optical design is shown in Fig. 2; the intermediate object plane is located at the front focal plane of a lens which collimates the light from the scene. The light is dispersed by a diffraction grating which laterally shifts each wavelength layer in the datacube along the direction of dispersion. The second lens then images the dispersed scene onto the digitial micromirror device (DMD). The DMD is a rectangular grid of programmable tilting mirrors. The mirror orientations are binary and can individually tilt to reflect the light into the second half of the system or to a beam dump. When the dispersed hyperspectral cube is incident upon the DMD, any spatial location (m, n) in the sheared hyperspectral datacube that is reflected into the beam dump effectively removes a column (in the wavelength direction) from the sheared hyperspectral cube which is reflected into the second arm. After passing through the third lens which recollimates the light, a second grating reverses the dispersion of the first grating on the hyperspectral datacube. The fourth lens then images the hyperspectral datacube on to the monochrome CCD camera. Read out of the CCD effectively flattens the hyperspectral datacube, with some spectral elements missing from some locations due to the code at the DMD.

For the AFSSI-C, calibration is crucial in the implementation of its measurement scheme. To understand the specific role of calibration we will briefly describe the AFSSI-C measurement and classification algorithms. A full description of the algorithms used is found in [13, 14]. The AFSSI-C uses adaptive feature-specific measurements in order to classify each spatial location's spectrum.

At first, the classification algorithm assumes each candidate spectrum at each spatial location has equal probability. After each measurement the camera intensity readout for a given location in the flattened datacube I_{mn}^{meas} is compared to the predicted intensity I_{mn}^{pred} for each possible spectrum at location (m, n). Using Sequential Hypothesis Testing (SHT), we monitor likelihoods rather than probabilities. A likelihood is the probability of a hypothesis (i.e. spectrum A is present) given a series of measurements. The ratio of likelihoods is computed after each measurement to update the probabilities of a given hypothesis being true. These probabilities are folded into the design process of the next DMD code. These probabilistically weighted measurements kernels are based on the system response to the candidate spectrums. Calibration allows one to acquire the as-measured spectral library which incorporates variations in the system response. Because classification takes into account the entire measurement history, any errors in this calibration will have a cumulative effect on performance.

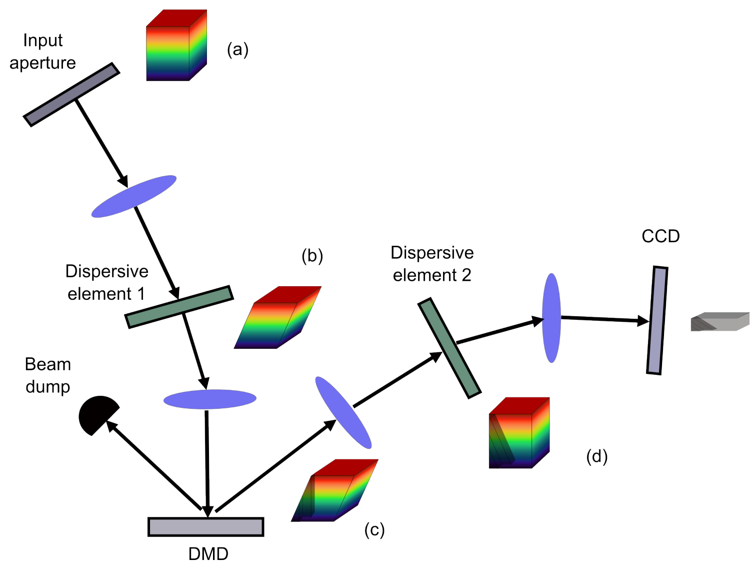

Figure 2. A diagram of the AFSSI-C and how the hyperspectral cube propagates through the optical architecture. (a) Light from a scene is represented by a hyperspectral cube. (b) The first diffraction grating disperses the colors of light, effectively shearing the cube along the dispersion direction. (c) The DMD leaves rectangular columns of rejected light in the sheared hypercube. (d) The second diffraction grating undoes the dispersion of the first grating. (e) The read out of the CCD flattens the datacube.

3. CALIBRATING THE ADAPTIVE, FEATURE-SPECIFIC SPECTRAL IMAGING CLASSIFIER (AFSSI-C)

The direct approach to calibration is to determine the impulse response of a system to each degree of freedom of the input signal [15, 16]. For hyperspectral imaging, we must measure the impulse response at each spatial location at every wavelength. This implies over 10^6 measurements for a $128 \times 128 \times 76$ hyperspectral datacube. In addition, this would be technically challenging, as it would require a narrowband filter or source for each spectral channel in the system.

Our first attempt at measuring the spectral library was an adaptation of the tunable filter approach, operating under the assumption that the spectra are spatially invariant over a column of monitor pixels. We display a given RGB value for a column of pixels on the monitor, which increases SNR compared to a single pixel and reduces measurement time, then physically sweep across the DMD mirror columns and measure the integrated intensity for each mirror column group. If we have 128 spatial columns in the hyperspectral cube and sweep across 76 mirror groups for each of the four candidate spectra, this still requires over 3.8×10^4 measurements.

To further reduce the number of measurements, we have developed a single shot approach to measuring the spectrum for each candidate in the library. Rather than sequentially sweeping across the columns of the DMD, we measure the entire spectrum at once by turning on a diagonal pattern of DMD mirrors and read out the CCD array, see Fig. 3. The CCD image is then segmented into groups of rows corresponding to the diagonal elements of the DMD pattern. The total signal in a given group of rows represents the intensity for that spectral channel. Coupled with the assumption that the spectra do not vary with position, it only requires c CCD exposures to measure the entire library, where c is the number of colors.

The most naive solution is to assume the system response is spatially invariant across the entire field of view, which allows the measurement of one column to represent all spatial locations. However, any real variation across the monitor will manifest as classification errors in the experiment. Alternatively, we can measure the system response of a column at a few locations across the monitor and interpolate between them to predict the spectrum at any location. This essentially treats calibration as a parameterized forward model and assumes that any spatially-dependent variations will be smooth. However, this technique still assumes the system response is shift invariant vertically, which will result in some locations being degraded by aberrations such as measurement error. Moreover, the impulse response at different locations is degraded by aberrations, such Scheimpflug distortion, and vignetting [17].

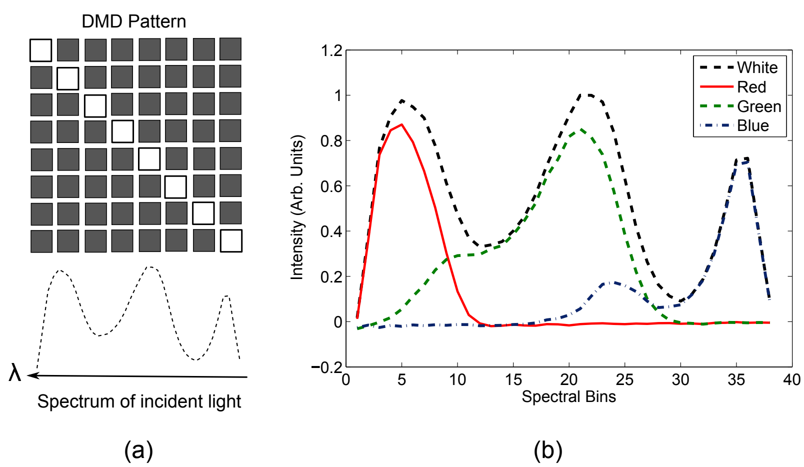

Figure 3. (a) A diagonal DMD pattern is used to measure the spectrum of a column displayed on the monitor. (b) The resulting spectrum for a white, red, green, and blue column.

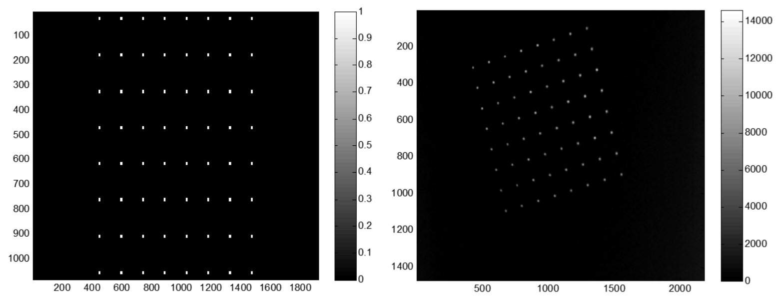

(a) (b)

Figure 4. (a) An 8×8 grid of dots is displayed on the monitor. (b) An image from the CCD of the 8×8 grid of dots. The centroid pixel coordinates of the dots in the image and their corresponding pixel coordinates on the object are used to estimate a projective transform between object and image.

For these calibration procedures, and for the experiment itself, we must understand how the image of the detector is related geometrically to the object on the monitor. For example, any distortion in the optical design will cause magnification variation as a function of field height. Scheimpflug distortion, which results from a tilted object plane (the DMD), produces different magnification variation in the transverse and longitudinal directions. This leads to a square object imaging to a trapezoid on the detector. In addition, the reflection axis of the micromirrors is along their diagonal and causes the detector (and the entire second arm) to be out of plane, which produces a rotation on the detector. Indeed the image of a vertical column on the monitor appears rotated on the detector. To account for these effects, we borrow an approach from the computer vision community which allows us to estimate the monitor scene given the detector readout.

We assume the distortion aberrations can be described geometrically as a projective transform [18]. The projective transform is a non-linear projection from object to image space. It is a more general form of affine transforms which includes linear transforms such as scale, rotation, and shear and the non-linear transform translation. In computer vision the projective transform is often used to simulate a change in the camera's perspective of an object. The projective transform is described by a 3×3 matrix which can be used to relate any set of pixel coordinates on the monitor $(m, n)_{\text{obj}}$ to a set of pixel coordinates on the CCD readout $(m, n)_{\text{img}}$.

The estimation of the transform matrix is requires control point pairs to determine how pixels from the object scene correspond to pixels in the image scene. Figure 4 demonstrates a set of points with known locations, $(m, n)_{\text{obj}}$, on the monitor and the corresponding set of points with measured locations, $(m, n)_{\text{img}}$, on the monitor. A minimum of four point pairs is needed to estimate the transformation matrix, however we use an 8×8 grid of points to improve accuracy. Since all the points are displayed simultaneously, a simple function iteratively locates the centroid within a small window on the image. For each control point, the window is moved to an approximate location and the centroid is calculated.

By applying this transformation matrix to every pixel in a detector image, we can reconstruct the object as it existed in undistorted object space. An example of this process can be seen in Fig. 5, where other optical effects such as vignetting and blur have reduced the quality of the final image. We can apply this projective transform to the measurement in our spectral calibration procedures to remove geometric distortions.

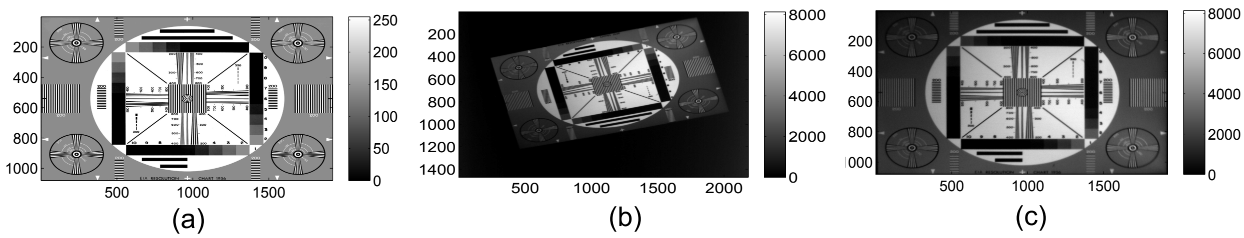

Figure 5. (a) The original image display on the monitor. (b) The image from the CCD. (c) Using the projection transformation, we reconstruction an estimate of (a) which corrects for geometrical distortions but does not correct for other effects such as vignetting, defocus, or diffraction that may degrade the point spread function.

4. CONCLUSION

CS techniques offer promising reductions in the number of measurements needed to acquire a signal. CS algorithms require knowledge of the system response to each degree of freedom in the input signal, but direct calibration for high dimensional CS measurement systems such as the AFSSI-C is impractical. We demonstrated a particular implementation of a parameterized forward model to drastically reduce the calibration time in the AFSSI-C from 1.2×10^6 measurement steps in the direct method to $(1 + kc)$ measurement steps in the single shot method, where c is the number of candidate spectra in the library and k is the number of locations (≈ 5) at which we perform the calibration. The constant 1 is required because one CCD exposure is needed to need estimate the projective transform.

Calibration continues to be a challenging issue; the lowest classification error, reached by the AFSSI-C for a $128 \times 128 \times 76$ hyperspectral datacube is approximately 5% [19]. We believe that in order to have the lowest classification error possible further advances in calibration must be achieved. In the future we hope to investigate alternative calibration techniques such as matrix completion [20].

REFERENCES

1. Gat, N., "Imaging spectroscopy using tunable filters: a review," in AeroSense 2000, 50–64, International Society for Optics and Photonics 2000.

2. Chang, C.-I., Hyperspectral imaging: techniques for spectral detection and classification, Springer 2003.
3. Cloutis, E., "Review article hyperspectral geological remote sensing: evaluation of analytical techniques," International Journal of Remote Sensing **17**(12), 2215–2242 1996.
4. Blackburn, G. A., "Hyperspectral remote sensing of plant pigments," Journal of Experimental Botany **58**(4), 855–867 2007.
5. Borengasser, M., Hungate, W. S., and Watkins, R. L., Hyperspectral remote sensing: principles and applications, CRC PressI Llc 2008.
6. Candès, E. J. and Wakin, M. B., "An introduction to compressive sampling," Signal Processing Magazine, IEEE **25**(2), 21–30 2008.
7. Stojnic, M., Xu, W., and Hassibi, B., "Compressed sensing of approximately sparse signals," in Information Theory, 2008. ISIT 2008. IEEE International Symposium on, 2182–2186, IEEE 2008.
8. Donoho, D. L., "Compressed sensing," Information Theory, IEEE Transactions on **52**(4), 1289–1306 2006.
9. Duarte, M. F., Davenport, M. A., Takhar, D., Laska, J. N., Sun, T., Kelly, K. F., and Baraniuk, R. G., "Single-pixel imaging via compressive sampling," Signal Processing Magazine, IEEE **25**(2), 83–91 2008.
10. Lustig, M., Donoho, D. L., Santos, J. M., and Pauly, J. M., "Compressed sensing mri," Signal Processing Magazine, IEEE **25**(2), 72–82 2008.
11. Gehm, M., John, R., Brady, D., Willett, R., Schulz, T., et al., "Single-shot compressive spectral imaging with a dual-disperser architecture," Opt. Express **15**(21), 14013–14027 2007.
12. Neifeld, M. A. and Shankar, P., "Feature-specific imaging," Applied Optics **42**(17), 3379–3389 2003.
13. Jansen, P., Dunlop, M., Golish, D., and Gehm, M., "Adaptive feature-specific spectral imaging," in SPIE Defense, Security, and Sensing, 83650B–83650B, International Society for Optics and Photonics 2012.
14. Dinakarababu, D., Golish, D., and Gehm, M., "Adaptive feature specific spectroscopy for rapid chemical identification," Optics Express **19**(5), 4595–4610 2011.
15. Gehm, M., "Calibration–an open challenge in creating practical computational- and compressive-sensing systems," in SAMPTA, 2013.
16. Townsend, D., Poon, P., Wehrwein, S., Osman, T., Mariano, A., Vera, E., Stenner, M., and Gehm, M., "Static compressive tracking," Opt. Express **20**(19), 21160–21172 2012.
17. Greivenkamp, J. E., Field guide to geometrical optics, SPIE Press Bellingham, Washington 2004.
18. Hartley, R. and Zisserman, A., Multiple view geometry in computer vision, vol. 2, Cambridge Univ Press 2000.
19. Dunlop, M., Poon, P., Vera, E., Golish, D., and Gehm, M., "The adaptive-feature specific imaging classifier (AFSSI-C)," in International Telemeteric Conference, 2013.
20. Candès, E. J. and Recht, B., "Exact matrix completion via convex optimization," Foundations of Computational Mathematics **9**(6), 717–772 2009.