Invited Paper

Adaptive spectroscopy: Towards adaptive spectral imaging

M.E. Gehm and J. Kinast

University of Arizona Department of Electrical and Computer Engineering, Tucson, AZ, 85743

ABSTRACT

Spectral imaging is an emerging tool for defense and security applications because it provides compositional information about the objects in a scene. The underlying task—measuring a 3-D dataset using a 2-D detector array—is challenging, and straightforward approaches to the problem can result in severe performance tradeoffs. While a number of ingenious (non-adaptive) solutions have been proposed that minimize these tradeoffs, the complexity of the sensing task suggests that adaptive approaches to spectral imaging are worth considering. As a first step towards this goal, we investigate adaptive spectroscopy and present initial results confirming dramatic cost/performance gains for a particular implementation.

Keywords: spectral imaging, adaptive optics, computational sensing

1. INTRODUCTION

Spectral imaging is a rapidly-expanding sensing modality for a variety of scientific and engineering applications primarily as a result of the supplementary information it provides about the chemical composition of the materials being imaged. While traditional imagers produce a two-dimensional spatial array of scalar values representing the intensity of a scene, a spectral imager produces a two-dimensional spatial array of vectors which contain the spectral information for the respective spatial locations. This data structure is known as the *spectral datacube* because of its three-dimensional nature. The addition of spectral data can provide valuable information in a variety of contexts ranging from defense and security applications[1,2] to biochemistry[3,4] to environmental monitoring,[5,6] to astrophysics.[7]

This manuscript discusses some fundamental challenges to traditional spectral imaging methods and describes a variety of computational sensing approaches that address these challenges. Later sections present the possibility that, for a variety of spectral imaging tasks, reconstructing the complete datacube is not necessary, and that system designs that directly exploit the spectral information would be suitable. In this context, the appeal of adaptive approaches to spectral imaging and spectroscopic exploitation tasks is discussed. The paper concludes with a detailed discussion of dynamic range matching spectroscopy—an implementation of adaptive spectroscopy that is designed for dealing with a common spectroscopic problem.

2. FUNDAMENTAL CHALLENGES IN SPECTRAL IMAGING

Straightforward approaches to spectral imaging (those based on simple extensions to traditional imaging or spectroscopy techniques) quickly ran afoul of two major challenges, both of which are fundamentally related to the dimensionality of the spectral imaging datacube. The first difficulty arises from the sheer number of signal elements in the datacube. For systems that involve a large number of spectral channels, the number of signal elements is obviously significantly larger than in traditional images with the same spatial resolution. This has a non-negligible affect on the signal to noise ratio (SNR) of the reconstructed datacube. For a given exposure time, there is some finite number of source photons that can be collected. After measurement acquisition, some methodology must be applied to apportion the photons among the signal elements in the datacube. By virtue of the increased number of signal elements, the overall number of photons assigned to any given signal element of the datacube is necessarily reduced from the result in a traditional imager. Small photon numbers then impact the signal to noise ratio of the estimated datacube in two ways. First, the number may prove small compared to additive noise provided by the measurement system (e.g. readout noise of a detector array). This difficulty can be mitigated with high-end detectors (with low readout noise) or by clever measurement schemes, as will be discussed below. Second, the low photon number leads to an inherent SNR penalty arising from the

Visual Information Processing XVII, edited by Zia-ur Rahman, Stephen E. Reichenbach, Mark Allen Neifeld,
Proc. of SPIE Vol. 6978, 69780I, (2008) · 0277-786X/08/$18 · doi: 10.1117/12.779174

underlying statistical fluctuations of the field (i.e. shot noise). Barring sources with exotic statistics, this penalty is unavoidable.

The second difficulty inherent in spectral imaging arises from the dimensionality mismatch between the datacube (3D) and a standard detector array (2D). As we consider datacubes of increasing size, the number of signal elements quickly outpaces the number of available detector elements. Using the standard mantra that solving a system of n unknowns requires at least n equations, we see that this mismatch will prove problematic for traditional reconstruction methods. Ultimately, the only possibilities are to minimize the size of the datacube or to employ some form of scanning approach in which a sequence of data acquisitions produces a total number of measurements (number of detector elements multiplied by number of acquisitions) that equals or exceeds the number of signal elements in the desired data cube. Scanning approaches, of course, incur a considerable time cost and limit the utility of the method for dynamic sources. Recent approaches have shown that, at least in some cases, the requirement that there be a number of measurements equal to or greater than the number of signal elements is also avoidable given clever measurement schemes. This will be discussed below.

2.1 SOLUTIONS TO THESE CHALLENGES

Advances in the field of spectral imaging over the past decade have largely focused on system designs that tackle the above challenges. The SNR penalty arising from small photon counts coupled with additive noise can be effectively overcome through the use of signal multiplexing. In multiplexed systems, measurements are not made on individual signal elements, but instead on combinations of elements. By combining multiple signal elements together, the measurement involves a larger number of photons, and the impact of additive noise is greatly reduced. The big challenges in multiplexed systems are in selecting the particular combinations of signal elements to measure and in designing an optical system that is capable of implementing these combinations.

One of the first attempts at a multiplexed spectral imager was Mooney's direct-view system.[8] In this approach, an intensity image of the scene was collected through a dispersive element. Whereas a standard approach (e.g. one based on a sequence of measurements through narrowband spectral filters) would collect only a small fraction of available light in a given acquisition, the direct-view system collects nearly all available light (losing only the small portion which misses the detector after interaction with the dispersive element). The system is clearly multiplexed, as any particular detector element receives light from a variety of position/wavelength points in the source datacube. The dimensionality mismatch remains, however, and the number of measurements in a single acquisition is far smaller than the number of signal elements. The direct-view system resolves this difficulty by rotating the dispersive element, and collecting a sequence of acquisitions at different disperser orientations. The technique is effectively a tomographic one, as a given acquisition is equivalent to a projection through the source datacube at an angle with respect to the wavelength axis. The magnitude of the angle is set by the linear dispersion of the system. As is common in tomographic systems, there is a conical, unsampled region of Fourier space (the "missing cone" or "limited angle" problem[9]) that impacts the final resolution of the reconstruction. Mooney later introduced algorithmic approaches for estimating the missing regions, reducing, though not eliminating, the resolution penalty.[10]

Not long after the introduction of the direct-view system, Descour and Dereniak introduced a system that also is based on tomographic sampling of the source datacube.[11] Their system, dubbed CTIS (Computed-Tomography Imaging Spectrometer), uses a holographic diffraction grating rather than a prism to produce the system dispersion. As with the direct-view system, there is a dimensionality mismatch between the datacube and the detector array. In contrast to the scanning approach used in the direct-view method, the CTIS approach implements all of the different disperser orientations into different spatial locations on the detector array, sacrificing datacube size (specifically spatial field-of-view) for all-at-once acquisition. As a result, the CTIS approach is applicable to time-varying sources whereas the direct-view method is not. Like the direct-view method, CTIS suffers from the missing cone problem; however the same algorithmic post-processing techniques can minimize the effect. Of the tomographic approaches, the CTIS technique is more widely adopted, and a variety of implementations and applications have appeared.[12–14]

Recent years have also seen a number of other, non-tomographic multiplexing approaches. The question of which particular signal element combinations should be formed for optimal performance can be formulated mathematically. For the case of only additive noise, it can be shown that the optimal combinations are related to mathematical objects known as Hadamard matrices.[15] A variety of systems have chosen to implement multiplexing schemes based on these techniques. A number of groups, most notably those led by Fateley and Riesenberg, have developed techniques that utilize digital micro-mirror (DMM) arrays to achieve the particular signal element combinations required by the Hadamard approach.[16,17] In contrast, the Brady group at Duke introduced a system that implements the combinations with a static aperture pattern.[18] In all cases, the dimensional mismatch exists and must be addressed. In the DMM systems, a sequence of acquisitions is made, changing the pattern of mirror orientations for each acquisition, while in the Duke approach, a linear translation is introduced between acquisitions in the sequence. All these Hadamard approaches have an overall throughput of 50%, resulting in a factor of two reduction in signal strength compared to the tomographic techniques. However, unlike the tomographic methods, the Hadamard techniques do not suffer from the missing cone problem, and pay no resolution penalty in the reconstruction.

For several years, there has been significant activity in an area known as "compressive sensing." Compressive sensing essentially overturns the belief that accurate solutions can only be found for fully- or over-determined linear systems. The underlying argument is that the majority of signals that arise in nature are sparse in some representation (e.g. many images are sparse in the wavelet basis; this is the basic concept of the JPEG2000 algorithm). The sparsity of the representation indicates that the signal has far fewer degrees-of-freedom than might be inferred from a naive counting of signal elements (e.g. an image is compressible because it can be represented by far fewer wavelet coefficients than there are pixels in the image). The implication is therefore that it is possible to achieve a valid signal reconstruction from a set of measurements that is smaller than the number of signal elements in the reconstruction, *provided* that the number of measurements exceeds the number of non-zero coefficients in the sparse representation of the signal. Originally, forming the reconstruction was thought to be a computationally intractable problem, but a number of mathematical results recently demonstrated that a simpler computation (which is amenable to standard linear programming techniques) can in many cases find the same solution.[19–22] This result has energized the sensing community and resulted in a significant number of experimental efforts. In the spectral imaging arena, the Brady group at Duke has utilized this concept to tackle the dimensionality mismatch problem. They recently demonstrated a multiplexed spectral imager that has no need for scanning of any type. The shortfall in detector elements with respect to signal elements in the datacube is overcome by using the detector elements to implement a compressive sensing approach to spectral imaging.[23] Early results are promising, but the reconstruction quality is lacking compared to other spectral imaging approaches. How much of this shortfall is an inherent result of the compressive sensing remains an open question.

The fundamental challenges to spectral imaging, then, have largely been addressed in one way or another by modern computational sensing approaches. The low signal levels with respect to additive noise are easily dealt with via multiplexing techniques (although some approaches introduce resolution penalties). The signal-to-noise penalty arising from the shot noise is not directly solved via multiplexing (multiplex approaches preserve the overall amount of shot noise). However, it is interesting to contemplate the fact that multiplexing *distributes* the shot noise evenly across the signal rather than concentrating it at the strongest signal elements—an effect that may prove to be useful.[24] Finally, the dimensionality mismatch between signal and detector can be resolved through the addition of compressive sensing techniques (though possibly with a performance penalty).

3. SPECTRAL IMAGING AS AN EXPLOITATION PROBLEM

As the previous section highlights, the majority of novel system design efforts in the spectral imaging field over the past decade have been focused toward addressing the challenges that are inherent to any attempt to estimate the spectral datacube. In this section, we argue that this effort, while clearly of great utility, misses the fact that, for most spectral imaging *applications*, estimation of the datacube is neither required, nor necessarily desired. The claims upon which we rest this argument are the following:

People rarely need or want a "pretty spectrum" The human visual cortex contains biological structures that have evolved over millions of years for sophisticated image processing and exploitation (identification, etc.) tasks. In contrast, human spectral processing is limited to color vision with relatively poor spectral resolution. Further, there is no native ability for analyzing spectral data presented as a list of numbers or a plot. Although such a skill can be learned (and is, by spectroscopists) the performance level is sill insignificant to what can be achieved by sophisticated signal processing methods. So, unlike images, which can be presented natively to an observer (the canonical "pretty picture") for advanced (and instinctive) processing, spectral data is largely taken *only* for the purpose of undergoing computational analysis.

Spectral image datacubes are rarely interpreted as a *gestalt* For a wide array of spectral imaging applications, the underlying goal is *not* finding complicated spatio-spectral structures within the cube, but rather the more mundane mapping of spectral information onto a spatial scene (although truly spatio-spectral approaches certainly have their place). Frequently, the desire in spectral imaging is to *supplement* spatial information about the scene with information extracted from the spectral content (most commonly, the chemical composition as a function of spatial position).

Based on these claims, we argue that, for a large number of spectral imaging applications, the proper approach is to directly apply spectral exploitation methods. However, much as the computational sensing methodology improved the performance of spectral imagers when the goal was datacube estimation, we believe that computational sensing approaches can improve the performance of spectral exploitation algorithms. In short, we argue that an important research direction is the development of novel system designs that have been optimized for a given spectral exploitation task by incorporating part or all of that task into the physical process of making a measurement.

Computational sensing approaches to any problem tend to involve the incorporation of some coding scheme into the measurement process, and the design of that coding scheme is then optimized for the particular task at hand. Of course, having optimized the coding scheme for a task, one can then consider giving the system the ability to algorithmically reoptimize itself based on information gathered from prior measurements. In this case, the system has become *adaptive*. In general, adaptive systems tend to have higher performance than non-adaptive systems, albeit with increased complexity and cost. In addition, we find that many exploitation approaches that we wish to consider are inherently iterative—an approach that lends itself well to adaptive systems. For all these reasons, we are interested in *adaptive spectral imaging*.

This is a relatively unexplored area. There has been some related work,[25] although in that system, the proposed architecture is again geared toward direct acquisition of the datacube. As mentioned above, for many tasks there is a natural separation of spatial information from spectral exploitation, we believe an important initial step is to explore what role adaptive system designs can play in spectroscopy (as opposed to the more complicated spectral imaging). By better understanding the methods by which systems can be constructed to perform adaptive spectral exploitation tasks, we gain information that we can later use when designing adaptive spectral imagers.

4. ADAPTIVE SPECTROSCOPY

As with adaptive spectral imaging, little research has been done in the area of adaptive spectroscopy (either for spectral acquisition or exploitation). This fact is particularly surprising given the successes of adaptive approaches to optical imaging over the past several decades.[26] The primary application of adaptive techniques to imaging is, of course, wavefront correction. This approach, first proposed in the early 1950's[27] and first successfully implemented in various forms over the next several decades, measures the distortion of a star or artificial guide star near the observation site. This measured distortion characterizes the aberrations introduced by propagation through the turbulent atmosphere. Adaptive structures within the light path are then used to cancel these measured aberrations and provide an undistorted view of the target. This approach has completely revolutionized telescope design to the point where few new major installations are constructed without adaptive optics.

As successful as wavefront correction has been, it involves what is, in reality, a very limited form of adaptability. The changes to the imager structure are used solely to counteract the changing aberrations introduced by the atmosphere—an intermediate medium that is not of interest to the observer at all. A richer form of adaptability would consider making changes to the optical structure based on *information* gained about the target of interest from previous measurements. That is, moving away from a naïve concept of observation to one based on *observation-as-measurement*. In this framework, the question to ask is "Given the existing prior knowledge about the source, what is the most *informative* measurement that could be made?" It is not surprising that the answer to this question could be something other than just merely canceling the aberrations introduced by the atmosphere. In recent years, several research groups have begun to explore this richer, information-centric form of adaptability. One major application has been extended dynamic range imaging.[28] In this approach, adaptive optics are used to accurately capture the details of a scene containing a dynamic range greater than the dynamic range of the detector array. Without an adaptive approach, the resulting image would be under- or overexposed (and possibly both). Yet another use of adaptability in imaging has been to explore sequential projective measurement to optimize the signal-to-noise ratio in a photon-starved environment.[29]

We are currently exploring several concepts in adaptive spectroscopy. The simplest is the spectral equivalent of the extended dynamic range imaging mentioned above. This approach is not, strictly speaking, associated with a spectral exploitation task; the method is concerned with acquisition of the spectrum, not its interpretation. However, it is an application that is readily suited to an adaptive approach, so it serves as an excellent starting point for our explorations. In terms of adaptive spectral exploitation approaches, we have primarily focused on library-based spectral identification tasks—matching the source to a library of known spectra (identification) or determining which mixture of known spectra match the source (chemometrics). We believe that adaptive computational spectrometers based on these ideas will outperform the combination of traditional instruments and post-processing algorithms. However, our work on the dynamic range matching is currently the most advanced, and the results from that work form the remainder of this manuscript.

4.1 DYNAMIC RANGE MATCHING SPECTROSCOPY

Dynamic range mismatches between the source spectrum and the instrument detector array are a common problem in spectroscopic measurement. In a traditional, non-adaptive instrument, the longest possible exposure duration, t_{max}, is that which just avoids bringing the contribution from the strongest spectral feature into saturation on the detector. Longer exposures will produce saturation and result in charge migration into neighboring detector channels (anti-blooming circuits can reduce, but not eliminate this effect), contaminating the measurements on those channels. Reading the signal from a detector, however, invariably includes a noise contribution. For signals with large dynamic ranges (defined as the ratio of the strongest to weakest features of interest), using an exposure duration t_{max} leaves the weakest features at levels which are indistinguishable from the detector noise. If the experimenter has available a time t_{total} for the measurement, with $t_{total} > t_{max}$, the only solution is to take a sequence of measurements of length t_{max} and average the results. The averaging has the effect of increasing the signal-to-noise ratio of the measurement. However, for any given detection metric, there is a sufficiently large source dynamic range where the signal-to-noise improvement from averaging is still insufficient to extract the weak features from the noise. When this is the case, the dynamic range of the source is said to exceed the dynamic range of the detector.

What is needed, then, is a method which allows the system to measure the strong spectral channels and then *turn them off* (or down) so that subsequent exposures may use longer exposure durations that can draw weak signals above the detector noise contribution. This process can then continue iteratively for as long as is alloted for the measurement, continually identifying channels that are nearing saturation, reducing or eliminating them entirely, and further increasing the exposure time. This is a task easily accomplished with adaptive components. How this might be achieved using common reflective or transmissive components is shown in Fig. 1 (for clarity, re-imaging optics are not shown).

4.1.1 THEORY

In the preceding discussion, a measurement scheme was outlined in which the exposure time for a given iteration of the measurement varied (in a controlled fashion) as the measurement progressed. This scheme, which would be problematic for a traditional spectrometer, is suitable for a device with an adaptive filter. Of course, the performance of the system is critically dependent on the manner in which the source spectrum affects which channels are active and what exposure durations are used. In this section we will provide a more detailed discussion of these topics.

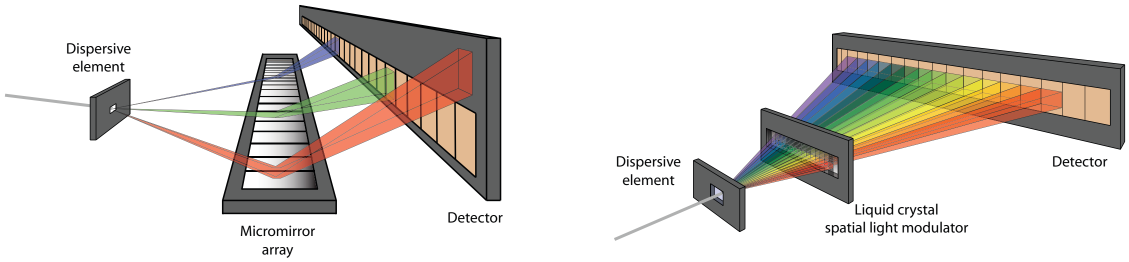

Figure 1. Reflective and transmissive methods of on/off channel control. In both cases the red channel (closest to the reader) has been turned off.

As the measurement process involves multiple signal acquisitions, we will use a superscript (k) to indicate the iteration of the measurement sequence. We assume that each detector channel has a dedicated filter that operates independently of all other channels. We represent the condition of the j^{th} filter channel by $f_j^{(k)} \in \{0, 1\}$. The exposure time is denoted $t^{(k)}$, and the measurement in detector channel j is given by $m_j^{(k)}$. If we include dark, readout, and shot noise, then a single measurement for a specific channel is given by

$$m_j^{(k)} = t^{(k)} \left(f_j^{(k)}\, \dot{s}_j + \dot{d}_j \right) + r_j^{(k)}. \tag{1}$$

In the preceding equation, $\dot{s}_j$ is the photon flux (including shot noise) from the source spectrum associated with channel j, $\dot{d}_j$ is the dark noise rate, and $r_j^{(k)}$ is the readout noise. Prior to the first measurement, the filter is set to transmit all light. That is, $f_j^1 = 1$. For subsequent measurements, if the anticipated value of $m_j^{(k)}$ will exceed the detector's saturation threshold for the chosen value of $t^{(k)}$, then $f_j^{(k)}$ is set to zero for that measurement. Once all measurements are complete, we can use a weighted average to generate an estimate of the spectrum's value in a particular channel

$$M_j^{(k)} = \frac{\sum_k \left[m_j^{(k)} - t^{(k)}\, \dot{d}_j \right]}{\sum_k f_j^{(k)}\, t^{(k)}}. \tag{2}$$

In constructing Equation 2, we have assumed that the readout noise contributes only zero mean additive noise. Moreover, it is unlikely that one would characterize dark noise rates on a channel-by-channel basis, in which case one can let $\dot{d}_j \to \dot{d}$, where $\dot{d}$ is the global estimate of the dark noise rate for the detector in use.

To this point, we have discussed time budget usage in general terms. Here, we provide a more detailed treatment. If we consider the case where we have a total time budget t_{total} and the exposure time is doubled between one measurement and the next, we have the scenario depicted in the lower portion of Figure 2. (For comparison, the top portion of the figure displays exposure durations for a traditional spectrometer with the same time budget).

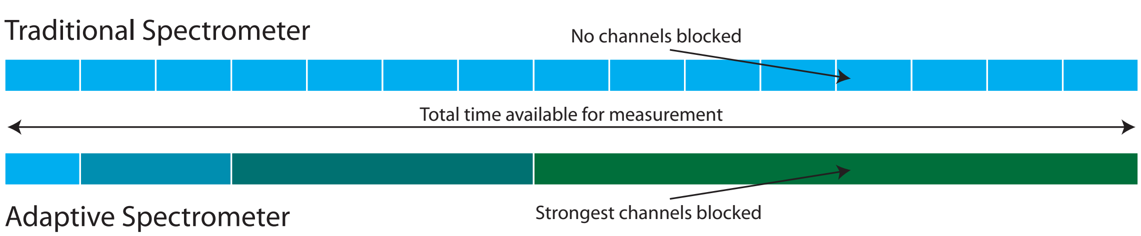

Figure 2. Schematic representation of how t_{total} is segmented in the traditional (top) and adaptive (bottom) instruments. Each block represents a separate exposure. In the adaptive instrument, blocking strong spectral channels allows longer durations in later exposures, improving the SNR of those channels.

Figure 2 depicts a scenario in which the sum of the exposure times for individual measurements coincides exactly with the value of t_{total}. In reality, this condition is unlikely. Consequently, some method for dealing with "misfit" remaining time must be developed. In the simulations that follow this section, we use the following sequence to use the allotted time budget:

1. An exposure of duration t_{max} is acquired.
2. Exposures of duration $2^1\, t_{max}$, $2^2\, t_{max}$, $2^3\, t_{max}$, and so on are acquired until doubling the exposure time once more would exceed the total time budget.
3. Acquire an exposure of duration $2^{\alpha}\, t_{max}$, where α is the largest integer that will not exceed the time budget.
4. Repeat step 3 until the remaining time is less than t_{max}. (The value of α is permitted to vary).
5. Once the remaining time budget has dwindled below t_{max}, simply take the longest possible exposure to fully exhaust the time budget.

Prior to each measurement in the above sequence, knowledge of the target spectrum gleaned from previous measurements is used to estimate which channels need to be turned on/off given the selected exposure time for the subsequent measurement. Using this model, simulations were conducted to assess the performance of a simple adaptive spectrometer.

4.1.2 SIMULATION

We now compare the simulated performance of a simple adaptive spectrometer to that of a traditional spectrometer. All simulations included shot noise, dark current (and its associated shot noise), as well as detector readout noise. For comparison, the simulations considered the performance of a set of four representative detectors: a high-end scientific camera (Andor DU434), a consumer-grade astrophotography camera (SBIG ST-7XME), a midlevel, industrial-grade board-level CCD (Lumenera Lu070), and a webcam (generic). The performance characteristics and cost of these detectors are shown in Table 1.

Type	Bit-depth	Dark rate (e^-/s)	Read noise (e^- RMS)	Cost
Scientific	16	0.001	7.5	$25,000
Astrophotography	16	1	15	$2,500
Midlevel Industrial CCD	12	10	10	$1,000
Webcam	8	10	10	$200

Table 1. Performance characteristics and cost of the representative detectors used in the simulation.

A summary of the simulation results is shown in Fig. 3. Here we take a signal-to-noise-ratio of 5 (defined as feature height divided by standard deviation of the noise) as our detection metric for the weak feature. The four solid lines show the nominal detection/non-detection dividing line for the four representative detectors when used in a traditional spectrometer. As expected, the more advanced detectors, by virtue of their better noise characteristics, are able to detect the weak spectral features at shorter measurement times and higher source dynamic ranges. The diamond data points on the plot represent the nominal detection/non-detection boundary for the webcam when used in the adaptive spectrometer. Regions below and to the right of the contour are system parameters for which the weak spectral feature is successfully detected. Note that the performance gain with increasing measurement time is significantly greater in the adaptive case. In fact, the webcam (cheapest detector) achieves performance parity with the most advanced detector at measurement durations on the order of 10–15 s.

In Fig. 4, a source spectrum with dynamic range $5{\times}10^3$ is simulated with a total measurement duration of 11 seconds (corresponding to the approximate crossing point in Fig. 3). For clarity, the source spectrum is highly artificial, containing a single strong and a single weak feature that are well separated. The results, however, are

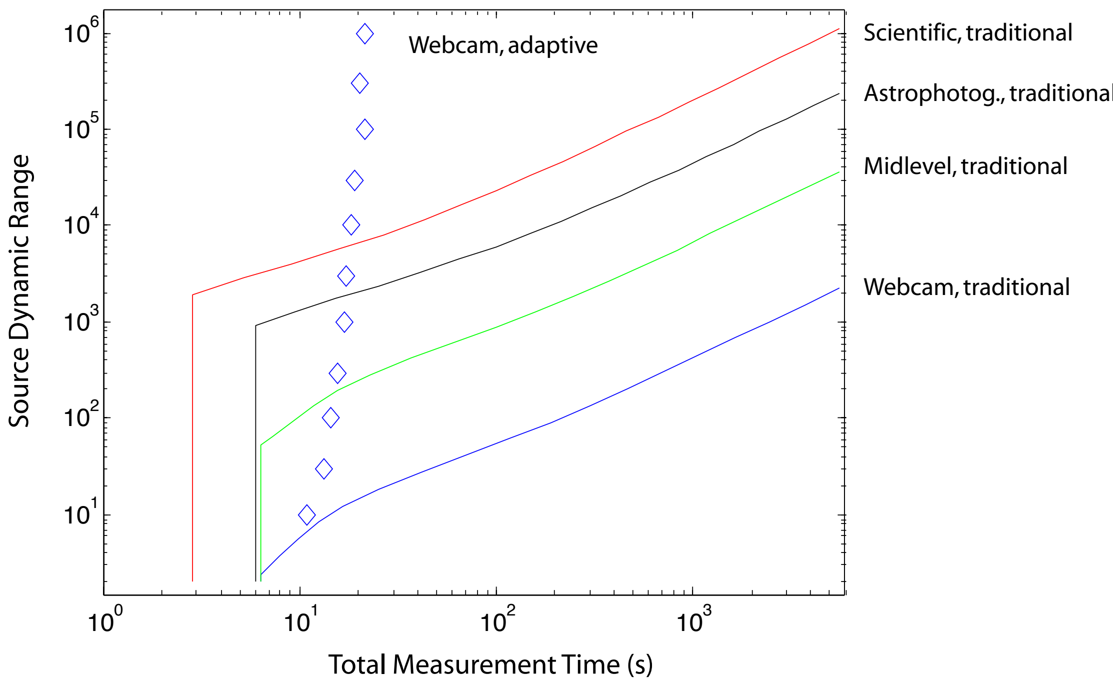

Figure 3. Result summary from the simulation. Solid lines represent the detection/non-detection boundary for the traditional spectrometer with different detectors. The diamond data points represent the detection/non-detection boundary for the adaptive spectrometer with the webcam detector. For any contour, regions below and to the right indicate system parameters which result in successful detection of the weak spectral feature. Note that the adaptive system becomes comparable with the most-advanced traditional system for measurement durations on the order of 10-15 seconds.

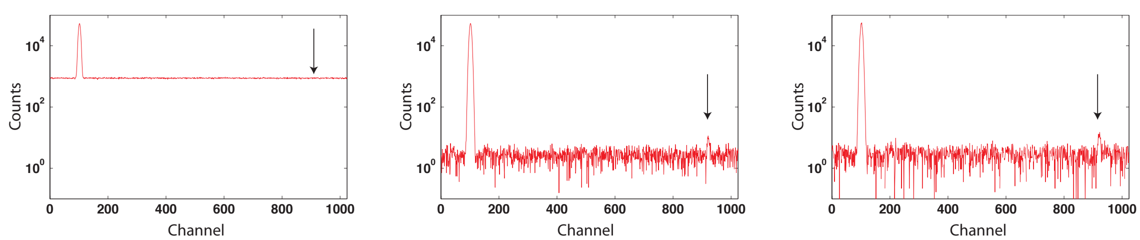

Figure 4. Source spectrum as measured by different instruments. (Left) Traditional spectrometer with webcam detector. Note that the weak peak is obscured by the noise baseline. (Middle) Traditional spectrometer with scientific-grade detector. The weak peak is clearly visible. (Right) Adaptive spectrometer with webcam detector. The weak peak is again clearly visible, despite the poor noise performance of the detector.

applicable to arbitrary spectral shapes. On the left is the spectrum as measured by the webcam in a traditional instrument. The strong peak is clearly visible, but the weak feature is obscured by the noise baseline. In the center is the spectrum as measured by the scientific camera in a traditional instrument. Both features are visible. Finally, on the right is the spectrum as measured by the webcam in a rudimentary adaptive spectrometer of the type described above. Again, note that both features are visible and the result is comparable to the traditional instrument with the high-end detector—while using a detector that is approximately two orders of magnitude less expensive.

The results of the simulation provide clear and strong support for the idea that inclusion of active components into the interior optical structure of spectrometers can have transformative effects on the price/performance of spectroscopic systems.

4.1.3 EXPERIMENTAL PROGRESS

We are currently constructing an experimental prototype to validate the predictions of this simulation and to serve as a testbed for future adaptive spectroscopy research. Adaptivity is provided by a digital micro-mirror (DMM) array that is located in what would be the traditional output plane of a dispersive spectrometer. The surface of the mirror array is then re-imaged by a subsequent set of optics onto a detector array. A schematic of the prototype systems is shown in Fig. 5. We expect to have initial experimental results shortly.

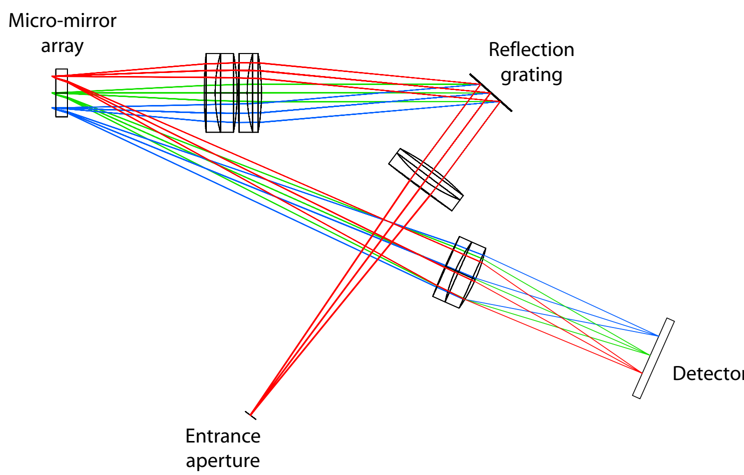

Figure 5. Optical schematic of the prototype dynamic range matching spectrometer currently under construction. The output of a spectrograph impinges on a digital micro-mirror array which is then re-imaged onto a detector array. Individual spectral channels can be turned off by changing the orientation of the appropriate mirrors in the array.

5. CONCLUSIONS

In conclusion, spectral imaging is a particularly challenging sensing modality as a result of the dimensionality of the resulting data. Computational sensing approaches have largely met these challenges, but in a way that focuses on accurate reconstruction of the complete spectral datacube without regard to how that data will ultimately be used. For a variety of cases, systems designs that more directly work to perform spectroscopic exploitation tasks in the hardware level may prove useful. Adaptive techniques appear especially well-suited to these tasks. Despite tremendous successes in imaging, adaptive techniques have made few inroads in the spectroscopic and spectral imaging modalities, and this appears to be a worthwhile area for exploration.

A particular implementation of adaptive spectroscopy—dynamic range matching spectroscopy—has served as a starting point in this exploration. Simple theory predicts, and simulations confirm, dramatic gains are possible in both system cost and performance. Experimental efforts are underway to develop an experimental prototype and initial results are expected soon.

REFERENCES

[1] Muise, R. and Mahalanobis, A., "Recent results of integrated sensing and processing using a programmable hyperspectral imaging sensor," *Proc. SPIE* **6245**, 624502 (2006).

[2] Hinnrichs, M., Jensen, J., and McAnally, G., "Handheld hyperspectral imager for standoff detection of chemical and biological aerosols," *Proc. SPIE* **5268**, 67–78 (2003).

[3] Denk, W., Strickler, J., and Webb, W., "Two-photon laser scanning fluorescence microscopy," *Science* **248**(4951), 73–76 (1990).

[4] Schultz, R., Nielsen, T., Zavaleta, J., Ruch, R., Wyatt, R., and Garner, H., "Hyperspectral imaging: A novel approach for microscopic analysis," *Cytometry* **43**, 239 – 247 (2001).

[5] Townsend, P. A., Foster, J. R., R. A. Chastain, J., and Currie, W. S., "Application of imaging spectroscopy to mapping canopy nitrogen in the forests of the central appalachian mountains using hyperion and aviris," *Geoscience and Remote Sensing, IEEE Transactions on* **41**(6), 1347–1354 (2003).

[6] Smith, W., Zhou, D., Harrison, F., Revercomb, H., Larar, A., Huang, A., and Huang, B., "Hyperspectral remote sensing of atmospheric profiles from satellites and aircraft," *Proc. SPIE* **4151**, 94 – 102 (2001).

[7] Lin, R. P., Dennis, B. R., and (Eds.), A. O. B., [*The Reuven Ramaty High-Energy Solar Spectrscopic Imager (RHESSI) - Mission Description and Early Results*], Kluwer Academic Publishers, Dordrecht (2003).

[8] Mooney, J., Vickers, V., An, M., and Brodzik, A., "High-throughput hyperspectral infrared camera," *J. Opt. Soc. Am. A* **14**, 2951–2961 (1997).

[9] Barrett, H., "Limited-Angle Tomography for the Nineties," *Journal of Nuclear Medicine* **31**(10), 1688 (1990).

[10] Brodzik, A. and Mooney, J., "Convex projections algorithm for restoration of limited-angle chromotomographic images," *Journal of the Optical Society of America A* **16**(2), 246–257 (1999).

[11] Descour, M. and Dereniak, E., "Computed-tomography imaging spectrometer: experimental calibration and reconstruction results," *Appl. Opt.* **34**, 4817–4826 (1995).

[12] Volin, C., Garcia, J., Dereniak, E., Descour, M., Hamilton, T., and McMillan, R., "Midwave-infrared snapshot imaging spectrometer," *Appl. Opt* **40**, 4501–4506 (2001).

[13] Sabatke, D., Locke, A., Dereniak, E., Descour, M., Garcia, J., Hamilton, T., and McMillan, R., "Snapshot imaging spectropolarimeter," *Optical Engineering* **41**, 1048 (2002).

[14] Hagen, N. and Dereniak, E., "New grating designs for a CTIS imaging spectrometer," *Proc. SPIE* **6565**, 65650N (2007).

[15] Harwit, M. and Sloane, N., [*Hadamard transform optics*], Academic Press, New York (1979).

[16] DeVerse, R., Hammaker, R., and Fateley, W., "Hadamard transform Raman imagery with a digital micromirror array," *Vibrational Spectroscopy* **19**(2), 177–186 (1999).

[17] Riesenberg, R. and Dillner, U., "HADAMARD imaging spectrometer with micro slit matrix," *Proc. SPIE* **3753**, 203–213 (1999).

[18] Gehm, M. and Brady, D., "High-throughput hyperspectral microscopy," *Proc. SPIE* **6090**, 609007 (2006).

[19] Candes, E. and Romberg, J., "Practical signal recovery from random projections," *Signal Processing, IEEE Transactions on* (2005).

[20] Candes, E., Romberg, J., and Tao, T., "Stable signal recovery from incomplete and inaccurate measurements," *Comm. Pure Appl. Math* **59**(8), 1207–1223 (2006).

[21] Donoho, D., "Compressed Sensing," *Information Theory, IEEE Transactions on* **52**(4), 1289–1306 (2006).

[22] Tsaig, Y. and Donoho, D., "Extensions of compressed sensing," *Signal Processing* **86**(3), 549–571 (2006).

[23] Gehm, M., John, R., Brady, D., Willett, R., and Schulz, T., "Single-shot compressive spectral imaging with a dual-disperser architecture," *Optics Express* **15**(21), 14013–14027 (2007).

[24] Brady, D. Private communication.

[25] Christensen, M., Euliss, G., McFadden, M., Coyle, K., Milojkovic, P., Haney, M., van der Gracht, J., and Athale, R., "Active-eyes: an adaptive pixel-by-pixel image-segmentation sensor architecture for high-dynamic-range hyperspectral imaging," *Appl. Opt* **41**(29), 6093–6103 (2002).

[26] Hardy, J., [*Adaptive Optics for Astronomical Telescopes*] (1998).

[27] Babcock, H., "The Possibility of Compensating Astronomical Seeing," *Publications of the Astronomical Society of the Pacific* **65**, 229– (Oct. 1953).

[28] Nayar, S. and Branzoi, V., "Adaptive Dynamic Range Imaging: Optical Control of Pixel Exposures over Space and Time," in [*IEEE International Conference on Computer Vision (ICCV)*], **2**, 1168–1175 (Oct 2003).

[29] Ke, J., Baheti, P., and Neifeld, M., "Applications of adaptive feature-specific imaging," in [*Visual Information Processing XVI*], *Proc. SPIE* **6575** (May 2007).