ACTIVE-EYES: an adaptive pixel-by-pixel image-segmentation sensor architecture for high-dynamic-range hyperspectral imaging

Marc P. Christensen, Gary W. Euliss, Michael J. McFadden, Kevin M. Coyle, Predrag Milojkovic, Michael W. Haney, Joeseph van der Gracht, and Ravindra A. Athale

The ACTIVE-EYES (adaptive control for thermal imagers via electro-optic elements to yield an enhanced sensor) architecture, an adaptive image-segmentation and processing architecture, based on digital micromirror (DMD) array technology, is described. The concept provides efficient front-end processing of multispectral image data by adaptively segmenting and routing portions of the scene data concurrently to an imager and a spectrometer. The goal is to provide a large reduction in the amount of data required to be sensed in a multispectral imager by means of preprocessing the data to extract the most useful spatial and spectral information during detection. The DMD array provides the flexibility to perform a wide range of spatial and spectral analyses on the scene data. The spatial and spectral processing for different portions of the input scene can be tailored in real time to achieve a variety of preprocessing functions. Since the detected intensity of individual pixels may be controlled, the spatial image can be analyzed with gain varied on a pixel-by-pixel basis to enhance dynamic range. Coarse or fine spectral resolution can be achieved in the spectrometer by use of dynamically controllable or addressable dispersion elements. An experimental prototype, which demonstrated the segmentation between an imager and a grating spectrometer, was demonstrated and shown to achieve programmable pixelated intensity control. An information theoretic analysis of the dynamic-range control aspect was conducted to predict the performance enhancements that might be achieved with this architecture. The results indicate that, with a properly configured algorithm, the concept achieves the greatest relative information recovery from a detected image when the scene is made up of a relatively large area of moderate-dynamic-range pixels and a relatively smaller area of strong pixels that would tend to saturate a conventional sensor.

OCIS codes: 110.2970, 110.3000, 110.4280, 230.3990, 230.4040, 230.4110.

1. Background and Motivation

The advent of hyperspectral imaging has been marked by an explosion in the quantity of data collected from a given scene. Given the limited number of photons that may be extracted from the scene and the scene's inherent limited information content, this increase has not been matched by the quality of data collected. In an effort to learn more about the hyperspectral phenomenology of scenes, sensors with hundreds of fixed wavebands have been developed. This has been at the expense of signal-to-noise ratio (SNR) in the wavebands. Currently, the hyperspectral-imaging problem is described as a data-mining problem, in which the vast datacube, with one spectral and two spatial dimensions, is harvested for subsequent extraction of useful information. This involves combining highly correlated data and projecting the datacube along arbitrary surfaces. This process is inherently wasteful, because the limited sensor readout bandwidth has been used to quantize more data than contain information. Most bits of the datacube are either later combined or

M. P. Christensen (mchriste@ap-photon.com), G. W. Euliss, M. J. McFadden, K. M. Coyle, and P. Milojkovic are with Applied Photonics, Incorporated, Fairfax, Virginia 22030-2309. M. W. Haney is with the University of Delaware, Newark, Delaware 19716. When this research was performed, J. van der Gracht and R. A. Athale were with George Mason University, Fairfax, Virginia 22030-4444. J. van der Gracht is now with Holospex, Incorporated, 6470 Freetown Road, Suite 200-104, Columbia, Maryland 21044.

Received 31 January 2002; revised manuscript received 3 June 2002.

0003-6935/02/296093-11$15.00/0

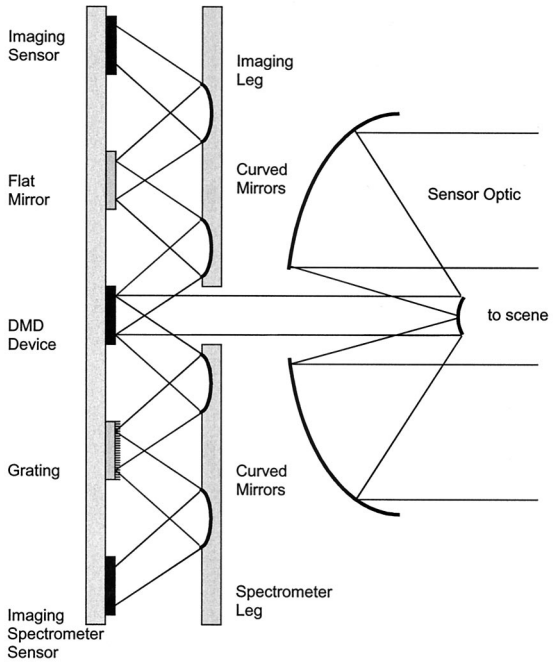

Fig. 1. Schematic diagram of ACTIVE-EYES architecture concept. Reflective sensor optics image the scene on the DMD device, which in turn segments the image into one of two legs. The imaging leg (shown on the top) relays the image to a broadband sensor, whereas the spectrometer leg introduces a grating to effect a starfield spectrometer.

discarded—only a small fraction of the bits contain useful information.

To illustrate this mismatch between data and information in a hyperspectral data cube, consider a situation with an expected datacube of 1024×1024 spatial elements with 200 bands per element. This is 200 MByte of data per frame, far exceeding the information content of typical scenes. If a scene is estimated to have less than 64 high-resolution spectra of interest and 64 spatial regions of interest (each 32×32 pixels large), then the total information content is ~78 KByte per frame. This represents a discrepancy of the order of 3 between information and data. Furthermore, in tactical applications such as battlefield awareness and battle damage assessment, these sensors are mounted on size–weight–power–downlink-bandwidth-limited platforms, and timeliness of information is a critical element. This situation embodies a technical rationale for simultaneous spatial and spectral processing that incorporates flexibility and adaptability for processing subsections of the data cube. The ACTIVE-EYES (adaptive control for thermal imagers via electro-optic elements to yield an enhanced sensor) approach, described in Section 2, provides the flexibility, adaptability, and efficiency to exploit the disparity between information and data.

2. ACTIVE-EYES Concept

The ACTIVE-EYES concept, shown in Fig. 1, integrates micro-opto-electro-mechanical (MOEM) technology with a spectrometer and a conventional imager for optimally extracting multidomain information in an imaging sensor. Similar digital micromirror device (DMD) architectures have been proposed and analyzed.[1–3] Blooming prevention through the use of the DMD to extend the dynamic range of an imaging system by removal of light from saturated regions was described.[1] The application of the DMD as part of a multiobject spectrograph, in which only a part of an image is spectrally analyzed and the rest of the light is rejected, was presented.[2] Finally, the dual use of the DMD device as a combination coronograph and multiobject spectrograph was proposed, and the timing, video-processing, and diffraction effects associated with the DMD were analyzed.[3] The ACTIVE-EYES approach described in this paper combines image amplitude control and spectral analysis in a manner that minimizes the number of wasted photons and maximizes the overall information extracted from the photon stream.

Using micromirror deflection to direct and modulate the intensity of a broadband image at individual pixels, the ACTIVE-EYES sensor adaptively guides the incoming light energy to optical sensor arrays that extract only the desired critical information for subsequent postprocessing. The deflection approach is based on MOEM micromirror switches. The ACTIVE-EYES concept achieves efficiency by adaptively, on a pixel-by-pixel basis, and in real-time, directing photons to the required spatial or spectral analysis sections. Using the high-rate temporal modulation and multistate nature of micromirrors, the ACTIVE-EYES concept processes spatial and spectral data simultaneously. Since the DMD device, sensor arrays, and any adaptive elements are all contained within the same multichip module plane (as shown in Fig. 1), the approach is potentially compact and compatible with conventional IR sensor foreoptics and detector packaging. Furthermore, a range of dynamically controlled or accessed dispersive structures could be integrated with the architecture to allow a programmable trade-off between spectral resolution, spatial resolution, and spectral coverage. By intelligent adaptation of the front-end to process only the proper mix of spatial and spectral data, the ACTIVE-EYES concept potentially provides orders-of-magnitude enhancement in the extraction of useful information through the front-end sensor by discarding large subsections of low-value elements of the spatial–spectral data cube before detection.

Figure 1 is a schematic diagram of the ACTIVE-EYES architecture. The foreoptics consist of a conventional telescope designed for the desired wavelength band and could be all-reflective to facilitate large optical bandwidths. An image of the scene is presented to the two-dimensional adaptive image segmenter by the foreoptics. The adaptive image segmenter performs two tasks: (i) image segmentation and redirection to specific optical channels and (ii) optical intensity control of individual pixel channels. Two optical trains are shown in this figure, above and below the DMD segmenter. The seg-

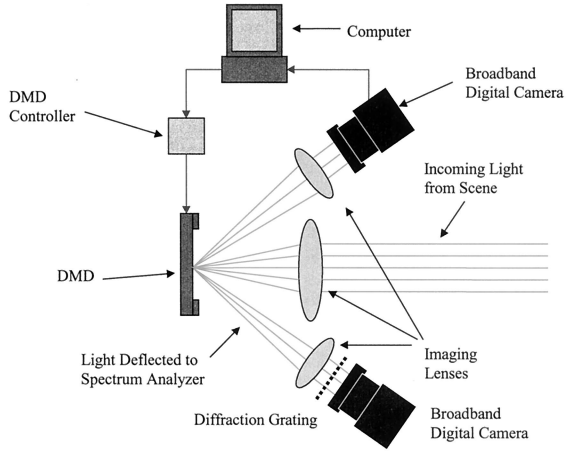

Fig. 2. Schematic depiction of ACTIVE-EYES experimental setup showing incoming light directed into one of two legs: one for broadband imaging and one for spectral analysis.

menter redirects specific individual image points or small subregions to the spectrometer channel. The remainder of the image is passed to a conventional imaging system. Any desired ratio of image/spectrometer energy from any given pixel can be implemented by rapid time modulation of the micromirror's position within a sensor integration time. Thus selected regions in the image may be analyzed by an imaging spectrometer while conventional imaging monitors the remainder of the image. One possible application of this particular design is an attentive spectrometer. In such a situation, the conventional image is scanned for potentially interesting features, perhaps as indicated by a particularly intense thermal signature. Each interesting feature is diverted to the spectrometer by deflection of a small part of the image into the spectrometer channel. The rest of the image can be continually monitored for new and potentially interesting events by observation of the image intensity channel. For example, in a damage-assessment situation the broadband imager could be used to locate burning debris. A subregion containing pixels from the flames of the debris could then be directed to the spectrometer leg where a high-resolution spectral analysis could be performed to determine what agents were being burned. The ACTIVE-EYES concept potentially provides orders-of-magnitude enhancement in the extraction of useful information through the front-end sensor. This is achieved by intelligent adaptation of the front-end to process only the proper mix of spatial and spectral data and to discard large non-information-rich subsections of the data cube before detection.

3. Experimental Results

To validate the ACTIVE-EYES concept, an initial laboratory prototype, based on off-the-shelf optical components, was constructed. Figure 2 is a schematic diagram of the experimental setup. The primary goal of this initial experiment was to demonstrate the two-legged adaptive image-segmentation element of the architecture. A commercially available Texas Instruments micromirror device segments an incoming scene image into two legs. One leg images the scene onto a CCD camera, whereas the other leg can be used to examine the spectra of selected regions in the image.

The prototype, shown in Fig. 3, was configured such that the Texas Instruments micromirror device was controlled to prevent CCD (charge-coupled device) blooming in the spatial domain in one leg while the other leg was used as a starfield spectrometer to examine the spectra of the regions of the image. The Texas Instruments DMD consists of 848×600 individually addressable mirrors, each 16 μm on a side. The mirrors can be tilted in one of two orientations separated by an angle of 20°. This allows for some of the light incident on the DMD to be deflected at 20° in one direction while the rest of it is deflected at 20° in the other. The two legs of the adaptive imaging system are defined by this deflection angle. With this setup it is possible to view simultaneously a scene imaged into the DMD without any saturation of the CCD camera detector elements, as well as to view the spectral content of the bright points in that image. The commercially available DMD controller[4] provides for one or multiple patterns to be displayed on the DMD in sequence at variable frame rates of the order of several kilohertz. Each frame is in the form of a black-and-white white bitmap file. Figure 4 shows a bitmap image of the Applied Photonics logo on the left; on the right, what the image looks like after it has been loaded onto the DMD.

When the light level is high enough to exceed the photoelectron capacity of a CCD element, the pixel becomes saturated, leading to information loss about the actual light level and other possible image-corrupting effects, such as blooming across other pixels. There are several ways to overcome this. In most cases saturation is reduced with a controllable optical iris, which limits the amount of light from the scene that actually reaches the detector array in the camera. Alternatively, the integration time of the CCD can be reduced to avoid saturation. Each of these methods is effective for alleviating the problem of CCD saturation, but both of them reduce detected light from *all* parts of the scene, when perhaps only one small section of the scene is causing saturation. Alternatively, a system using active pixel-by-pixel image segmentation allows saturated pixels to be controlled individually without decreasing the light detected from other pixels in the scene. The ACTIVE-EYES architecture is one way for the CCD array to detect the maximum amount of available light from every part of the scene, without any part of the CCD becoming saturated.

The objective delineated above is achieved by use of the algorithm described by the flow chart in Fig. 5. The CCD camera is positioned to form an image of the DMD, which itself has a scene imaged onto it by the foreoptics of the broadband optical sensor. A frame grabber captures a frame from the camera, which is

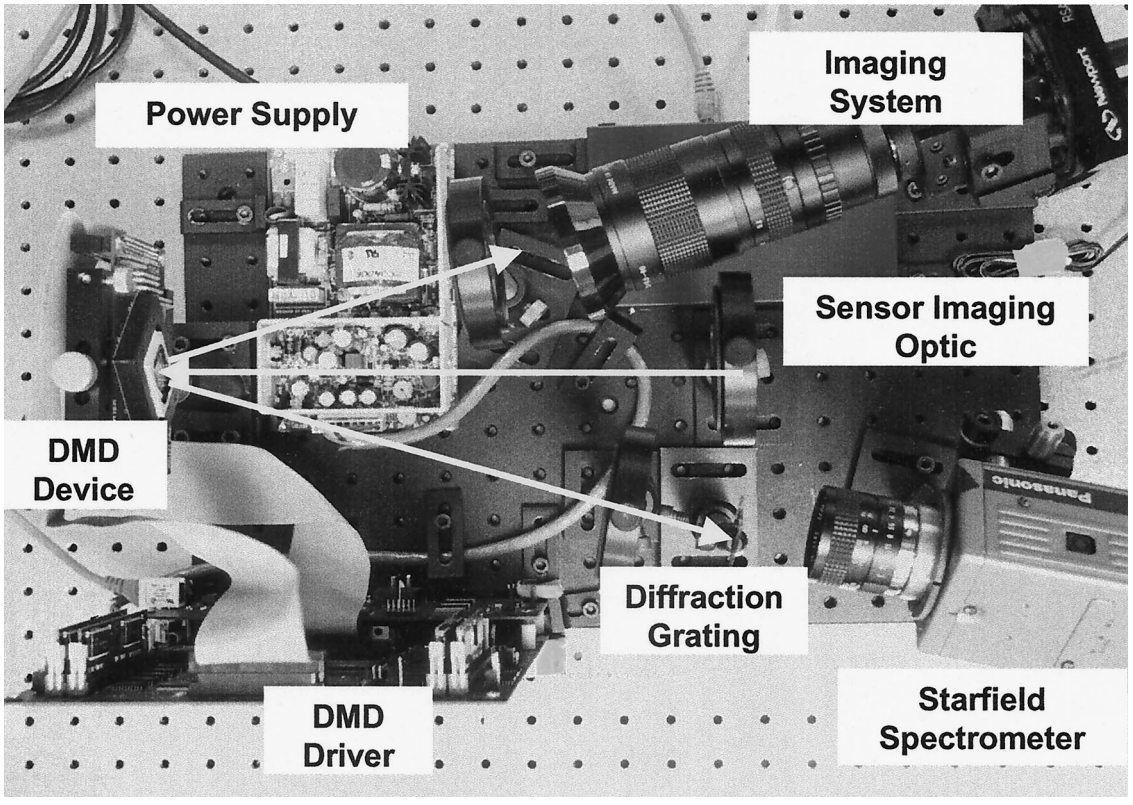

Fig. 3. Photograph of the ACTIVE-EYES experimental setup.

then analyzed to see whether any of the pixels are above some predefined threshold. The image is modified such that the pixel with a value above the threshold are made black while all the others are made white. The colors black or white in the bitmap file determine the direction that the micromirror at the location corresponding to each pixel will be tilted when the file is loaded onto the DMD. The image resulting from the threshold operation is added to the queue of frames to be displayed on the DMD in sequence with a repetition rate many times the frame rate of the CCD. In this way the mirrors in the areas of the image that were saturated now deflect toward the spectrometer for half the time, after the first iteration. Another frame is captured, and the same procedure is performed on it so that any areas of the image that are still saturated will again be deflected from the camera. In each successive iteration the duty cycle of the pixels above threshold is reduced by a factor of 2. This results in eliminating the saturation in the areas of the image where it occurs, without unnecessarily deflecting any light

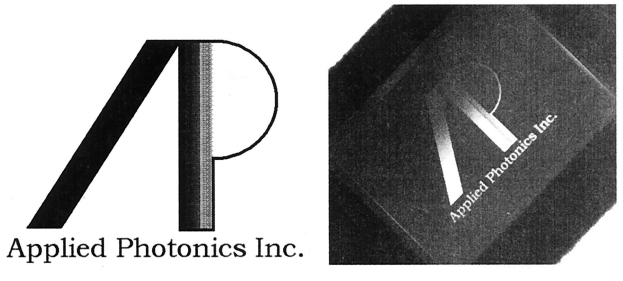

Fig. 4. Example test image as it appears in its original bitmapped form and as displayed on the DMD.

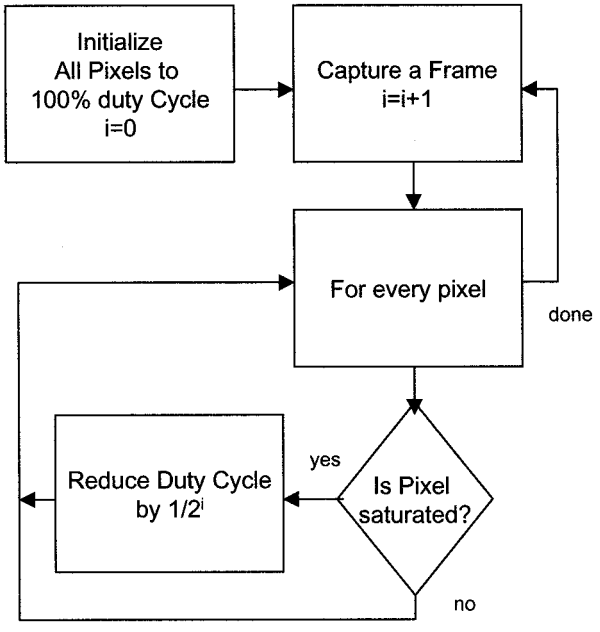

Fig. 5. Flow chart describing algorithm for eliminating saturation in the sensor by iterative reduction of the DMD duty cycle on a pixel-by-pixel basis.

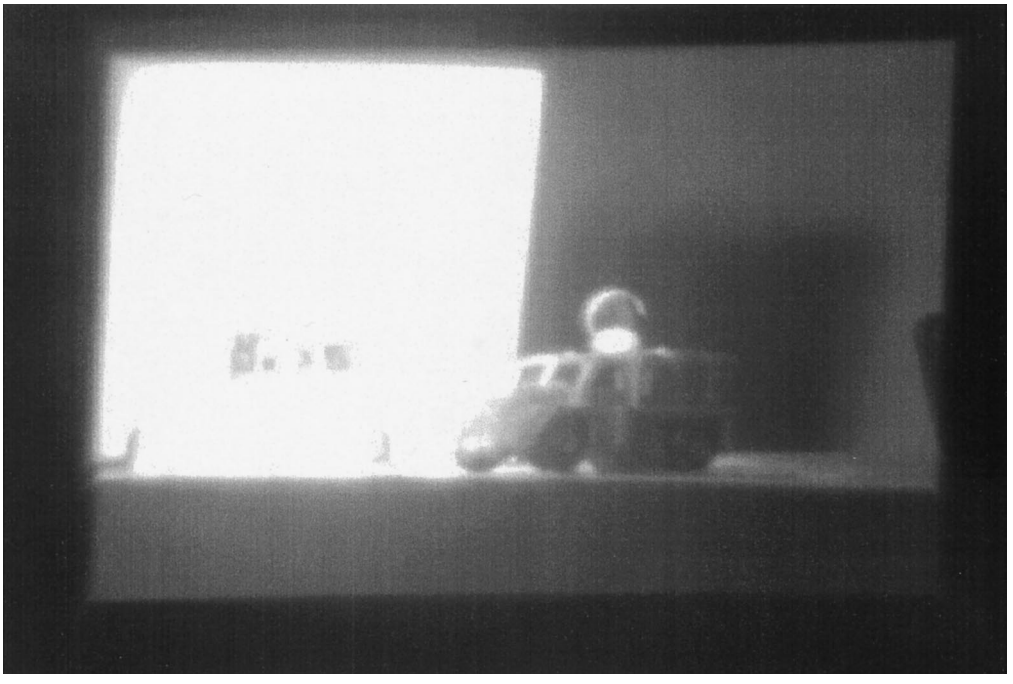

Fig. 6. Example of a highly saturated scene with all micromirrors at 100% duty cycle.

away from the more dimly lit areas. At the same time, since the bright points are deflected to the spectrometer for most of the cycles, the spectra of these points can be viewed on the other image-segmentation leg.

The capabilities outlined above were demonstrated in the following experiment. A test object scene, consisting of metal objects (model trucks) with mirrors attached and highly reflecting white paper with printed words, was used. The objects were illuminated with a strong white-light source, and various color filters were used to evaluate the spectral processing section. The scene was deliberately illuminated to the point of saturating the CCD on the image viewing side of the image segmenter. The saturation is evident in the unprocessed scene image shown in Fig. 6, which has a large washed-out region of saturation resulting from the highly reflective paper and a smaller saturated region that is due to glinting from the metal model truck. A small section of the large washed-out region and the glint from the truck were diverted to the starfield spectrometer leg of the architecture. The resulting spectra are shown in Fig. 7. The spectrum of the washed-out region contains the continuum of visible wavelengths, whereas the spectrum from the toy truck glint contains two bands.

Once the spectral analysis was completed, the DMD was available to correct the nonuniform illumination of the scene. After the algorithm described above was run through several iterations, the CCD was no longer saturated in any region, thus allowing objects that were previously not discernable in the scene to be readily identified. This same procedure could be applied to a broadband IR sensor application. Figure 8 shows the imaging leg scene after the saturation-control algorithm was implemented.

In many cases it is true that the bright spots in the scene are interesting. That is, a bright spot in an image may be indicative of metal glinting, a fire, or some other kind of plume. For this reason we might want to analyze the spectrum of the light from each of these bright spots so that we can identify it by its spectral signature. Fortunately, the byproduct of the saturation-control procedure described above is that most of the light from the bright regions of the scene is deflected from the imaging leg. In the other leg a spectrometer is used to analyze the spectrum of each bright spot in the scene, thereby differentiating between broadband reflection of the paint from the truck and a narrow band (filtered white light) from the reflective feature on its roof. In this way the image-segmentation architecture was used to prevent CCD saturation and analyze bright regions in a scene.

4. Analysis

An imaging system can be thought of as an array of communication channels. Spatial information is transferred from the scene to the restored image, contained in the values assigned to each pixel by the imager. The amount of information is a function of the accuracy associated with those pixel values. When a pixel value exceeds the saturation level of the associated detector, the accuracy of the imaging system is lost in that pixel. Hence the information capacity of the imaging system is degraded. Furthermore, when a pixel value lies below the noise level in the associated detector, the value of that pixel becomes uncertain and the information regarding the

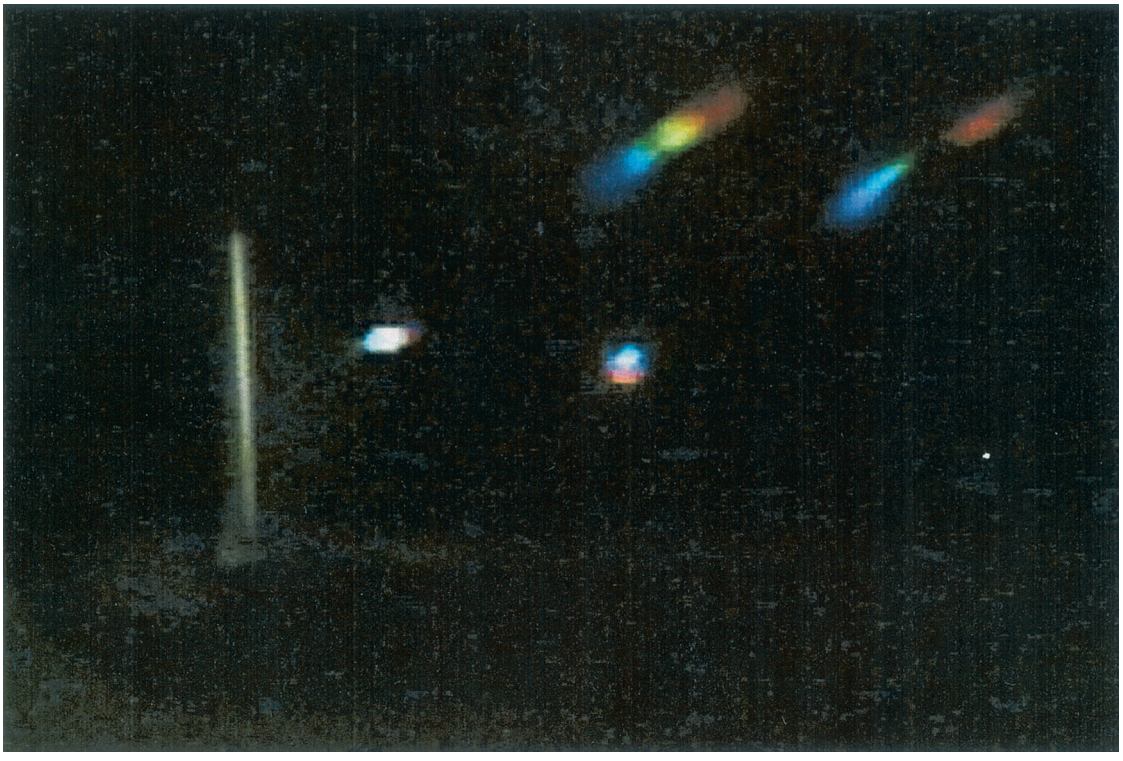

Fig. 7. Screen capture from starfield spectrometer leg showing the spectra from two glints from the trucks scene shown in Fig. 6.

scene in that pixel cannot be accurately recovered. This suggests that competing results must be considered when the maximum pixel value of the collected image exceeds the dynamic range of the photodetector array. Globally attenuating the image will recover pixel values that are above the detector saturation level, but this may come at the expense of pushing low-valued pixels below the noise floor.

The preliminary experimental results presented in Section 3 show that the concept of using a DMD for pixel-by-pixel local attenuation holds promise for extending the effective dynamic range of the sensor and

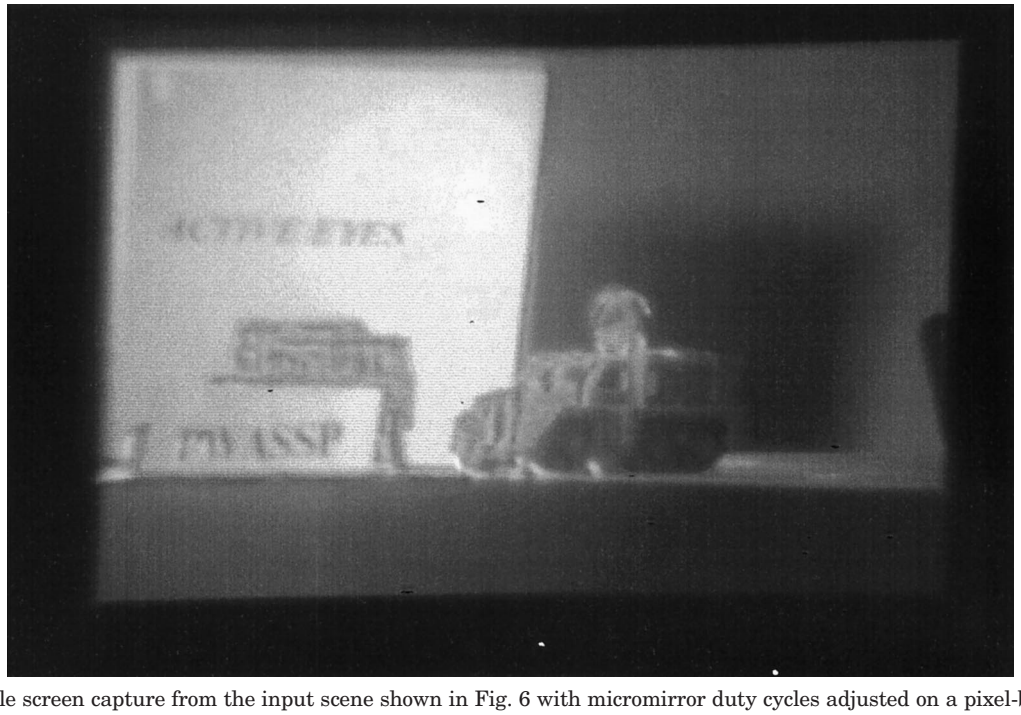

Fig. 8. Example screen capture from the input scene shown in Fig. 6 with micromirror duty cycles adjusted on a pixel-by-pixel basis to extend the dynamic range of the CCD (the interlacing of the CCD camera is evident in the stripes in the saturated areas).

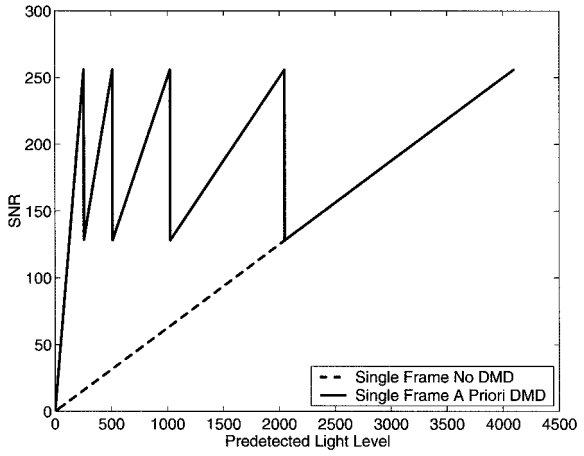

Fig. 9. Plot of DMD sensor contrast enhancement showing SNR for input pixels requiring four effective doublings of the maximum detectable light level. This graph assumes *a priori* knowledge of the required attenuation.

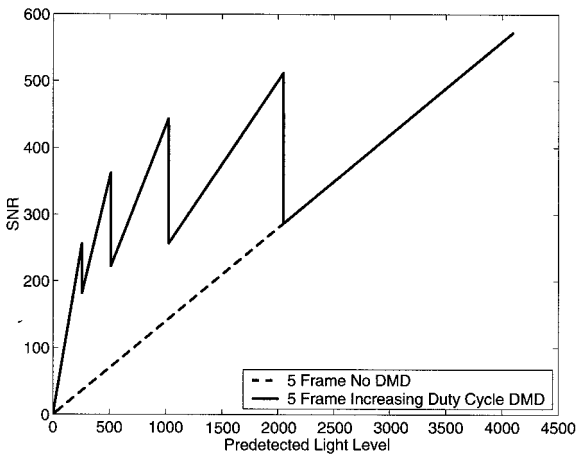

Fig. 10. Plot of DMD sensor contrast enhancement showing SNR for pixel measurements assuming an iterative approach beginning with maximum attenuation and decreasing pixels that are beneath the noise floor on each successive frame for five frames. This is compared with a global attenuation of 1/16 with five integrated frames.

recovering information otherwise lost to saturation. To quantify the potential benefits of the pixelated DMD-based saturation-control approach, analyses were performed comparing local attenuation to a more conventional approach using global attenuation. For comparison, an 8-bit camera was assumed with a noise floor of 1 corresponding to the least-significant bit (LSB). The pixelated attenuation is interesting only when the maximum pixel value of the collected image exceeds the dynamic range of the sensor; if this is not the case, then the pixelated approach has no benefit. When the maximum pixel value of the collected image exceeds the dynamic range of the sensor, the global attenuation approach has two options: (1) Accept the saturated pixels and assign those pixels a single value in the restored image, or (2) attenuate the entire image prior to detection so that all detected values are beneath saturation. In either case, some of the spatial information contained in the predetected image is lost. For the following analysis we consider the second of these two options.

With *a priori* knowledge of the magnitude of each pixel in the input scene, the iterative part of the algorithm described in Section 3 is not necessary. When we estimate the noise floor as 1 (the LSB of the camera), Fig. 9 shows the detector SNR corresponding to a range of pixel values in the predetected image (i.e., predetected pixel values). The light levels in the predetected image are also normalized to the LSB of the camera. For the range of light levels between 2048 and 4095 the local approach and the global approach are identical, because they both derive from an attenuation of 1/16. For all lower light levels, the local approach shows an advantage that is due to its ability for selectively reducing the attenuation in these regions appropriately. This advantage, however, relies on an accurate prediction of the input levels to choose the appropriate attenuation. Otherwise, the initial image will be attenuated more than is necessary or will still contain saturated regions. As described in Section 3, an iterative algorithm can quickly lock in on the correct attenuation factor, but this will require at least two, and perhaps several, frames. These same frames could otherwise be summed to improve the SNR of the global approach. Figure 10 takes into account the requirement for multiple frames. The plot of Fig. 10 shows the cumulative effect of five collected frames (necessary to go from 1/16 attenuation down to no attenuation by factors of 1/2). Note that the SNR of the global approach is increased over that of the single frame shown in Fig. 9. Again, the local attenuation has improved the SNR of the measurements over that of the global approach, but the improvement is less in the regions with fewer integrated frames. If this approach is reversed to begin with a 100% duty cycle on all micromirrors and then decrease the duty cycle of saturated pixels, the DMD cannot perform as well as the global approach for high input levels (2048–4095). This is shown in Fig. 11. Note that the SNR for the lower light levels (0–511) is increased over that of Fig. 10.

Ideally, the initial conditions of the algorithm would be tailored to the expected histogram of the predetected image. If the anticipated histogram were concentrated in the lower range of possible values, the approach shown in Fig. 11 would be preferable. In contrast, if the histogram were expected to contain mostly values in the upper half of possible values (2048–4095), then the approach of Fig. 10 would likely be a more suitable choice. If, however, the expected histogram were contained mostly in the intermediate range of light levels, initial attenuation should be based on that range of light levels. The pixel attenuation in subsequent frames could be increased or decreased accordingly. Beginning the iterative process in a densely populated region of the

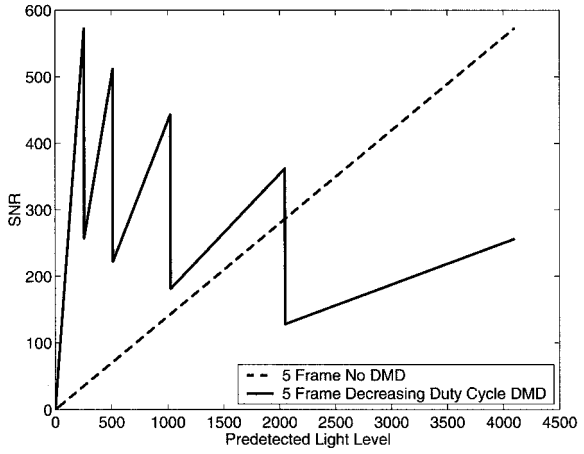

Fig. 11. Plot of DMD sensor contrast enhancement showing SNR for pixel measurements assuming an iterative approach beginning with minimum attenuation and increasing attenuation for pixels that are saturated on each successive frame for five frames. This is compared with a global attenuation of 1/16 with five integrated frames.

histogram offers two potential advantages. First, it increases the SNR of, and therefore the amount of information collected by, a larger number of pixels. The other advantage occurs if the histogram is centrally concentrated. The number of iterations required for implementing the process is then minimized. Figure 12 shows such an approach in which the initial attenuation was for a range of values from 512 to 1023.

It should be noted that this increase in measured light levels is different than dynamic range, because the SNR of the measurement never exceeds the inherent measuring capability of the sensor. Therefore the described procedure enables an 8-bit photodetector array to measure a predetected image with a maximum light level of 4096 times the LSB. However, it does not enable an 8-bit camera to pro-

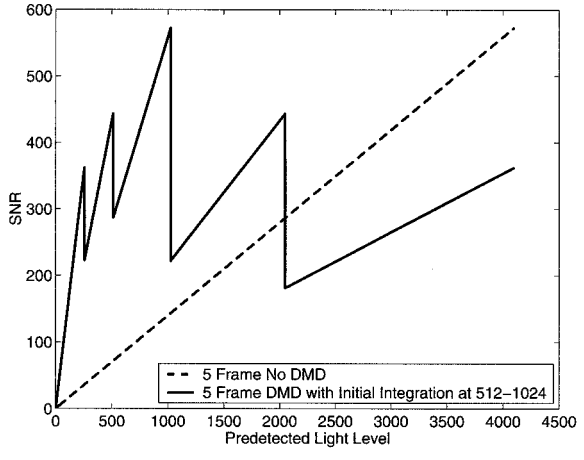

Fig. 12. Plot of DMD sensor contrast enhancement showing SNR for pixel measurements assuming an iterative approach beginning with an intermediate attenuation (appropriate for pixel values between 512 and 1024) and decreasing attenuation for pixels that are beneath the noise floor and increasing attenuation for pixels above saturation on each successive frame for five frames. This is compared to a global attenuation of 1/16 with five integrated frames.

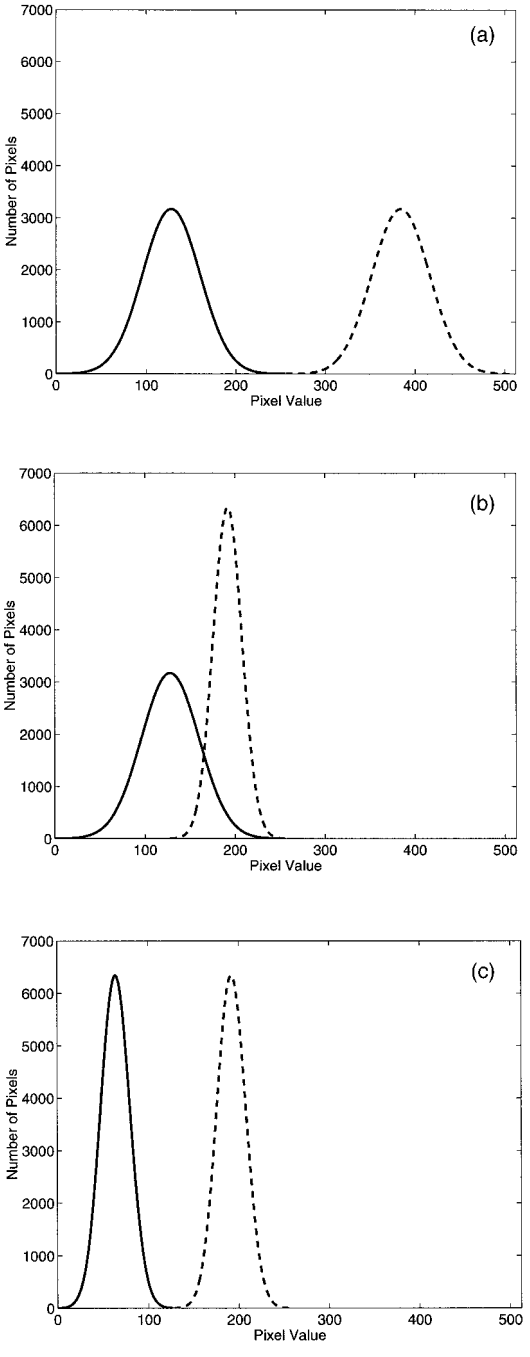

Fig. 13. (a) Bimodal scene histogram for Example 1. The dashed curve corresponds to pixels that will appear saturated when imaging with an 8-bit camera. (b) Modified scene histogram for Example 1 resulting from local attenuation by a factor of 2. Previously saturated pixels are now within the dynamic range of the 8-bit camera, but their SNR has been lowered; previously unsaturated pixels do not experience a reduction in SNR. (c) Modified scene histogram for Example 1 resulting from global attenuation by a factor of 2. Previously saturated pixels are now within the dynamic range of the 8-bit camera, but all pixels in the image now have a reduced SNR.

Table 1. Statistical Pixel and Scene Information Data from Quantitative Examples

	Example				
	1	2a	2	2b	2c
Mean pixel value					
Unsaturated region	128	128	128	128	128
Saturated region	384	3968	3968	3968	3968
% of Pixels saturated	50%	10%	25%	50%	75%
Information recovered (original)	0.50	0.90	0.75	0.50	0.25
Information recovered (using global attenuation)	0.98	0.81	0.81	0.81	0.81
Information recovered (using DMD contrast enhancement) (increasing duty cycle)	0.99	0.97	0.94	0.90	0.85
Information recovered (using DMD contrast enhancement) (decreasing duty cycle)	0.98	0.96	0.90	0.79	0.69
Refer to Figure(s)	13	Not shown	14	Not shown	Not shown

duce an image with 12 bits of gray-scale resolution. It is therefore not the same as having a 12-bit camera for the price of an 8-bit camera. Furthermore, if this approach is extended to large light levels, the uncertainty of the DMD duty cycle will limit the accuracy of the detection. For example, the performance of the DMD controller system used in the above experiments would limit the architecture to powers-of-2 attenuation. A custom device could be more flexible.

5. Quantitative Examples

We can apply the information theoretic model of imaging introduced by Fellgett and Linfoot[5] and extensively applied to sampled imaging systems by Huck *et al.* (Ref. 6, for example) to produce quantitative examples of the procedures described above. These examples are intended to provide a comparative assessment of global attenuation versus the selective local attenuation that can be implemented with the DMD. It is not implied that the procedures described in these examples are optimized. In the following examples we consider specific cases in which the predetected image is described by a bimodal distribution consisting of two Gaussian distributions with different means. We assume that the modes of the histogram represent discrete isoplanatic regions so that the analysis described in Ref. 6, for example, can be applied individually to those regions. Furthermore, we assume an 8-bit imaging system designed so that the diffraction-limited spot size that is due to the optical aperture is matched to an individual photodetector in the fully filled imaging array, indicating an undersampled image. Global attenuation is assumed to be implemented by means other than aperture reduction, e.g., by reduction of integration time or by use of a neutral-density filter, to avoid introducing changes in the modulation transfer function of the system.

In the context of this section, signal is defined in terms of the standard deviation in the predetected image over the isoplanatic regions, rather than in terms of an individual pixel value. We then define the SNR as the ratio between that standard deviation and the standard deviation in photodetector noise. We assume a SNR of 128 in the predetected image for both isoplanatic regions, and the mean size of the spatial detail in the scene is assumed equal to the dimensions of a photodetector. The information density[6] is a measure of the mutual information per pixel in an image. In the following examples we define the average information density as the combined average of the information density over the two isoplanatic regions. An upper limit on the average information density would then correspond to the information density in the absence of saturation and would, after inclusion of the effects of any image summation, represent the maximum available information in the predetected image. We use this upper limit as a normalization reference.

In the first example the predetected image is characterized by the histogram in Fig. 13(a). The two modes are identical in size and shape, indicative of two equally sized isoplanatic regions in the image with similar spatial statistics. The maximum pixel value (normalized to the detector noise level) in the predetected image is 512, exceeding the dynamic range of the imaging system by a factor of 2, thus requiring two frames be collected for implementing the local attenuation. In the first frame, no spatial information is collected from the saturated region, but the unsaturated region is collected and stored. Following a factor-of-2 local attenuation of the saturated region, the histogram in Fig. 13(b) results. The second frame is collected for the entire scene. The previously unsaturated region that was stored from the first frame is summed with the second frame to improve the SNR in that region. Had the image been globally attenuated by a factor of 2, the histogram in Fig. 13(c) would result. Two frames are required for implementing the local attenuation

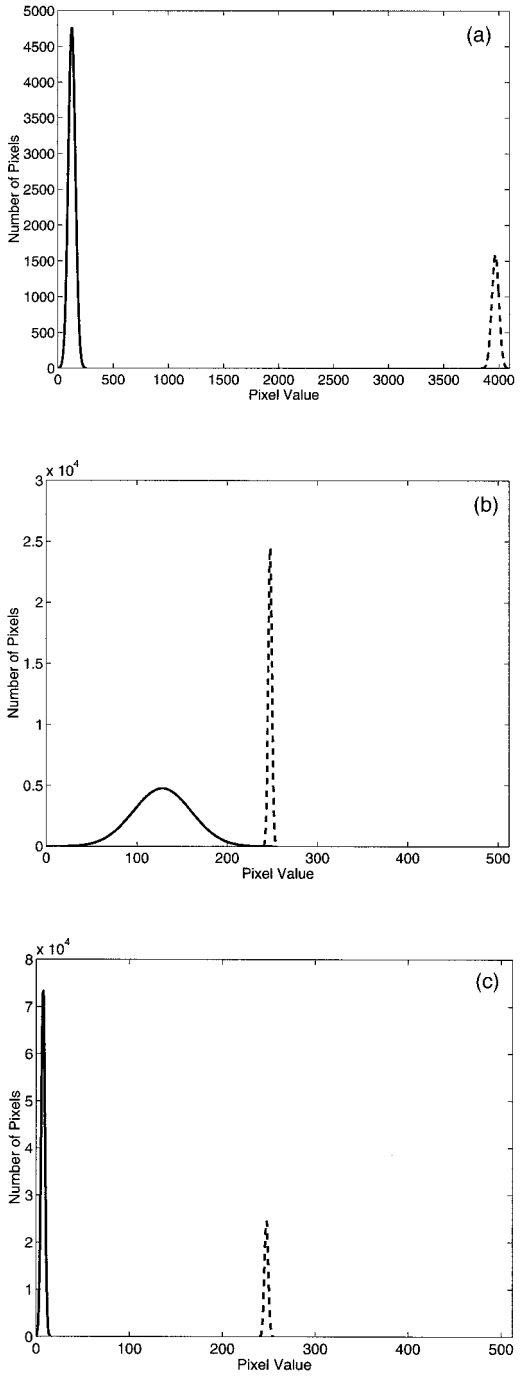

Fig. 14. (a) Bimodal scene histogram for Example 2. The dashed curve represents pixels that will appear saturated when imaging with an 8-bit camera. (b) Modified scene histogram for Example 2 resulting from local attenuation by a factor of 16. Previously saturated pixels are now within the dynamic range of the 8-bit camera, but their SNR has been lowered; previously unsaturated pixels do not experience a reduction in SNR. (c) Modified scene histogram for Example 2 resulting from global attenuation by a factor of 16. Previously saturated pixels are now within the dynamic range of the 8-bit camera, but all pixels in the image now have a reduced SNR.

approach. Therefore two frames are also assumed in the global attenuation case and summed for fair comparison. In each case the average information density over the entire image is determined. The results of the averaged information density for the local attenuation and the global attenuation are both shown in Table 1. All results are normalized to the upper limit on the available information density defined above. Half the available information density is lost in the saturated image. The amount of available information collected by local attenuation is 0.98 compared with 0.98 recovered by global attenuation. For this particular situation, both approaches are equally successful at recovering most of the information lost to saturation.

The second example is illustrated by the histogram in Fig. 14(a). In this case the histogram is still represented by two Gaussians with different means, but the saturated region is one-third the size of the unsaturated region. Furthermore, the maximum light level of the predetected image is much larger than the dynamic range of the imaging system. The example then requires five frames, and the resulting histogram is shown in Fig. 14(b), whereas the globally attenuated histogram is shown in Fig. 14(c). As in the previous example, the globally attenuated case takes advantage of the summation of the five frames that are required for implementing the local attenuation procedure. In this case a quarter of the available information density is lost in the saturated image. Local attenuation recovered part of this with 0.90 of the available information while the globally attenuated image collected only 0.81 of the available information. The local attenuation performed significantly better than global attenuation in this example, but the gain was less significant than in the first example.

In general, with the approaches and assumptions described above, local attenuation and global attenuation appear to perform equivalently well when the saturated and the unsaturated regions occupy equal areas. As the saturated region becomes smaller, or as the maximum pixel value of the predetected image increases, the local attenuation approach improves in comparison with the global attenuation. Alternatively, the algorithm used for this data can be altered as described in Section 4 for iteratively increasing the duty cycle of the pixels that have detected values below the middle of the dynamic range of the sensor. For example, see Fig. 10. Results of these examples for both increasing and decreasing DMD duty-cycle algorithms, along with some additional examples, are summarized in Table 1. As Table 1 indicates, the amount of information gained by means of implementing the local attenuation may not always justify the added complexity. However, the above analysis does not take into account the application-specific importance of the information. Furthermore, it should be noted that the dynamic-range analyses that we have presented would also apply with minor modifications to other pixel-level control schemes such as on-chip CMOS (complementary metal-oxide semiconductor) integration.[7]

6. Conclusion

The ACTIVE-EYES concept, using a DMD for image segmentation and processing of hyperspectral data, has been presented. An experimental prototype using a DMD array to select between a gray-scale camera and an imaging spectrometer has been demonstrated. The resulting architecture was analyzed in terms of the SNR of the measured data and the spatial information that was retained. The analysis showed that the pixel-by-pixel attenuation of the ACTIVE-EYES concept restored spatial detail that would have been lost in a saturated or globally attenuated sensor. The results indicate that the ACTIVE-EYES architecture holds promise for efficiently extracting information in hyperspectral-imaging applications.

However, in the case of saturated pixels the importance of the additional information to the specific application must be carefully considered. Future research may focus on the integration of the experimental prototype and deriving optimum algorithms in an information-efficiency sense.

The authors gratefully acknowledge the support of the Defense Advanced Research Projects Agency (DARPA) through contract F33615-00-1-1625, which is monitored by the Air Force Research Laboratory (AFRL).

References

1. J. Castracane and M. Gutin, "DMD-based bloom control for intensified imaging systems," in *Diffractive and Holographic Technologies, Systems, and Spatial Light Modulators IV*, I. Cindrich, S. H. Lee, and R. L. Sutherland, eds., Proc. SPIE **3633,** 234–242 (1999).
2. K. Kearney, M. Corio, and Z. Ninkov, "Imaging spectroscopy with digital micromirrors," in *Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications*, M. M. Blouke, N. Sampat, G. M. Williams, and T. Yeh, eds., Proc. SPIE **3965,** 11–20 (2000).
3. K. Kearney and Z. Ninkov, "Characterization of a digital micromirror device for use as an optical mask in imaging and spectroscopy," in *Spatial Light Modulators*, R. L. Sutherland, ed., Proc. SPIE **3292,** 81–92 (1998).
4. Rochester Microsystems, 400 Air Park Drive, Suite 60, Rochester, New York.
5. P. B. Fellgett and E. H. Linfoot, "On the assessment of optical images," Philos. Trans. R. Soc. London **247,** 269–407 (1955).
6. F. O. Huck, C. L. Fales, and Z. Rahman, "An information theory of visual communication," Philos. Trans. R. Soc. London Ser. A **354,** 2193–2248 (1996).
7. V. M. Brajovic, R. Miyagawa, and T. Kanade, "Temporal photoreception for adaptive dynamic range image sensing and encoding," Neural Netw. **11,** 1149–1158 (1998).