Programmable Imaging using a Digital Micromirror Array[*]

Shree K. Nayar and Vlad Branzoi
Department of Computer Science
Columbia University
New York, NY 10027
{nayar, vlad}@cs.columbia.edu

Terry E. Boult
Department of Computer Science
University of Colorado
Colorado Springs, Colorado 80933
tboult@cs.uccs.edu

Abstract

In this paper, we introduce the notion of a programmable imaging system. Such an imaging system provides a human user or a vision system significant control over the radiometric and geometric characteristics of the system. This flexibility is achieved using a programmable array of micro-mirrors. The orientations of the mirrors of the array can be controlled with high precision over space and time. This enables the system to select and modulate rays from the light field based on the needs of the application at hand.

We have implemented a programmable imaging system that uses a digital micro-mirror device (DMD), which is used in digital light processing. Although the mirrors of this device can only be positioned in one of two states, we show that our system can be used to implement a wide variety of imaging functions, including high dynamic range imaging, feature detection, and object recognition. We conclude with a discussion on how a micro-mirror array can be used to efficiently control field of view without the use of moving parts.

1 A Flexible Approach to Imaging

In the past few decades, a wide variety of novel imaging systems have been proposed that have fundamentally changed the notion of a camera. These include high dynamic range, multispectral, omnidirectional, and multi-viewpoint imaging systems. The hardware and software of each of these devices are designed to accomplish a particular imaging function and this function cannot be altered without significant redesign.

In this paper, we introduce the notion of a programmable imaging system. Such a system gives a human user or a computer vision system significant control over the radiometric and geometric properties of the system. This flexibility is achieved by using a programmable array of micro-mirrors. The orientations of the mirrors of the array can be controlled with very high speed. This enables the system to select and modulate scene rays based on the needs of the application at hand. The end result is a single imaging system that can emulate the functionalities of several existing specialized systems as well as new ones.
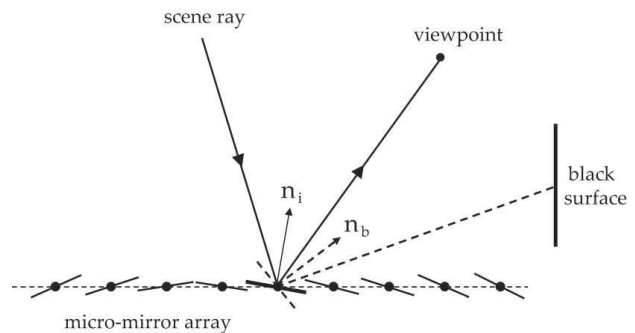
Figure 1: The principle underlying programmable imaging using a micro-mirror array. If the orientations of the individual mirrors can be controlled with high precision and speed, scene rays can be selected and modulated in a variety of ways, each leading to a different imaging system.

The basic principle behind the proposed approach is illustrated in Figure 1. The system observes the scene via a two-dimensional array of micro-mirrors, whose orientations can be controlled. The surface normal $\mathbf{n}_i$ of the i^{th} mirror determines the direction of the scene ray it reflects into the imaging system. If the normals of the mirrors can be arbitrarily chosen, each mirror can be programmed to select from a continuous cone of scene rays. In addition, each mirror can also be oriented with normal $\mathbf{n}_b$ such that it reflects a black surface (with zero radiance). Let the integration time of the image detector be T. If the mirror is made to point in the directions $\mathbf{n}_i$ and $\mathbf{n}_b$ for durations t and $T - t$, respectively, the scene ray is attenuated by t/T. As a result, each imaged scene ray can also be radiometrically modulated with high precision.

Since the micro-mirror array is programmable, the above geometric and radiometric manipulations can be varied to create a wide range of transformations of the light field of the scene to captured images. We will show that the radiometric modulation enables us to achieve several functions including high dynamic range imaging, optical feature detection, and object recognition using appearance matching. In addition, we show how the orientations of the mirrors can be programmed to control the field of view and resolution of an imaging system.

2 Imaging with a Micromirror Device

Ideally, we would like to have full control over the orientations of our micro-mirrors. Such devices are being developed for adaptive optical processing in astronomy
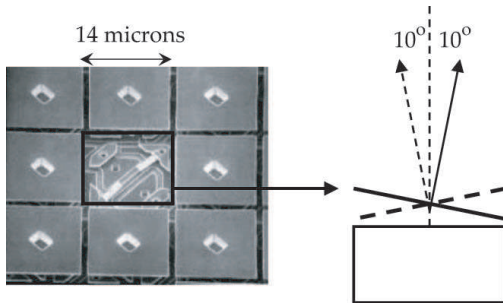
Figure 2: Our implementation of programmable imaging uses a digital micro-mirror device (DMD). Each mirror of the device is roughly 14×14 microns and can be oriented with high precision and speed at $+10$ or -10 degrees.

[20]. However, at this point in time, they do not have the physical properties and programmability that we need for our purpose. To implement our ideas, we use the digital micro-mirror device (DMD) that was introduced in the 1980s by Hornbeck at Texas Instruments [7],[8]. The DMD is a micro-electro-mechanical system (MEMS) that has evolved rapidly over the last decade and has found many applications [3]. It is the key enabling technology in many of today's projection systems [9]. The latest generation of DMDs have more than a million mirrors, each mirror roughly 14×14 microns in size (see Figure 2). From our perspective, the main limitation of current DMDs is that the mirrors can only be oriented in one of two directions; $-10°$ or $+10°$ about one of the mirror's diagonal axes (see Figure 2). However, the orientation of the mirror can be switched from one state to the other in a few *microseconds*, enabling modulation of incident light with very high precision.

Figure 3 shows the optical layout of the system we have developed using the DMD. The scene is first projected onto the DMD plane using an imaging lens. This means that the cone of light from each scene point received by the aperture of the imaging lens is focused onto a single micro-mirror. When all the mirrors are oriented at $+10°$, the light cones are reflected in the direction of a re-imaging lens which focuses the image received by the DMD onto a CCD image detector. Note that the DMD in this case behaves like a planar scene that is tilted by $20°$ with respect to the optical axis of the re-imaging lens. To produce a focused image of this tilted set of source points, one needs to tilt the image detector according to the well-known Scheimpflug condition [18].

3 Prototype System

It is only recently that developer kits have begun to appear that enable one to use DMDs in different applications. When we began implementing our system this option was not available. Hence, we chose to re-engineer an off-the-shelf DMD projector into an imaging system by reversing the path of light; the projector lens is used to form an image of the scene on the DMD rather than illuminate the scene via the DMD. Figure 4(a) shows a
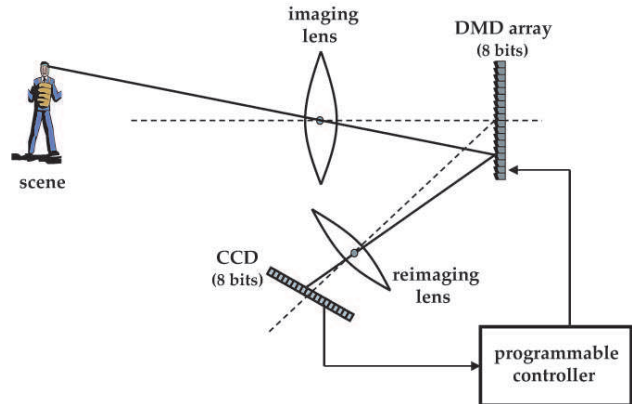

Figure 3: Imaging using a DMD. The scene image is focused onto the DMD plane. The image reflected by the DMD is re-imaged onto a CCD. The programmable controller captures CCD images and outputs DMD (modulation) images.
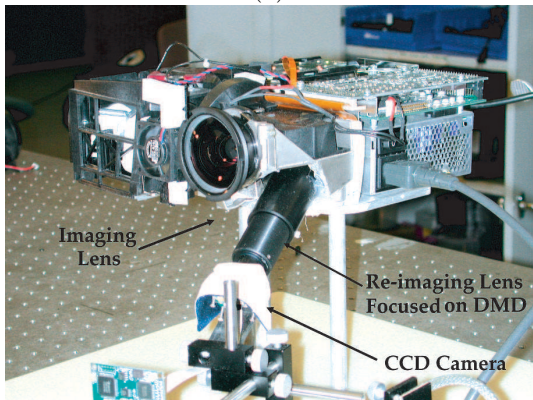
partly disassembled Infocus LP 400 projector. This projector uses one of the early versions of the DMD with 800×600 mirrors, each 17×17 microns in size. The modulation function of the DMD is controlled by simply applying an 8-bit image (VGA signal) to the projector input. We had to make significant hardware changes to the projector. First, the projector lamp had to be blocked out of the optical path. Then, the chassis had to be modified so that the re-imaging lens and the camera could be attached to the system. The final system is shown in Figure 4(b). The CCD camera used is an 8-bit monochrome Sony XC-75 model with 640×480 pixels. The processing of the camera image and the control of DMD image is done using a Dell workstation with a 2.5 GHz Pentium 4 processor.

DMDs have previously been used in imaging applications, but for very specific tasks such as recording celestial objects in astronomy. For instance, in [12] the DMD is used to mask out bright sources of light (like the sun) so that dimmer regions (the corona of the sun) can be imaged with higher dynamic range. In [11], a DMD is used to mask out everything except a small number of stars. Light from these unmasked stars are directed towards a spectroscope to measure the spectral characteristics of the stars. In [2] and [1], the DMD is used to modulate brightness values at a pixel level for high dynamic range multispectral imaging and removal of image blooming, respectively. These works address rather specific imaging needs. In contrast, we are interested in a flexible imaging system that can perform a wide range of functions.

An extensive off-line calibration of the geometric and radiometric properties of the system was conducted (see [15] for details). The geometric calibration involves determining the mapping between DMD and CCD pixels. This mapping is critical to controlling the DMD and interpreting the images captured by the CCD. The geometric calibration was done by using a bright scene
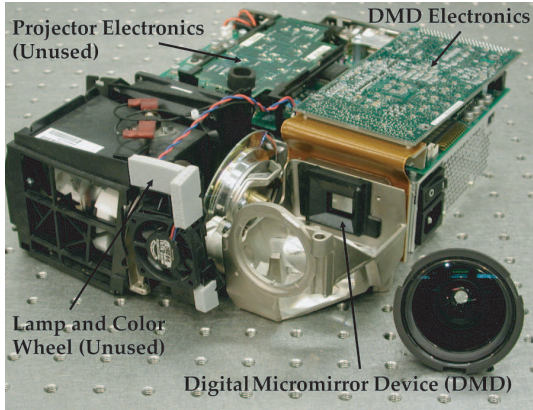
(a)

(b)

Figure 4: (a) A disassembled Infocus LP 400 projector that shows the exposed DMD. (b) In this re-engineered system, the projector lens is used as an imaging lens that focuses the scene on the DMD. The image reflected by the DMD is re-imaged by a CCD camera.

with uniform brightness. Then, a large number of square patches were used as input to the DMD and recorded using the CCD, as shown in Figure 5. In order to scan the entire set of patches efficiently, binary coding of the patches was done. The centroids of corresponding patches in the DMD and CCD images were fit to a piece-wise, first-order polynomial. The computed mapping was found to have an RMS error of 0.6 (CCD) pixels.

Figure 6 shows two simple examples that illustrate the modulation of scene images using the DMD. One can see that after modulation some of the scene regions that were previously saturated produce useful brightness values. Note that the captured CCD image is skewed with respect to the DMD modulation image. This skewing is due to the required tilt of the CCD discussed above and can be undone using the calibration result. In our system, the modulation image can be controlled with 8 bits of precisions and the captured CCD images have 8 bits of accuracy. Hence, the measurements made at each pixel have roughly 16 bits of information.

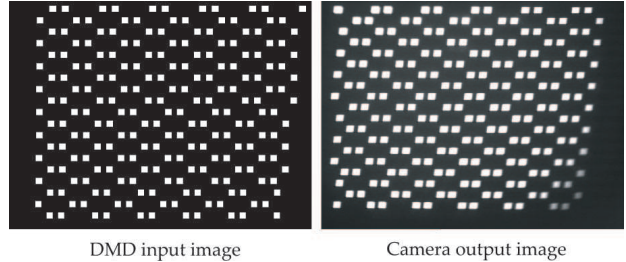DMD input image   Camera output image

Figure 5: Geometric calibration of the imaging system. The geometric mapping between the DMD and the CCD camera is determined by applying patterned images to the DMD and capturing the corresponding CCD images.
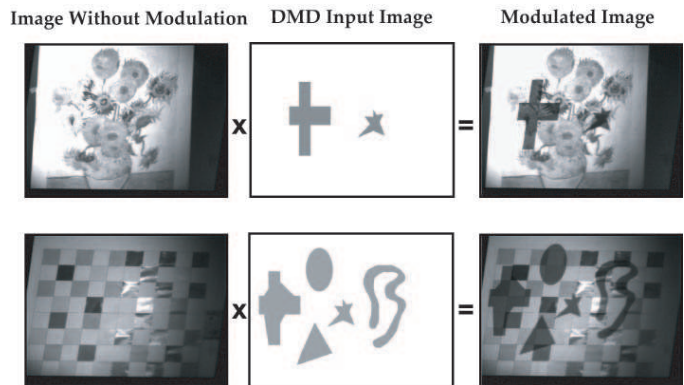

Figure 6: Examples that show how image irradiance is modulated with high resolution using the DMD.

4 High Dynamic Range Imaging

The ability to program the modulation of the image at a pixel level provides us a flexible means to implement several previously proposed methods for enhancing the dynamic range. In this section, we will describe three different implementations of high dynamic range imaging using the programmable imaging system.

4.1 Temporal Exposure Variation

We begin with the simplest implementation, where the global exposure of the scene is varied as a function of time. In this case, the control image applied to the DMD is spatially constant but changes periodically with time. An example of a video sequence acquired in this manner is shown in Figure 7, where 4 modulation levels are cycled over time. It has been shown in previous work that an image sequence acquired in this manner can be used to compute high dynamic range video when the motions of scene points between subsequent frames is small [4] [10]. Alternatively, the captured video can be subsampled in time to produce multiple video streams with lower frame-rate, each with a different fixed exposure. Such data can improve the robustness of tasks such as face recognition, where a face missed at one exposure may be better visible and hence detected at another exposure.
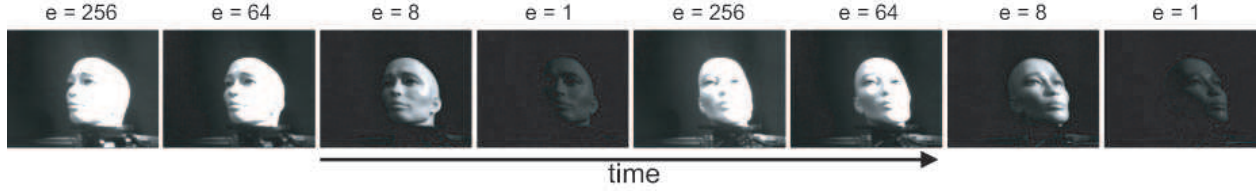
e = 256   e = 64   e = 8   e = 1   e = 256   e = 64   e = 8   e = 1

time

Figure 7: Spatially uniform but temporally varying DMD inputs can be used to generate a video with varying exposure (**e**). Using a DMD in this case produces high quality data compared to changing the exposure time or the camera gain.

Videos of the type shown in Figure 7 can also be obtained by changing the integration time of the detector or the gain of the camera. Due to the various forms of camera noise, changing integration time or gain compromises the quality of the acquired data. In our case, since the DMD can be controlled with 8 bits of accuracy and the CCD camera produces 8-bit images, the captured sequence can be controlled with 16 bits of precision.

4.2 Spatio-Temporal Exposure Variation

In [16], the concept of spatially varying pixel exposures was proposed, where an image is acquired with a detector with a mosaic of neutral density filters. The captured image can be reconstructed to obtain a high dynamic range image with a slight loss in spatial resolution. Our programmable system allows us to capture an image with spatially varying exposures by simply applying a fixed (checkerboard-like) pattern to the DMD. In [17], it was shown that a variety of exposure patterns can be used, each trading-off dynamic range and spatial resolution in different ways. Such trade-offs are easy to explore using our system.

It turns out that spatially varying exposures can also be used to generate video streams that have higher dynamic range for a human observer, without post-processing each acquired image as was done in [16]. If one uses a fixed pattern, the pattern will produce a very visible modulation that would be distracting to the observer. However, if the pattern is varied with time, the eye becomes less sensitive to the pattern and a video with a larger range of brightnesses is perceived by the observer. Figure 8(a) shows the image of a scene taken without modulation. It is clear that the scene has a wide dynamic range and an 8-bit camera cannot capture this range. Figure 8(b) shows four consecutive frames captured with spatially varying exposures. The exposure pattern uses 4 different exposures (e_1, e_2, e_3, e_4) within each 2×2 neighborhood of pixels. The relative positions of the 4 exposures are changed over time using a cyclic permutation. In the images shown in Figure 8(b), one sees the spatial patterns introduced by the exposures (see insets). However, when this sequence is viewed at 30 Hz, the pattern is more or less invisible (the eye integrates over the changes) and a wider range of brightnesses are visible.
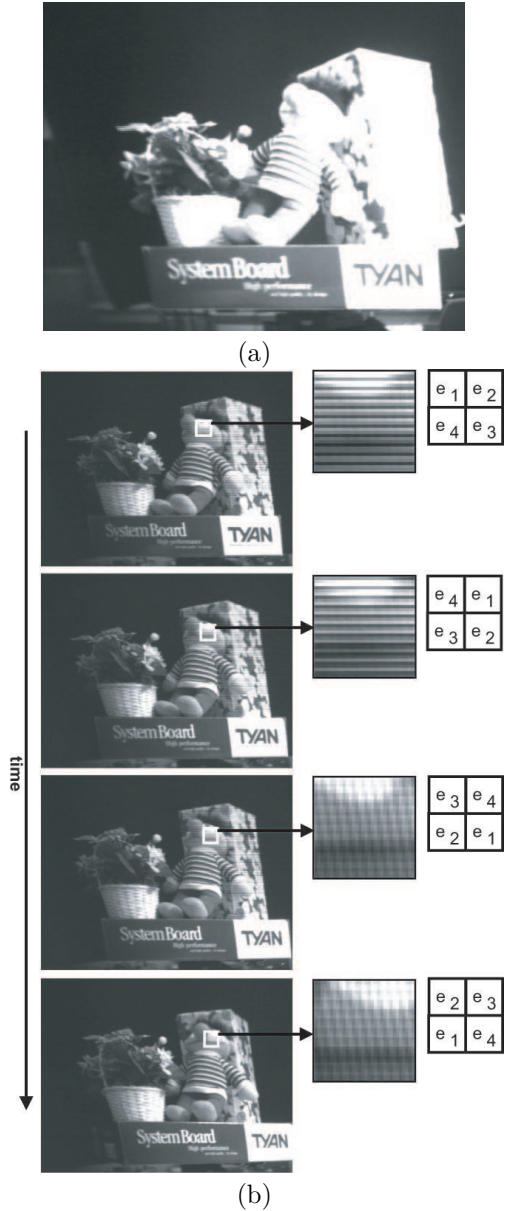

(a)

(b)

Figure 8: (a) A scene with a wide range of brightnesses captured using an 8-bit (low dynamic range) camera. (b) Four frames of the same scene captured with spatio-temporal exposure modulation using the DMD. When such a video is viewed at frame-rate, the observer perceives a wider dynamic range without noticing the exposure changes.
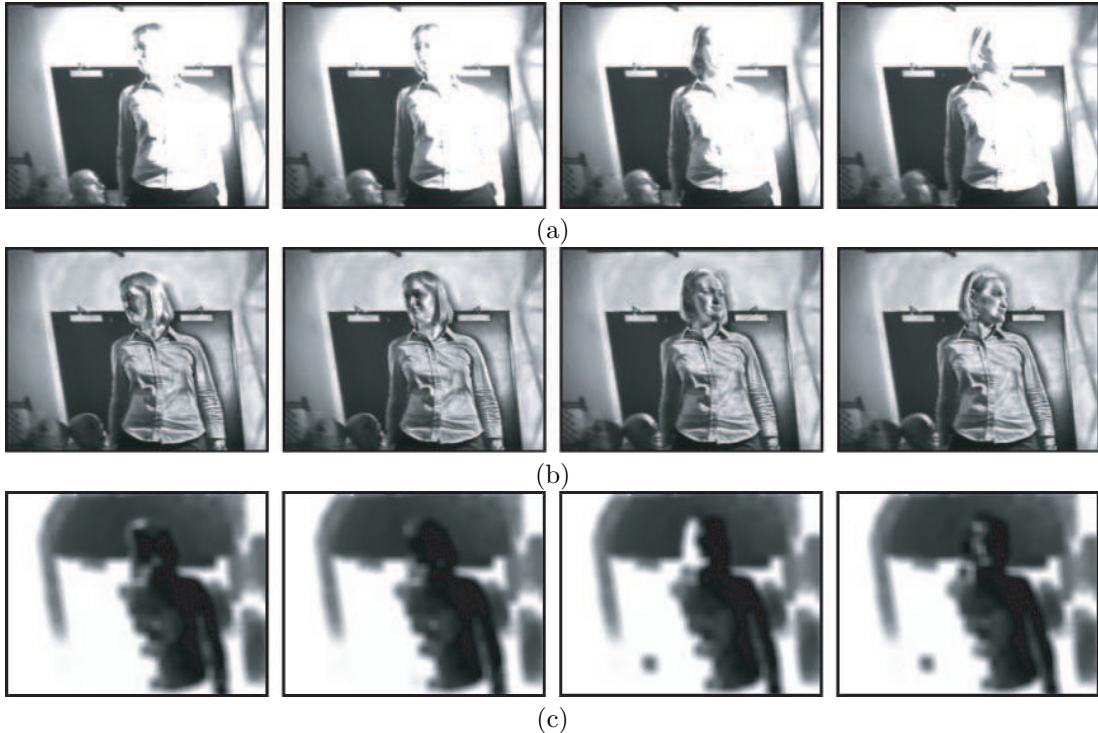
Figure 9: (a) Video of a person taken under harsh lighting using a conventional (8-bit) camera. (b) The raw output of the programmable system when the DMD is used to achieve adaptive dynamic range. (c) The modulation images applied to the DMD. The raw camera output and the DMD modulation can be used to compute a video with very high dynamic range.

4.3 Adaptive Dynamic Range

Recently, the method of adaptive dynamic range was introduced in [14], where the exposure of each pixel is controlled based on the scene radiance measured at the pixel. A prototype device was implemented using an LCD attenuator attached to the front of the imaging lens of a camera. This implementation suffers from three limitations. First, since the LCD attenuator uses polarization filters, it allows only 50% of the light from the scene to enter the imaging system even when the attenuation is set to zero. Second, the attenuation function is optically defocused by the imaging system and hence pixel-level attenuation could not be achieved. Finally, the LCD attenuator cells produce diffraction effects that cause the captured images to be slightly blurred.

The DMD-based system enables us to implement adaptive dynamic range imaging without any of the above limitations. Since the image of the scene is first focused on the DMD and then re-imaged onto the image detector, we are able to achieve pixel-level control. In addition, the fill-factor of the DMD is very high compared to an LCD array and hence the optical efficiency of the modulation is closer to 90%. Because of the high fill-factor, the blurring/diffraction effects are minimal. In [2] and [1], a DMD has been used to address this problem, but these previous works do not adequately address the real-time spatio-temporal control issues that arise in the case of dynamic scenes. We have implemented a con-

trol algorithm very similar to that described in [14] for computing the DMD modulation function based on each captured image. Results from this system are shown in Figure 9. The first row shows a person under harsh lighting imaged without modulation (conventional camera). The second row shows the output of the camera and the third row shows the corresponding modulation (attenuation) images applied to the DMD. As described in [14], the modulation image and the acquired image can be used to compute a video stream that has an effective dynamic range of 16 bits.

5 Intra-Pixel Optical Feature Detection

The field of optical computing has developed very efficient and powerful ways to apply image processing algorithms (such as convolution and correlation) [5]. A major disadvantage of optical computing is that it requires the use of *coherent* light to represent the images. This has proven cumbersome, bulky, and expensive. It turns out that programmable modulation can be used to implement a special class of image processing tasks directly to the *incoherent* optical image captured by the imaging lens, without the use of any coherent sources. In particular, one can apply convolution at an *intra-pixel* level very efficiently. By intra-pixel we mean that the convolution mask is being applied to the distribution of light energy within a single pixel rather than a neighborhood of pixels. Intra-pixel optical processing leads to very efficient algorithms for finding features such as edges, lines, and corners.

Consider the convolution $f * g$ of a continuous optical image f with a kernel g whose span (width) is less than, or equal to, a pixel on the image detector. We can rewrite the convolution as $f * (g^+ - g^-)$ where g^+ is made up of only the positive elements of g and g^- has only the absolute of the negative elements of g. We use this decomposition since incoherent light cannot be negatively modulated (the modulation image cannot have negative values). An example of such a decomposition for the case of a first-derivative operator is shown in Figure 10(a). As shown in the figure, let each CCD pixel correspond to 3×3 DMD pixels; i.e. the DMD has three times the linear resolution of the CCD. Then, the two components of the convolution (due to g^+ and g^-) are directly obtained by capturing two images with the modulation images shown in Figure 10(b). The difference between these images gives the final result $(f * g)$.

Figure 10(c) shows the four *optically processed* images of a scene obtained for the case of the Sobel edge operator. The computed edge map is shown in Figure 10(d). Since our DMD has only 800×600 elements, the edge map is of lower resolution with about 200×150 pixels. Although four images are needed in this case, it can be applied to a scene with slowly moving objects where each new image is only used to update one of the four component filter outputs in the edge computation. Note that all the multiplications involved in the convolutions are done in optical domain (at the speed of light).

6 Optical Appearance Matching

In the past decade, appearance matching using subspace methods has become a popular approach to object recognition [19][13]. Most of these algorithms are based on projecting input images to a precomputed linear subspace and then finding the closest database point that lies in the subspace. The projection of an input image requires finding its dot product with a number of vectors. In the case of principal component analysis, the vectors are the eigenvectors of a correlation or covariance matrix computed using images in the training set.

It turns out that optical modulation can be used to perform all the required multiplications in optical domain, leaving only additions to be done computationally. Let the input image be m and the eigenvectors of the subspace be $e_1, e_2, ...e_k$. The eigenvectors are concatenated to obtain a larger (tiled) vector $B = [e_1, e_2, ...e_k]$ and k copies of the input image are concatenated to obtained the (tiled) vector $A = [m, m, ...m]$. If the vector A is "shown" as the scene to our imaging system and the vector B is used as the modulation image, the image captured by the camera is a vector $C = A.*B$, where $.*$ denotes an element-by-element product of the two vectors. Then, the image C is raster scanned to sum up its k tiles to obtain the k coefficients that correspond to the subspace projection of the input image. This coefficient vector is compared with stored vectors and the closest match reveals the identity of the object in the image.
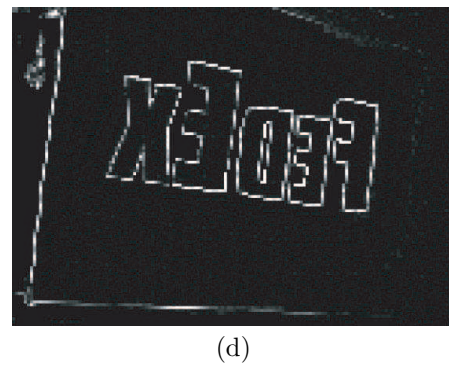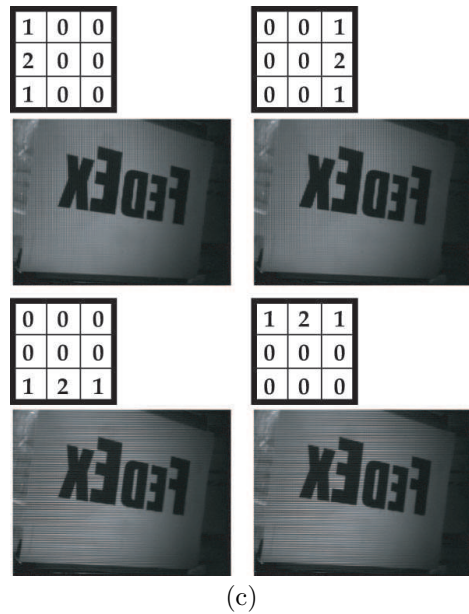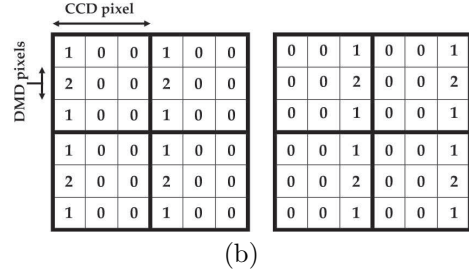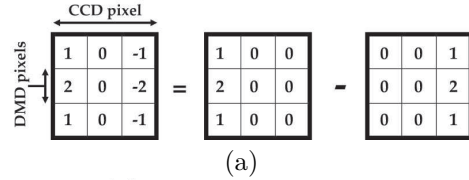
(a)

(b)

(c)

(d)

Figure 10: (a) Decomposition of a convolution kernel into two positive component kernels. (b) When the resolution of the DMD is higher than that of the CCD, intra-pixel convolution is done by using just two modulation images and subtracting the resulting CCD images. (c) Four images that result from applying the four component kernels of a Sobel edge operator. (d) The edge map computed from the four images in (c).
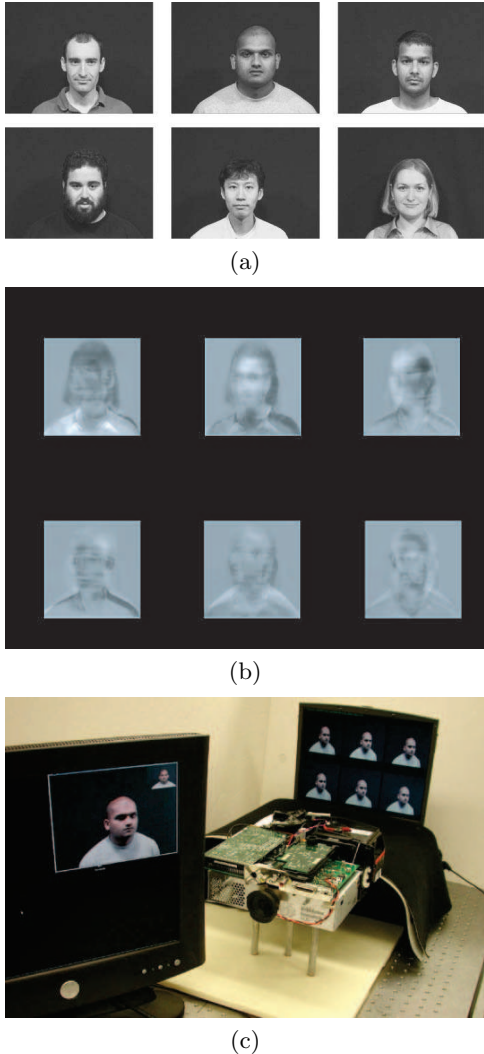
(a)

(b)

(c)

Figure 11: (a) People used in the database of the recognition system (different poses of each person are included). (b) The 6 most prominent eigenvectors computed from the training set, tiled to form the modulation image. (c) A tiling of the input (novel) image is "shown" to the imaging system. Simple summation of brightness values in the captured image yields the coefficients needed for recognition.

We have used our programmable system to implement this idea and develop a real-time face recognition system. Figure 11(a) shows the 6 people in our database; 30 poses (images) of each person were captured to obtain a total of 180 training images. PCA was applied and the 6 most prominent eigenvectors are tiled as shown in Figure 11(b) and used as the DMD modulation image. During recognition, the output of the video camera is also tiled in the same way as the eigenvectors and displayed on a screen that sits in front of the imaging system, as shown in Figure 11(c). The 6 parts of the captured image are summed to obtain the 6 coefficients. A simple nearest-neighbor algorithm is applied to these coefficients to recognize the person in the input image.

7 Programmable Field of View

Thus far, we have mainly exploited the radiometric flexibility made possible by our imaging system. The use of a programmable micro-mirror array also allows us to very quickly alter the field of view and resolution characteristics of an imaging system[1]. Quite simply, a planar array of mirrors can be used to emulate a *deformable mirror* whose shape can be changed almost *instantaneously*.

To illustrate this idea, we do not use the imaging system in Figure 4 as the imaging optics would have to be substantially altered to facilitate field of view manipulation. Instead, we consider the case where the micro-mirror array does not have an imaging lens that focuses the scene onto it but instead directly reflects the scene into the camera optics. This scenario is illustrated in Figure 12(a), where the micro-mirror array is aligned with the horizontal axis and the viewpoint of the camera is located at the point P at height h from the array.

If all the mirrors are parallel to the horizontal axis, the array behaves like a planar mirror and the viewpoint of the system is simply the reflection P' of the point P. The field of view in this case is the solid angle subtended by the mirror array from P. Now consider the micro-mirror located at distance d from the origin to have tilt θ with the horizontal axis, as shown in Figure 12(a). Then, the angle of the scene ray imaged by this mirror is $\phi = 2\theta + \alpha$, where $\alpha = \tan^{-1}(d/h)$. It is also easy to show that the viewpoint of the system corresponding to this particular mirror element has the (x and y) coordinates $Q_x(d) = d - \sqrt{(h^2 + d^2)}\cos\beta$ and $Q_y(d) = -\sqrt{(h^2 + d^2)}\sin\beta$, where $\beta = (\pi/2) - \phi$. If all the micro-mirrors have the same tilt angle θ, then the system has a *locus of viewpoints* (caustic) that is a segment of the line that passes through P and Q.

If the mirrors of the array can be controlled to have any orientation within a continuous range, one can see that the field of view of the imaging system can be varied over a wide range. For instance, if the mirrors at the two end-points of the array (at $-a$ and b) have orientations $-\theta$ and θ, and the orientations of mirrors in between vary smoothly between these two values, the field of view of the camera is enhance by 4θ. As we mentioned, the DMDs that are currently available can only have one of two mirror orientations (+10 or −10 degrees). Therefore, if all the mirrors are initially inactive (0 degrees) and then oriented at 10 degrees, the field of view remains the same but its orientation changes by 20 degrees. This very case is illustrated by the images shown in Figure 12(b), where the left image shows one view of a printed sheet of paper and the right one shows a rotated view of the same.

One can see that both the images are blurred. This is be-

---

[1]This approach to controlling field of view using a mirror array is also being explored by Andrew Hicks at Drexel University [6].
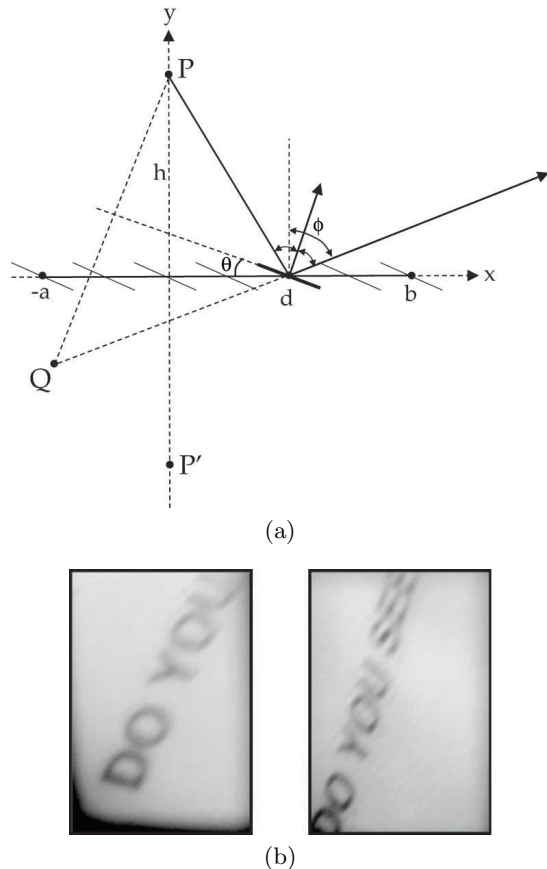
(a)

(b)

Figure 12: (a) The field of view of an imaging system can be controlled almost instantly by using a micro-mirror array. Here, the scene is being reflected directly by the array into the viewpoint of the camera. (b) Two images of the same scene taken via a DMD with all the mirrors at 0 degrees (left image) and 10 degrees (right image).

cause we are imaging the scene directly through a DMD without using a re-imaging lens and hence many mirrors lie within the light cone that is imaged by a single pixel. Since the mirrors are tilted, the surface discontinuities at the edges of the mirrors cause diffraction effects. These effects can be minimized if one has greater control over the orientations of the micro-mirrors. We expect such devices to become available in the future.

8   Conclusion

Programmable imaging using a micro-mirror array is a general and flexible approach to imaging. We have shown that this approach enables one to substantially alter the design of an imaging system with ease. We believe this concept is timely. Significant advances are being made in MEMS technology that are expected to have direct impact on the next generation of digital micromirror arrays. When micro-mirror arrays allow greater control over the orientations of individual mirrors, programmable imaging can impact imaging applications in many fields of science and engineering.

References

[1] J. Castracane and M. A. Gutin. DMD-based bloom control for intensified imaging systems. In *Diffractive and Holographic Tech., Syst., and Spatial Light Modulators VI*, volume 3633, pages 234–242. SPIE, 1999.

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

[3] D. Dudley, W. Duncan, and J. Slaughter. Emerging digital micromirror device (DMD) applications. White paper, Texas Intruments, February 2003.

[4] R. Ginosar, O. Hilsenrath, and Y. Zeevi. Wide dynamic range camera. U.S. Patent 5,144,442, September 1992.

[5] J. W. Goodman. *Introduction to Fourier Optics*. McGraw-Hill, New York, 1968.

[6] R. A. Hicks. *Personal Communication*. 2003.

[7] L.J. Hornbeck. Bistable deformable mirror device. In *Spat. Light Mod. and Apps.*, volume 8. OSA, 1988.

[8] L.J. Hornbeck. Deformable-mirror spatial light modulators. In *Projection Displays III*, volume 1150, pages 86–102. SPIE, August 1989.

[9] L.J. Hornbeck. Projection displays and MEMS: Timely convergence for a bright future. In *Micromachined Devices and Components*, volume 2642. SPIE, 1995.

[10] S. B. Kang, M. Uyttendaele, S. Winder, and R. Szeliski. High dynamic range video. *ACM Trans. on Graph.*, 22(3):319–325, 2003.

[11] K. J. Kearney and Z. Ninkov. Characterization of digital micromirror device for use as an optical mask in imaging and spectroscopy. In *Spatial Light Modulators*, volume 3292, pages 81–92. SPIE, April 1998.

[12] F. Malbet, J. Yu, and M. Shao. High dyn. range imaging using a deformable mirror for space coronography. *Public. of the Astro. Soc. of the Pacific*, 107:386, 1995.

[13] H. Murase and S. K. Nayar. Visual Learning and Recognition of 3D Objects from Appearance. *International Journal of Computer Vision*, 14(1):5–24, January 1995.

[14] S. K. Nayar and V. Branzoi. Adaptive Dynamic Range Imaging: Optical Control of Pixel Exposures over Space and Time. *Proc. of International Conference on Computer Vision (ICCV)*, pages 1168–1175, 2003.

[15] S. K. Nayar, V. Branzoi, and T. E. Boult. A programmable imaging system. *Tech. Rep., Computer Science Dept., Columbia Univ., (in preparation)*, 2004.

[16] S. K. Nayar and T. Mitsunaga. High dynamic range imaging: Spatially varying pixel exposures. *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition 2000*, 1:472–479, June 2000.

[17] S. K. Nayar and S. G. Narasimhan. Assorted Pixels: Multisampled Imaging with Structural Models. *Proc. of Euro. Conf. on Comp. Vis. (ECCV)*, 4:636–152, 2002.

[18] W. J. Smith. *Mod. Opt. Eng.* McGraw-Hill, 1966.

[19] M. Turk and A. P. Pentland. Face recognition using eigenfaces. *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 586–591, 1991.

[20] R. K. Tyson. *Principles of Adaptive Optics*. Academic Press, 1998.