

Adaptive Dynamic Range Imaging: Optical Control of Pixel Exposures Over Space and Time

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

Abstract

This paper presents a new approach to imaging that significantly enhances the dynamic range of a camera. The key idea is to adapt the exposure of each pixel on the image detector, based on the radiance value of the corresponding scene point. This adaptation is done in the optical domain, that is, during image formation. In practice, this is achieved using a spatial light modulator whose transmittance can be varied with high resolution over space and time. A real-time control algorithm is developed that uses acquired images to automatically adjust the transmittance function of the spatial modulator. Each captured image and its corresponding transmittance function are used to compute a very high dynamic range image that is linear in scene radiance.

We have implemented a video-rate adaptive dynamic range camera that consists of a color CCD detector and a controllable liquid crystal light modulator. Experiments have been conducted in scenarios with complex and harsh lighting conditions. The results indicate that adaptive imaging can have a significant impact on vision applications such as monitoring, tracking, recognition, and navigation.

1 High Dynamic Range Imaging

A typical real-world scene has an enormous range of brightness (radiance) values. Digital cameras used in computer vision and computer graphics generally provide 8 bits of brightness information per color channel, at each pixel. That is, within each color channel, all radiance values in the scene are mapped to one of 256 image brightness values, irrespective of the complexity of the scene. Computer vision problems, such as the perception of shape, the estimation of motion, and the recognition of objects, are often under-constrained and thus inherently hard to solve. The low quality of image data produced by today's cameras makes each of these problems even harder. This fact is now widely accepted; the problem of acquiring "high dynamic range" (HDR) images has drawn the attention of researchers in vision and graphics.

This paper presents a novel approach to HDR imaging, called adaptive imaging, that produces unprecedented dynamic ranges. Before presenting this approach, a

quick review of previous work is in order. All of the previous methods seek to improve the dynamic range of an imaging system either by trading-off temporal or spatial resolution of the image detector, or by using multiple detectors.

Temporal Exposure Change: The most popular approach to HDR imaging is to sequentially capture multiple images of the scene using different exposures (for examples, see [7], [3], [11], [22], [18].) Low exposures provide useful data in bright scene regions and high exposures in dark scene regions. Therefore, the acquired images can be fused to obtain a single high quality image. In [12], [4], [13], and [17], this approach is taken a step further; the radiometric response of the imaging system is also recovered from the acquired images.

The above methods are of course suited only to static scenes; the imaging system, the scene objects, and their radiances must remain constant during the sequential capture of images under different exposures.

Mosaicing with Spatially Varying Filter: Recently, the concept of generalized mosaicing [19], [1] was introduced where a spatially varying neutral density filter is rigidly attached to the camera. When this imaging system is rotated, each scene point is observed under different exposures. The captured images are fused to obtain a wide-angle mosaic that has very high dynamic range. As with temporal exposure change, the disadvantage here is that the scene must remain constant while the imaging system rotates.

Multiple Image Detectors: The scene/camera motion restrictions faced by the above methods can be remedied by using multiple imaging systems. This approach has been taken by several investigators (for examples, see [5], [18]). Beam splitters are used to generate multiple copies of the optical image of the scene. Each copy is measured by an image detector whose exposure is preset by using an optical attenuator or by adjusting the exposure time of the detector.

This approach has the advantage of producing HDR images in real time. The disadvantage is that it is more expensive as it requires multiple detectors and precision optics for good alignment between the acquired images.

Multiple Sensor Elements Within a Pixel: A rather novel approach to HDR imaging uses a custom

detector (see [8] and [9] for examples) where each detector cell includes two sensing elements (potential wells) of different sizes (and hence sensitivities). When the detector is exposed to the scene, two measurements are made within each cell and they are combined on-chip before the image is read out.

This approach is expensive as it requires a sophisticated detector to be fabricated. Furthermore, spatial resolution is reduced by a factor of two since the two potential wells take up roughly the same space as two pixels in a conventional detector.

Spatially Varying Pixel Exposures: In [15], the notion of spatially varying pixel sensitivities was introduced. The idea is to assign different (fixed) exposures to neighboring pixels on the image detector. When a pixel is saturated in the acquired image it is likely to have a neighbor that is not, and when a pixel produces zero brightness it is likely to have a neighbor that produces non-zero brightness. In short, there are no large “holes” of saturation or darkness in the acquired image and an HDR image can be reconstructed from it. In this case, one is really trading-off spatial resolution for brightness resolution.

Smart Pixels: In [2], a novel solid state image sensor is described where each pixel on the device includes a computational element that measures the time it takes to attain full potential well capacity. Since the full-well capacity is the same for all pixels, the time to achieve it is proportional to scene radiance. This approach is attractive but faces the challenge of scaling to high spatial resolution while keeping fabrication costs under control.

In this category, also lie other custom detectors that are designed to adjust the exposure of each pixel (or a set of pixels) independently (see [6], [20]). As we shall see shortly, the development of such devices is really complementary to the adaptive imaging technique we propose in this paper.

2 Adaptive Dynamic Range Imaging

We now present the basic principles underlying adaptive dynamic range (ADR) imaging and show why it can achieve significant improvements in dynamic range over the existing HDR methods described in section 1. Then, we will describe a few different optical configurations that can be used to implement such an imaging system.

2.1 Pixels with Adaptive Exposures

We use the term adaptive “imaging” rather than “detector” or “sensor” as the adaptation we propose takes place in optical domain, that is, during image formation. The effect of the adaptation is a change in the incident irradiance of each pixel¹, based on the scene

¹Note that this is different from automatic gain control (AGC) or auto-exposure control (AEC) features found in commercial




Figure 1: The exposure of a pixel can adapt to scene radiance by using an optical attenuator whose transmittance is computed based on the brightness measured by the pixel.

radiance the pixel is meant to measure. This is done using a *controllable* optical attenuator, as shown in figure 1. The brightness I measured by a detector pixel is used to adjust the transmittance T of the attenuator cell corresponding to the detector pixel. Here, the subscripts t and $t + 1$ correspond to the points in time an image and the next one are captured.

The above adaptation technique ensures that the measurement “bits” of each pixel are most effectively used. Figure 2 shows two pixels that are subjected to scene radiance values $L^{(1)}$ and $L^{(2)}$. Let us assume that initially the transmittances of the attenuators of the two pixels are equal, i.e. $T_t^{(1)} = T_t^{(2)}$. For this attenuation level, it turns out that the measured brightness values are $I_t^{(1)}$ and $I_t^{(2)}$, where $I_t^{(1)}$ is a low value and $I_t^{(2)}$ is equal to the saturation level I_{sat} of the detector. Based on these measurements, the attenuations of the pixels are adjusted to $T_{t+1}^{(1)}$ and $T_{t+1}^{(2)}$, such that both the radiance values produce large but unsaturated brightness levels. *This ability to vary the exposures of individual pixels over time gives adaptive imaging a significant edge over conventional HDR techniques. In fact, since our approach is an optical one, in principle, it can be used to further enhance the dynamic range of any existing HDR imaging method.*

A couple of observations are worth making. First, it may appear that the initial measurement $I_t^{(1)}$ of the radiance $L^{(1)}$ should be good enough. However, the best brightness resolution for $I_t^{(1)}$ is obtained when the radiance $L^{(1)}$ falls in the most sensitive region of the radiometric response function. Since the response of the detector is known, it is possible to find the transmittance $T_{t+1}^{(1)}$ that would achieve this. When the image detector has a linear response function (as shown in figure 2), the best exposure is clearly the one that brings $L^{(1)}$ just beneath saturation. That is:

$$T_{t+1}^{(1)} = T_t^{(1)} \frac{(I_{sat} - \epsilon)}{I_t^{(1)}}, \quad (1)$$

cameras. These features do not permit the exposures of individual pixels to be controlled.




Figure 2: Adaptive dynamic range imaging is based on the ability to spatially and temporally control the exposure of the image detector. Here exposures (and hence the response functions) of two image pixels are varied based on the brightness values they measure. See text for details.

where, ϵ is a small number.

Our second observation pertains to $L^{(2)}$; the above procedure cannot be applied to $L^{(2)}$ since $I_t^{(2)}$ was saturated. That is, we have no way of finding the best exposure $T_{t+1}^{(2)}$ since we do not have a meaningful estimate of $L^{(2)}$. This is the most fundamental problem faced by the adaptive imaging technique. The solution is to develop a control algorithm that quickly brings $L^{(2)}$ beneath the saturation level I_{sat} . Once this happens, the optimal transmittance can be computed as in the case of $L^{(1)}$. The real-time control algorithm used to achieve this will be described in section 5.

2.2 Dynamic Range of an Adaptive Pixel

The dynamic range of a conventional camera can be defined as

$$DR = 20 \log \frac{I_{max}}{I_{min}}, \quad (2)$$

where I_{min} and I_{max} are brightness values corresponding to the minimum and maximum radiance values that can be measured by the detector [21]. The minimum brightness I_{min} is typically set to 1. Therefore, an 8-bit detector has a dynamic range of $20 \log 255 = 48.13$ decibels.

In the case of an ADR camera, the minimum gray level remains $I_{min} = 1$ but the maximum detectable gray level becomes $I_{max} (T_{max}/T_{min})$, where T_{max} and T_{min} are the maximum and minimum transmittances that can be achieved using the controllable attenuator shown in figure 1. Therefore, the dynamic range of an ADR

camera can be written as:

$$DR_{adr} = 20 \log \frac{I_{max} T_{max}}{I_{min} T_{min}}. \quad (3)$$

In short, the dynamic range of the ADR system is the sum of the dynamic ranges of the detector and the attenuator used. For example, if the transmittance of the attenuator can be controlled with 8 bits of precision and the image is captured using an 8-bit camera, then $DR_{adr} = 20 \log (255 \times 255) = 96.26$ decibels, which is equivalent to the dynamic range of a 16-bit camera.

This does not imply that such an ADR system has the uniform quantization of a 16-bit camera. In the ADR case, the quantization levels of the detector are uniformly spaced along the response function but the slope of the response function depends on the attenuation of the pixel. Therefore, the quantization of the complete system is non-uniform with larger radiances being measured with less precision than lower ones (see figure 2).

3 ADR Imaging with Spatial Modulator

To extend the concept of adaptive imaging from a single pixel to the full image plane, we use a two-dimensional spatial light modulator (SLM). An excellent survey of SLM's can be found in [16]. A large variety of SLM's have been developed that differ from each other in the mechanisms/devices they use to modulate incident light. Some of the popular ones are liquid crystals, electro-optical devices, acousto-optical devices, and digital micromirror devices. For our implementation, we chose a liquid crystal (LC) spatial attenuator [10], [14] because these are widely available and easy to control.

Figure 3 shows different ways of configuring an adaptive dynamic range system using a spatial light attenuator. In figure 3(a), the attenuator is placed adjacent to the detector so that the attenuation takes place just before the image is formed. This configuration is clearly the most compact but it poses serious optical challenges. Generally, detector pixels tend to be smaller than attenuator pixels. When this is the case, the boundaries between adjacent attenuator pixels can cause significant obstruction and diffraction of the imaged light rays. This, in turn, results in image degradation.

These problems can be alleviated by placing the attenuator in front of the imaging lens, as shown in figure 3(b). Here, the diffraction effects are reduced since each pixel on the detector receives light that passes through multiple attenuator pixels. The number of attenuator pixels that affect the brightness of a single pixel is determined by the aperture size of the imaging lens. In this case, during adaptation, detector pixels cannot be controlled independent of each other; couplings between local sets of detector pixels must be taken into account. In our implementation, we chose to use this configuration and

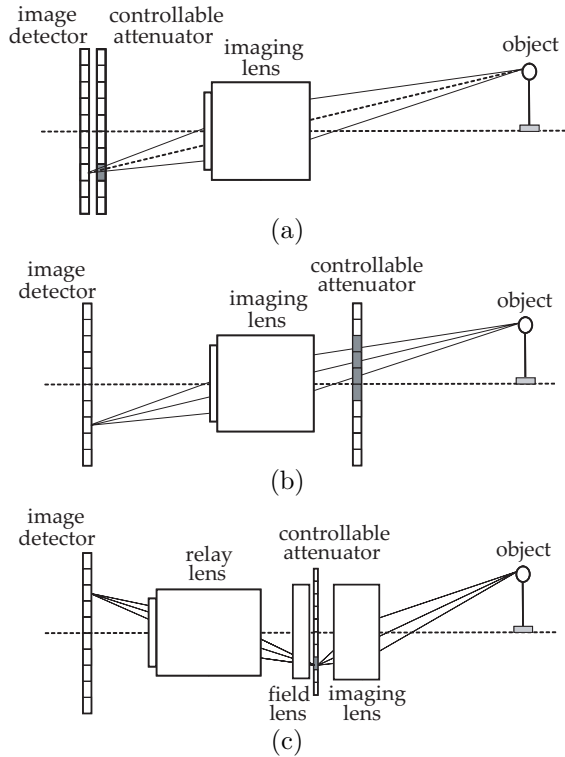


Figure 3: Adaptive dynamic range imaging is implemented using a controllable spatial light attenuator. The attenuator can be placed (a) adjacent to the image detector, (b) in front of the imaging lens, or (c) within the imaging optics such that the image is first formed on the attenuator and then reimaged onto the detector.

developed a control algorithm that takes the above spatial coupling into account.

If one really wants pixel-level control of image irradiance, the configuration shown in figure 3(c) can be used. In this case, an image of the scene is first formed at the plane of the attenuator itself. This intermediate image is then “reimaged” using a field lens and a relay lens onto the image detector.

4 The Imaging Model

As we mentioned in the previous section, for our implementation of the ADR camera we chose the configuration shown in figure 3(b). A simplified optical layout for this case is shown in figure 4. Our goal here is to develop a comprehensive model of this system. In section 5, we will use this imaging model to develop a real-time ADR algorithm.

4.1 Geometry

Since the attenuator is being used to control detector irradiance, we must know, *a priori*, the mapping between points on the attenuator and the detector. Both




Figure 4: Optical layout of an ADR imaging system that is based on the configuration shown in figure 3(b).

these devices are assumed to be planar² but they can be arbitrarily translated and/or rotated with respect to each other. In the case of the simple lens model shown in Figure 4, we see that both devices share the same center of projection (the center of the thin lens). Even in the case of a compound lens, it is easy to see that the mapping from the attenuator plane to the detector plane must be a perspective one. Therefore, the mapping between points on the two planes is given by a projective transformation, or a homography. Let us assume that corresponding points on the attenuator and the detector are denoted by homogeneous 3-vectors \mathbf{u} and \mathbf{x} , respectively. Then:

$$\mathbf{x} = \mathbf{H} \mathbf{u}, \quad (4)$$

where, \mathbf{H} is the projective transformation and is a non-singular 3×3 matrix. If the points are defined in local 2D coordinate systems, located on the planes of two devices, as (x, y) and (u, v) , we have:

$$x = \frac{h_{11}u + h_{12}v + h_{13}}{h_{31}u + h_{32}v + h_{33}}, \quad y = \frac{h_{21}u + h_{22}v + h_{23}}{h_{31}u + h_{32}v + h_{33}}, \quad (5)$$

where, h_{ij} are the elements of \mathbf{H} . Note that the above expressions are linear in h_{ij} and are unaffected by multiplication by a scale factor. Hence, one of the elements of \mathbf{H} can be set to an arbitrary constant and the remaining ones can be found using 4 or more pairs of corresponding points on the attenuator and the detector.

We have assumed here that the lens of the imaging system adheres to perspective projection and does not have severe radial or tangential distortions. If this is not the case, we advocate using either a high-order polynomial or a piece-wise linear function to represent the mapping.

4.2 Radiometry

Consider the image point (x, y) in figure 4. Let the radiance of the corresponding scene point be $L(x, y)$.

²The attenuator can be modeled as a thin, flat plastic sheet of constant thickness (less than 1 mm in the case of an LCD attenuator).

The light rays from the scene point that are captured by the imaging lens first pass through the attenuator. Therefore, image irradiance at (x, y) can be written as:

$$E(x, y) = L(x, y) T(u(x, y), v(x, y)) P(x, y), \quad (6)$$

where, T is the *effective optical transmittance* for the corresponding attenuator point (u, v) and the factor P includes the effects of the lens aperture and other radiometric lens effects such as fall-off and vignetting.

The effective transmittance is of great significance to us. Note that, if the imaging lens were replaced by a pinhole camera, only a single light ray would make it from the scene point to the detector. In this case, the effective transmittance $T(u, v)$ would be exactly equal to the *actual optical transmittance* $\tilde{T}(u, v)$ of the spatial light modulator used. In the case of a lens, however, each image point receives a bundle of light rays from its corresponding scene point (see figure 4). Therefore, the effective transmittance must be computed as an integral of the actual transmittance over the area of intersection between the light bundle and the attenuator.

If the aperture is circular and the attenuator is roughly parallel to the plane of the lens, the intersection of the light bundle with the attenuator is a circular disc of radius (see figure 4):

$$r = \frac{a(z-d)}{z}, \quad (7)$$

where, a is the diameter of the lens aperture, d is the distance of the attenuator from the lens, and z is the depth of the scene point. Since $d \ll z$, it is safe to assume that $r = a$, i.e. the radius r is independent of the location of the point in the scene. Therefore, the effective and actual transmittances can be related as:

$$T(u, v) = b(u, v; r) * \tilde{T}(u, v), \quad (8)$$

where, b is a blurring function which can be approximated with a pillbox function of radius r . In practice, we find that using a pillbox convolved with a Gaussian works better³.

Since we will be capturing digital images and outputting digital control signals, we must take into account the radiometric responses of the image acquisition system and the controllable attenuator. The image irradiance E is related to measured brightness value I as:

$$I(x, y) = f_c(E(x, y)), \quad (9)$$

where f_c is typically a monotonic, non-linear function. The non-linearity arise from the characteristics of the

³Due to vignetting, it is possible that the blurring function varies over the field of view of the imaging system. In practice, we have found these variations to be very small and hence we have assumed the blurring function to be spatially invariant.

image detector, the imaging electronics, and the image acquisition hardware. Similarly, the actual optical transmittance \tilde{T} of the spatial attenuator is related to the control signal C applied to the attenuator as:

$$\tilde{T}(u, v) = f_a(C(u, v)). \quad (10)$$

Here, f_a represents the effects of the attenuator controller and the transfer function of the optical elements of the attenuator. In our implementation of ADR imaging (section 6), we have used simple calibration procedures to determine the geometric (\mathbf{H}) and radiometric (r, b, f_c and f_a) characteristics of the system.

5 Spatio-Temporal Attenuator Control

We now develop a real-time control algorithm for spatially and temporally adapting the dynamic range of the imaging system.

5.1 Control Algorithm

At one level, the control scheme we seek has simple goals; to brighten dark scene points and to darken bright ones. From a technical viewpoint, however, we are faced with a control task that is hard for the following reasons.

- **Scene Dynamics:** Any control mechanism is most effective when it has at its disposal a good dynamical model of the system it seeks to control. In our case, our observations are limited to brightness images of the scene. In the case of any reasonable scene, it is hard to model/predict the dynamics of scene radiance values.
- **Saturation:** The radiance values that are of most interest to us are those that produce saturated brightness measurements. Given a saturated measurement, there is no obvious way to determine the optimal attenuation for the corresponding scene point.

- **Coupling due to Attenuator Blurring:** Finally, the most challenging problem is to develop an algorithm that is stable despite the coupling between adjacent measured image irradiances, which was discussed in section 4.

To ensure that the algorithm can keep up with scene changes, we make our control mechanisms simple and fast. We assume that the system is initialized with the highest exposure level so that points are as bright as can be and the control system's main objective is to remove saturation. Also, to avoid oscillations (instabilities), we introduce a bias in the system by being very aggressive in lowering saturated brightness values and relatively conservative in increasing overly attenuated values.

We will drop the spatial coordinates in our notations and describe the algorithm using a single image point. The same algorithm is applied to all image points. First, the measured image is mapped to a normalized irradiance image using the inverse of the radiometric response function of the imaging system, $E = f_c^{-1}I$,




Figure 5: The image irradiance levels and ranges that are used by the ADR control algorithm. See text for details.

where $0 \leq E \leq 255$. Figure 5 shows the key parameters (thresholds and ranges) used by the algorithm. Any irradiance above E_{sat} is considered to be saturated. Note that E_{sat} is set to be slightly lower than the actual saturation level of the imaging system. The goal of the algorithm is to try to bring all irradiance values within a “desirable range” which is $E_{des} \pm \Delta E$. This range is placed below the saturation level E_{sat} and in the most sensitive part (highest SNR) of the detector’s radiometric response function.

The complete control algorithm can be stated as a set of actions (denoted as \rightarrow) that are applied to the *effective* transmittance of a point based on where its irradiance lies with respect to the levels defined above:

$$\begin{aligned}
 E \geq E_{sat} &\rightarrow T_{t+1} = \alpha \frac{T_t}{2} + (1 - \alpha)T_t, \\
 E_{sat} > E \geq E_{des} + \Delta E &\rightarrow T_{t+1} = \alpha T_t \frac{E_{des}}{E} + (1 - \alpha)T_t, \\
 E_{des} + \Delta E > E \geq E_{des} - \Delta E &\rightarrow T_{t+1} = T_t, \\
 E_{des} - \Delta E > E &\rightarrow T_{t+1} = \beta T_t \frac{E_{des}}{E} + (1 - \beta)T_t. \quad (11)
 \end{aligned}$$

These actions can be summarized as follows. If a pixel is saturated, its transmittance is reduced by a large fraction. This is a very aggressive way of bringing the corresponding scene radiance within the dynamic range of the detector. If the irradiance is below saturation but above the desirable range ($E_{des} \pm \Delta E$), then the transmittance is decreased. If it is below the desirable range, the transmittance is increased. Note that this applies not only to dark scene points but also points that are initially saturated but fall below the desirable range due to repeated reduction. Finally, if a point lies within the desirable range, its transmittance is left unchanged. The parameters (gains) α and β lie between 0 and 1 and determine speeds of the transmittance control in the downward and upward directions, respectively. In our simulations [14] as well as our prototype implementation, good results were obtained by choosing $\alpha = 1$ and $\beta = 0.5$.

Note that the above algorithm only computes the desired effective transmittance. As we know from section 4.2, to obtain this transmittance, we need to apply to

the spatial modulator the actual transmittance function $\hat{T}(u, v)$. From equation 8, we see that, in principle, $\tilde{T}(u, v)$ can be obtained by deconvolving the effective transmittance with the blurring function $b(u, v; r)$. This operation is equivalent to high-pass filtering the effective transmittance function, which often results in actual transmittance values that lie well outside the dynamic range of the spatial modulator. Hence, we are forced to compromise. We settle for smoothly varying dynamic range adaptation rather than one that has pixel-level resolution. Instead of high-pass filtering, we apply a morphological type of operation to the effective transmittance function that “grows” the lower transmittance values:

$$\hat{T}_{t+1}(u, v) = \min_{-N \leq m, n \leq N} T_{t+1}(u + m, v + n), \quad (12)$$

where, N is the size of the growing filter. Note that this operation is consistent with our goal of being aggressive in darkening and conservative in brightening image irradiances. Lastly, this transmittance function is mapped to an attenuator control signal by using the inverse of the radiometric response of the attenuator:

$$\hat{C}_{t+1}(u, v) = f_a^{-1} \hat{T}_{t+1}(u, v). \quad (13)$$

To understand the behavior of our control scheme, we have performed a large number of simulations which we have detailed in [14]. These simulations have verified that our algorithm is able to converge to the desired image irradiance values very quickly and hence it is able to keep up with rapid changes in scene radiances due to the motion of objects and/or changes in illumination.

5.2 From ADR Image to Scene Radiance

The control algorithm brings all scene radiance values within the measurable range of the detector. To achieve this different pixels are assigned different effective exposures to radiance. In short, adaptive dynamic range performs an *optical compression*⁴ of the dynamic range of the scene. Such a compression is generally non-monotonic; image brightnesses are expected to have a different relative ordering from scene radiances.

This dynamic range compression can be very effective when the output of the ADR camera is used for monitoring by a human, for instance in video surveillance. In applications that are more quantitative in nature, precise estimates of scene radiance are of interest. In such cases, the measured image irradiance $E_t(x, y)$ and the applied transmittance $\hat{T}_t(u, v)$ can be used to compute an HDR image of the scene using equations 6 and 8:

$$L_t(x, y) = \frac{E_t(x, y)}{(\hat{T}_t(u, v) * b(u, v; r)) P(x, y)}. \quad (14)$$

⁴This compression is far more effective than any type of digital dynamic range compression applied to images after they are captured.

6 Implementation

We now describe the ADR camera system we have developed and present experimental results.

6.1 ADR Camera

The imaging system we have developed is shown in figure 6. A Sony DXC-107A color video camera is used with a 12.5 mm focal length Fujinon imaging lens. In front of the lens is attached a VLCD25-P LCD panel from Marshall Electronics. The accompanying LCD electronics allows us to input our transmittance function as an NTSC video signal. The output of the video camera is connected to a Matrox Meteor II digitizer board which captures images (640×480) at 30 Hz. The control algorithm is executed on a Dell Workstation that has a Pentium 4 processor and 1 Gb of RDRAM.




Figure 6: The adaptive dynamic range camera consisting of a color video camera, an imaging lens, an LCD spatial attenuator, and electronics to control the attenuator.

The imaging parameters of the ADR camera were obtained via a comprehensive calibration procedure. The details of the calibration are given in [14]. The calibration results in the following terms, each of which is defined in section 4: (a) the geometric mapping \mathbf{H} between the attenuator and detector; (b) the radiometric response f_c of the image acquisition system, for each of the three color channels; (c) the response function f_a of the attenuator; (d) the combined spatial fall-off function $P(x, y)$ for the lens and the attenuator; and (e) the blurring $b(u, v; r)$ of the attenuator transmittance due to the finite aperture of the imaging system.

6.2 Experimental Results

The ADR camera was tested using a variety of indoor and outdoor scenes. Two of these results are shown in figures 7 and 8. In each case, we recorded the video output of the ADR camera and the video input to the LCD attenuator. For comparison purposes, the same scene was also recorded with the ADR feature turned off (i.e. a conventional 8-bit camera). In all case, the automatic gain control (AGC) of the camera was left on, as this is



Figure 7: A scene with a bright sheet of paper. While the paper is saturated in the case of a conventional camera (despite the use of AGC), the crumpling of the paper and the text on it are clearly visible in the case of the ADR camera. Note the spatio-temporal adaptation of the attenuation function shown in the right column.

typically the case with surveillance cameras. We have shown the raw output of the ADR camera rather than the 16-bit HDR output as the HDR output would have to be dynamic range compressed to be displayed on an 8-bit monitor.

In Figure 7, we see that even with AGC turned on, a conventional video camera saturates the sheet of paper. In the case of the ADR camera, all radiances are measured without saturation. As a result, the text on the sheet of paper and the crumpling of the paper are clearly visible. The transmittance of the attenuator is shown alongside the ADR camera output to demonstrate how the ADR algorithm quickly adjusts transmittances over space and time to keep up with changing scene radiances.

Figure 8 shows a typical video surveillance setting where a person is seen emerging from a dark room into the brightly lit outdoors. In the case of the conventional camera (with AGC), the person is visible when inside the dark room but completely saturated when he enters the sunlight. In contrast, the ADR camera produces a very detailed image of the scene; the person, his activ-




Figure 8: A typical video surveillance scenario where people emerge from dark areas into brightly lit ones. In the case of a normal camera (with AGC), the person is virtually invisible (saturated) when he is outdoors. The ADR camera maintains the contrast of the person as well as objects and textures in the scene, irrespective of their lighting.

ities, and textures in the scene are easily recognizable. This example shows the value of ADR imaging in the context of video surveillance and monitoring.

Acknowledgments

This work was conducted at the Computer Vision Laboratory (CAVE) at Columbia University. It was supported in parts by an ONR Grant (No. N00014-03-1-0023) and a DARPA HID Contract (No. N00014-00-1-0916). The authors are grateful to Richard Longman at the Mechanical Engineering Department at Columbia University for the valuable discussions on feedback control systems.

References

- [1] M. Aggarwal and N. Ahuja. High Dynamic Range Panoramic Imaging. *Proc. of International Conference on Computer Vision (ICCV)*, 1:2–9, 2001.
- [2] V. Brajovic and T. Kanade. A Sorting Image Sensor: An Example of Massively Parallel Intensity-to-Time Processing for Low-Latency Computational Sensors. *Proc. of IEEE Conference on Robotics and Automation*, pages 1638–1643, April 1996.

- [3] P. Burt and R. J. Kolczynski. Enhanced Image Capture Through Fusion. *Proc. of International Conference on Computer Vision (ICCV)*, pages 173–182, 1993.
- [4] P. Debevec and J. Malik. Recovering High Dynamic Range Radiance Maps from Photographs. *Proc. of ACM SIGGRAPH 1997*, pages 369–378, 1997.
- [5] H. Doi, Y. Hara, Y. Kenbo, and M. Shiba. Image sensor. Japanese Patent 08-223491, August 1986.
- [6] R. Ginosar and A. Gnusin. A Wide Dynamic Range CMOS Image Sensor. *IEEE Workshop on CCD and Advanced Image Sensors*, June 1997.
- [7] R. Ginosar, O. Hilsenrath, and Y. Zeevi. Wide dynamic range camera. U.S. Patent 5,144,442, September 1992.
- [8] R. J. Handy. High dynamic range CCD detector/imager. U.S. Patent 4623928, November 1986.
- [9] M. Konishi, M. Tsugita, M. Inuiya, and K. Masukane. Video camera, imaging method using video camera, method of operating video camera, image processing apparatus and method, and solid-state electronic imaging device. U.S. Patent 5420635, May 1995.
- [10] E. Lueder. *Liquid Crystal Displays*. John Wiley and Sons, Chichester, U.K., 2001.
- [11] B. Madden. Extended Intensity Range Imaging. Technical Report MS-CIS-93-96, Grasp Laboratory, University of Pennsylvania, 1993.
- [12] S. Mann and R. Picard. Being ‘Undigital’ with Digital Cameras: Extending Dynamic Range by Combining Differently Exposed Pictures. *Proc. of IST’s 48th Annual Conference*, pages 442–448, May 1995.
- [13] T. Mitsunaga and S. K. Nayar. Radiometric Self Calibration. *Proc. of Computer Vision and Pattern Recognition ’99*, 1:374–380, June 1999.
- [14] S. K. Nayar and V. Branzoi. Adaptive dynamic range imaging using a spatial light modulator. *Technical Report, Computer Science, Columbia University*, 2002.
- [15] 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.
- [16] I. A. Neff, R. A. Athale, and S. H. Lee. Two Dimensional Spatial Light Modulators: A Survey. *Proceedings of the IEEE*, 78(5):826–855, May 1990.
- [17] M.A. Robertson, S. Borman, and R.L. Stevenson. Estimation theoretic approach to dynamic range improvement through multiple exposures. *Proc. of IEEE International Conference on Image Processing (ICIP)*, pages 159–163, 1999.
- [18] K. Saito. Electronic image pickup device. Japanese Patent 07-254965, February 1995.
- [19] Y. Schechner and S. K. Nayar. Generalized Mosaicing. *Proc. of International Conference on Computer Vision (ICCV)*, 1:17–24, 2001.
- [20] P. M. A. Serafini and C. G. Sodini. A Programmable Wide Dynamic Range CMOS Imager with On-Chip Automatic Exposure Control. <http://www-mtl.mit.edu/research/sodini/sodiniresearchreports.html>, 2000.
- [21] A. J. P. Theuwissen. *Solid State Imaging with Charge-Coupled Devices*. Kluwer Academic Press, Boston, 1995.
- [22] Y. T. Tsai. Method and apparatus for extending the dynamic range of an electronic imaging system. U.S. Patent 5309243, May 1994.