

Defocus techniques for camera dynamic range expansion

Matthew Trentacoste^a and Cheryl Lau^a and Mushfiqur Rouf^a and Rafał Mantiuk^a and Wolfgang Heidrich^a

^aUniversity of British Columbia, 2366 Main Mall, Vancouver, Canada

ABSTRACT

Defocus imaging techniques, involving the capture and reconstruction of purposely out-of-focus images, have recently become feasible due to advances in deconvolution methods. This paper evaluates the feasibility of defocus imaging as a means of increasing the effective dynamic range of conventional image sensors. Blurring operations spread the energy of each pixel over the surrounding neighborhood; bright regions transfer energy to nearby dark regions, reducing dynamic range. However, there is a trade-off between image quality and dynamic range inherent in all conventional sensors.

The approach involves optically blurring the captured image by turning the lens out of focus, modifying that blurred image with a filter inserted into the optical path, then recovering the desired image by deconvolution. We analyze the properties of the setup to determine whether any combination can produce a dynamic range reduction with acceptable image quality. Our analysis considers both properties of the filter to measure local contrast reduction, as well as the distribution of image intensity at different scales as a measure of global contrast reduction. Our results show that while combining state-of-the-art aperture filters and deconvolution methods can reduce the dynamic range of the defocused image, providing higher image quality than previous methods, rarely does the loss in image fidelity justify the improvements in dynamic range.

Keywords: high dynamic range, HDR, camera, deconvolution, aperture filter, defocus imaging

1. INTRODUCTION

The fact that the range of luminances found in the real world greatly exceeds the capabilities of imaging sensors is a fundamental problem encountered during the acquisition of digital images. Real world scenes contain values brighter and darker than the range that can be captured at any one time by conventional image sensors, and as a result, over-exposed and under-exposed pixels commonly occur in photographs. Conventional image sensors cannot match the dynamic range of the scene, and can only capture a subset of the luminances present. Although specialized high dynamic range image sensors can capture the range of luminances found in most real world scenes, they suffer from lower signal-to-noise ratios (SNR) or slower read-out speeds.¹

Photographers have contended with this problem since the advent of photography, and the most common solution is the concept of exposure to control the amount of light that falls on the sensor. While controlling the amount of light reaching the sensor by adjusting the aperture and the exposure time, photographers are able to select which subset of the scene luminances they wish to capture without undesirably over- or under-exposing the image. However, adjusting the exposure does not improve the limited dynamic range that can be acquired. The subset of luminances that can be accurately captured can be thought of as a slice through the entire range of luminances found in the scene. Adjusting the exposure can move the slice up and down the range of scene luminances, and pixels with luminances above the top of the slice are recorded as white, while pixels with luminances below the bottom of the slice are recorded as black. A correctly chosen exposure can minimize the number of over- and under-exposed pixels, resulting in a properly exposed image, but it is the dynamic range

Further author information: (Send correspondence to M.T.)

M.T.: E-mail: mmt@cs.ubc.ca, Telephone: 1-604-827-3979

C.S.: E-mail: cheryls@cs.ubc.ca

M.R.: E-mail: nasarouf@cs.ubc.ca

R.M.: E-mail: mantiuk@cs.ubc.ca

W.H.: E-mail: heidrich@cs.ubc.ca

of the sensor that controls the width of the slice. The fact remains that if the dynamic range of the scene exceeds that of the sensor, some pixels will not be recorded accurately.

Most existing techniques capture multiple slices of the luminance range and combine them into a single image representing a wider slice, and all of these methods require some tradeoff to extend the width of the slice. Multi-exposure high dynamic range reconstruction^{2,3} takes a sequence of slices distributed in time, trading off temporal resolution for a larger dynamic range. Similarly, placing an array of different neutral density filters onto the sensor⁴ can trade spatial resolution for a wider slice of the dynamic range. The best option is to develop new sensor technology^{1,5,6} that is directly capable of capturing a wider slice of the dynamic range, but these sensors are still a ways off from commercial availability and currently suffer from resolution and quality issues.

2. RELATED WORK

We now detail the research related to increasing the effective contrast an imaging sensor can capture by blurring the image it records. This method borrows from two existing fields of work: coded aperture imaging and deconvolution.

2.1 Coded aperture

Coded aperture imaging first appeared in the context of astronomical x-ray and gamma ray imaging as a solution to constraints in feasible optical systems for those telescopes. Compared to visible light, the high energy photons simply pass through media without refracting, making conventional lenses useless. At first, the standard practice was to use a pinhole aperture to produce a sharp resulting image. However, this approach blocks the majority of the energy from the source and has a very poor signal-to-noise ratio (SNR).

The first improvement over the pinhole aperture was the random aperture arrays proposed by Dicke.⁷ Instead of a single pinhole, they used a two-dimensional array of randomly positioned pinholes, resulting in numerous shifted copies of the image overlaid on top of each other, which they attempted to correct with a decoding step. This process significantly improved the SNR of the imaging system, increasing the amount of light captured by the number of pinholes in the array, but did so at the cost of the resolving power. It is impossible to completely undo the cumulative effects of the random array, and there was always some residual blur. This method was improved upon by structured patterns such as URA⁸ and MURA.⁹ These patterns retain the multiple aperture holes of the random array but are constructed in such a way that the placement of the holes has a unique and complete means of decoding the signal. This design permits the complete restoration of the original image, and the patterned array effectively acts as a single pinhole with the SNR of multiple apertures.

More recently, the principles of coded aperture imaging have been employed in visible light photography. Levin et al.¹⁰ and Veeraraghavan et al.¹¹ both make use aperture filters to accurately determine the amount of blur in order to refocus an image. Similarly, Raskar et al.¹² apply the same principle to the temporal domain, using a patterned opening and closing of the shutter across a single exposure in order to decode motion blur.

2.2 Deconvolution

The coded aperture design in visible light optical systems differs from that of systems for x-ray astronomy. Several of the underlying assumptions the optical setup no longer hold, and simple methods of restoring the original image cannot be used. Instead, deconvolution is used to reverse the effects of the blur introduced by the camera optics. In deconvolution, it is assumed that some desired image has been distorted (blurred) by some known function, and the goal is to recover that original image. Mathematically this is represented as:

$$f = f_0 \otimes k + \eta, \tag{1}$$

where the recorded image f is the result of the original image f_0 convolved by some pointspread function (PSF) k with additive noise η . Numerous solutions have been proposed over the years, from frequency-space methods such as Wiener filtering¹³ to iterative methods such as Richardson-Lucy deconvolution^{14,15} and expectation-maximization.

It is exceedingly difficult to recover the original image at good quality. The system of equations does not have an exact solution due to the presence of noise and image information is suppressed by the frequency response of the filter. Additionally, the system is ill-posed, resulting in an infinite number of possible solutions. Results often include ringing around edges, amplified noise and other artifacts. In the case of coded aperture imaging, the PSF k is known and specifically chosen to aid in conditioning the deconvolution. While this design reduces the size of the solution space, it is still no easy task, and additional measures are required to guide the process towards the correct solution.

New methods have incorporated *natural image statistics* into deconvolution algorithms to better recover the original image. The power spectra of images of real world scenes all roughly follow the same distribution. Individual images exhibit variation, but the overall trend carries strongly across all natural images. A number of recent papers have utilized deconvolution algorithms incorporating these assumptions to better recover the original image. Fergus et al.¹⁶ used deconvolution to remove motion blur from single photographs, while Levin et al.¹⁰ used a combination of coded aperture and enhanced deconvolution to recover depth and refocus images. More recently, Bando et al.,¹⁷ using work based on Figueiredo and Nowak¹⁸ and Bioucas-Dias,¹⁹ presented a means of recovering the original image from defocus blur without a coded aperture.

Recent work by Zhou and Nayar²⁰ analyzed which coded aperture patterns are best suited for deconvolution of defocus blur. This body of work provides a basis for how we can take an optically blurred image of reduced dynamic range, condition its frequency distribution with an apertured, then use deconvolution theory paired with natural image statistics to try and recover the original image.

3. OVERVIEW

The majority of the existing methods attempt to expand the dynamic range of the sensor to match a fixed range of real scene luminances incident upon the sensor. We investigate the opposite, reducing the dynamic range of the scene to better fit in the limited range of the sensor.

The method we investigate attempts this reduction in a two-part, combined optical and software approach. First, we optically blur the image to reduce the dynamic range of the scene incident on sensor. Then we restore the original image detail and dynamic range in software. Blurring is a convolution operation, where the energy that would fall on a single photosite of the sensor is spread over a local neighborhood, and reciprocally, that photosite receives some amount of energy from its neighborhood. This exchange reduces the difference between the pixel and its neighbors, thus producing an image with less local contrast.

The underlying assumption of this approach is that the most extreme luminance values are sufficiently spatially distributed such that the local contrast reduction from convolution will reduce the number of pixels outside the sensor's dynamic range. A good candidate for this method would be small point sources, such as streetlights at night, while a bad candidate would be large bright areas, such as daytime sky viewed through a windows.

We analyze, both in terms of possible dynamic range reduction and resulting image quality, the properties of the optical-digital system composed of a blurred image obtained by an aperture filter inserted into the optical path then restored by a deconvolution method. The results are quantified to determine whether any combination can produce a meaningful reduction of dynamic range required to capture by a sensor while maintaining acceptable image quality.

4. PHYSICAL SETUP

While the optical system in a camera lens is complex, the entire collection of lenses can be approximated as a pair of thick lenses with the aperture in between them. This pair of lenses focuses a bundle of rays coming from points in the scene to points on the sensor, while the size and shape of the aperture controls which rays in the bundle reach the sensor.

The optical system focuses the image by directing all the rays in a bundle that originate from a point on the focal plane in the scene to converge to the same point on the imaging sensor. If a point lies in front of or behind the focal plane, the bundle of rays do not converge to a point on the sensor. Instead, the sensor plane

will intersect the cone of light exiting the rear lens element, resulting in a circular pattern projected on the sensor. The amount of defocus determines the radius of the blur circle, and points further from the focal plane are proportionally further from focusing on the sensor and are more blurred since they intersect a bigger slice of the cone.




Figure 1. Placement of aperture filter in optical setup.

As shown in Figure 1, the aperture sits in the middle of the imaginary pair of optical elements we treat as the lens. Light rays pass directly through it. The circular blur pattern normally observed in out-of-focus images results from the circle shape of the aperture in a normal lens. With a different aperture shape, if a point is out of focus, the sensor still intersects with the cone of out of focus rays, but the aperture shape has blocked some rays traveling through it and the shape of the blur matches that pattern of the aperture.

5. CODED APERTURE

As mentioned in the previous section, defocus blur convolves the image of the scene by aperture shape. Thus, the Fourier transform of the captured image has the frequency characteristics correlated with the Fourier transform of the aperture shape. The goal of inserting an aperture filter into the optical system is to produce a blur pattern that has good frequency preservation properties.

Convolution is equivalent to multiplying the Fourier transform of the scene image by the aperture optical transfer function (OTF), the frequency space representation of the aperture shape. Therefore, in a noise-free case, deconvolution can be viewed as dividing the Fourier transform of the scene image by the aperture OTF. Zeros and very small values in Fourier transform of the filter result in division by zero and excessive error in the deconvolved image; these values are responsible for many of the artifacts observed.

The question of “what constitutes a good coded aperture pattern?” involves a number of criteria. We take this inquiry to specifically consider the aperture pattern that yields the highest quality reconstruction of the original image with the least number of artifacts. As mentioned above, the zeros in the frequency response of the filter prevent information from the scene being captured by the camera.

The ideal coded aperture pattern has a frequency spectrum with as few near-zero values as possible, with the locations of the existing near-zero values located as far away from the low frequency bins and each other as possible. Additionally, that property must hold for a number of different scales of the filter, as points at different depths in the scene will be blurred by different amounts.

There are several practical issues beyond these theoretical considerations. Sufficiently small holes in the aperture pattern will cause the light passing through them to diffract, causing ringing artifacts in the captured image. Another desired property is that the aperture pattern transmits as much light as possible. Given these considerations, we will now analyze several coded aperture patterns (shown in Figure 2) and discuss their relative merits:




Figure 2. Aperture filters evaluated.

Standard (circular) aperture: The circular aperture found in conventional lenses corresponds to a first-order Bessel function in frequency space. The Bessel function resembles a dampened `cos` function, with near-zero values around the numerous zero-crossings of the function. This function causes very poor results with conventional deconvolution algorithms but can be made to yield acceptable results with a better deconvolution algorithm utilizing natural image statistics. However, as far as coded aperture patterns go, the standard aperture is very poor, and the only advantage is that it has the highest light transmission of any pattern.

Circular Gaussian: The Gaussian function has two compelling reasons for its use: value of the function never reaches zero and the Fourier transform of a Gaussian function is another Gaussian. However, there are two equally significant caveats. First, while the Gaussian doesn't ever reach zero, it quickly reaches the noise floor of the camera. For the Fourier transform of a Gaussian function to be sufficiently broadband, the aperture pattern is nearly a pinhole. Second, while the Gaussian function has infinite extent, the actual filter shape would be a truncated Gaussian, the Gaussian multiplied by a box filter. The result would have the same zero-crossing issues as the standard aperture.

Veeraraghavan et al.¹¹ started with a MURA code, then improved it with an optimization function. They employed gradient decent to select the binary pattern that best fits the difference in configuration between conventional cameras and x-ray telescopes. Modelling such differences as the linear convolution that occurs in the optical systems of conventional cameras, they iteratively searched for the binary pattern that maximized the minimum value of the filter's frequency response.

Zhou and Nayar²⁰ improve upon the work of Veeraraghavan et al.¹¹ and formulate a definition of the quality of the aperture filter based directly on the quality of the deconvolved image as opposed to the properties of the filter power spectrum alone. They assume that the camera will be used to capture natural images, and they make use of the $1/f$ relation of frequencies in natural images to further enhance their filter choice. Additionally, they note that the amount of noise present in the captured image significantly affects the quality of the deconvolution, and they use a genetic algorithm to search for the optimal coded aperture for a series of different noise levels. Both this paper and Veeraraghavan et al.¹¹ assume that the radius of the blur is determined by another method.

Levin et al.¹⁰ conducted similar work but optimized a filter for a different set of criteria. In their work, they wanted to accurately determine the amount of defocus present at every pixel and use that to recover the depth. Opposed to the work of Veeraraghavan et al.¹¹ and Zhou and Nayar,²⁰ who either directly or indirectly optimized filter patterns with a minimal number of near-zero values, Levin et al.¹⁰ constructed a filter such that there was the maximum possible difference in the position of the zero-crossing in the frequency spectrum for different filter radii. This allowed them accurately recover the amount of blur at every pixel by looking at the missing information, which they were able to fill in with their deconvolution algorithm.

6. DECONVOLUTION

Deconvolution is used to restore the blurred image recorded by the sensor to its original sharpness and dynamic range. As discussed in Section 2.2, more recent deconvolution routines based on natural image statistics have made significant progress towards restoring the sharpness of blurred images.

While the use of a coded aperture filter improves the spatial frequency properties of the captured image, the fact remains that the deconvolution problem is extremely ill-posed. The blurred image could be the result of any one of an infinite number of images that produce the same result when convolved by the chosen filter. However, none of those images will be exactly the same as the perfect (not blurred) photograph of a real world scene.

Real world scenes all share some very specific properties that can be used to guide the result of the deconvolution towards a physically-plausible result. Specifically, natural images tend to consist of large areas of nearly-constant values with sharp divisions between them. Described in a more formal way, the derivatives of natural images have a heavy-tailed distribution with a narrower peak and longer tail than a Gaussian function. This distribution implies that most pixels have values close to zero, but some small number of pixels, especially those lying next to edges, have significantly larger values.

Both Bando and Nishita¹⁷ and Levin et al.¹⁰ have presented deconvolution algorithms that make use of these natural image statistics. They solve a system of equations that includes a weighting term that corresponds to the prior assumption of a natural-image distribution. Bando and Nishita¹⁷ modify the WaveGSM¹⁹ algorithm to operate in gradient space and perform expectation maximization on the resulting non-linear system of equations using the second-order stationary iterative method. Levin et al.¹⁰ approach the problem as finding the maximum likelihood explanation, employing iteratively reweighted least-squares to solve for the non-linear sparse prior term.

7. EVALUATION

Our goal is to determine whether any combination of aperture filter and deconvolution algorithm can produce a meaningful increase in effective dynamic range while maintaining acceptable image quality. Specifically, we compare filters by Veeraraghavan et al.,¹¹ Zhou and Nayer,²⁰ and Levin et al.¹⁰ in addition to the Gaussian and standard aperture. We evaluate deconvolution methods by Bando and Nishita¹⁷ and Levin et al.¹⁰ in addition to classical Wiener filtering¹³ and Richardson-Lucy^{14,15} methods.


In our evaluation, we measure the *reduction* in contrast between the blurred image and the original image as opposed to the expansion in contrast between the result of the deconvolution and the blurred image. Errors in images produced by deconvolution often appear as ringing artifacts around contrast edges, artificially increasing the contrast in those regions. This approach avoids mistaking those artifacts with meaningful increases in dynamic range. The quality of the deconvolution algorithm is still determined as the difference between the original image and the reconstructed image.

The structure of the image has an impact on the effectiveness of the algorithm. If the size of a bright or dark image feature exceeds the given filter diameter, there will be no reduction in dynamic range since energy is only exchanged between pixels within the local neighborhood. Concerning dynamic range reduction, our analysis considers both properties of the filter used as a measure of local contrast reduction as well as the distribution of intensity values in the image at different scales as a measure of global contrast reduction.

Our primary interest was validating the proof of concept, and we conduct our evaluation using synthetic results instead of real optical systems. This approach introduces less complexity in quantifying the performance of the method while still determining if the upper bound of the performance is of sufficient quality.

We present two of the images we used to evaluate the performance in this proposal, *Atrium Morning* and *Atrium Night*, shown in Figure 7. This pair was chosen to demonstrate how the performance of the system depends on the spatial distribution of luminance values in the image. While *Atrium Night* has a larger dynamic range than that of *Atrium Morning*, its extreme values are located in small bright light sources in contrast to the broad skylights in *Atrium Morning*.

Table 1 shows the radii of filters used in the evaluation and the change in dynamic range of the image as a result of being blurred by a standard aperture disk filter of different radii. The night image results in a larger reduction in dynamic range since the point light sources distribute energy over neighborhoods of significantly



Atrium Morning

Atrium Night

Figure 3. Sample images used in evaluation.

dimmer values. The tables and plots included in this section all reference the differences in dynamic range in terms of the photographic concept of exposure value (EV) stops. In photography, a change of 1 stop or 1 EV represents a unit change on the \log_2 scale, where $1 \text{ EV} = \log_2(L_1) - \log_2(L_2)$.

We evaluated the quality of the images reconstructed by the deconvolution algorithms in terms of *peak signal-to-noise ratio* (PSNR).²¹ Typical values for the PSNR in lossy image and video compression are between 30 and 50 dB, where higher is better.^{21,22} However, the images used in this evaluation have a larger dynamic range than that of conventional 8-bit images. From evaluating the quality of the results we have chosen 35 dB to be the lower bound on acceptable image quality for deconvolved images. These images have visible artifacts, but all of the features are still clearly visible. Similarly, we have chosen a dynamic range reduction of 2 stops to be the minimum acceptable reduction in dynamic range.

Our evaluation proceeded as follows. First, each image was convolved by each of the aperture filters at a number of different radii. The minimum and maximum values of the original image and the blurred image were

Radius	Atrium Morning			Atrium Night		
	min	max	reduction	min	max	reduction
Original	0.00	11.0		0.00	12.0	
1	0.00	10.8	0.200	0.452	12.0	0.452
2	0.00	10.6	0.424	0.622	12.0	0.622
3	0.00	10.3	0.716	1.163	11.8	1.34
4	0.02	10.0	1.00	1.436	11.4	1.99
5	0.08	9.94	1.14	1.589	11.4	2.23
6	0.15	9.92	1.24	1.731	11.2	2.51
8	0.31	9.83	1.48	1.890	10.8	3.13
9	0.40	9.79	1.61	1.950	10.5	3.41
11	0.66	9.71	1.94	2.08	10.3	3.74
13	0.86	9.67	2.19	2.18	10.1	4.13
16	1.04	9.59	2.45	2.26	9.61	4.65

Table 1. Amount of reduction in dynamic range as a function of radius of a standard aperture (disk) filter in pixels. The minimum and maximum values are the result of convolving the original image with a disc filter of the specified size, showing how much of the dynamic range reduction was from reducing the intensity of highlight regions versus increasing the intensity of shadow regions. All units are in terms of powers of two, referred to as exposure value (EV) stops.

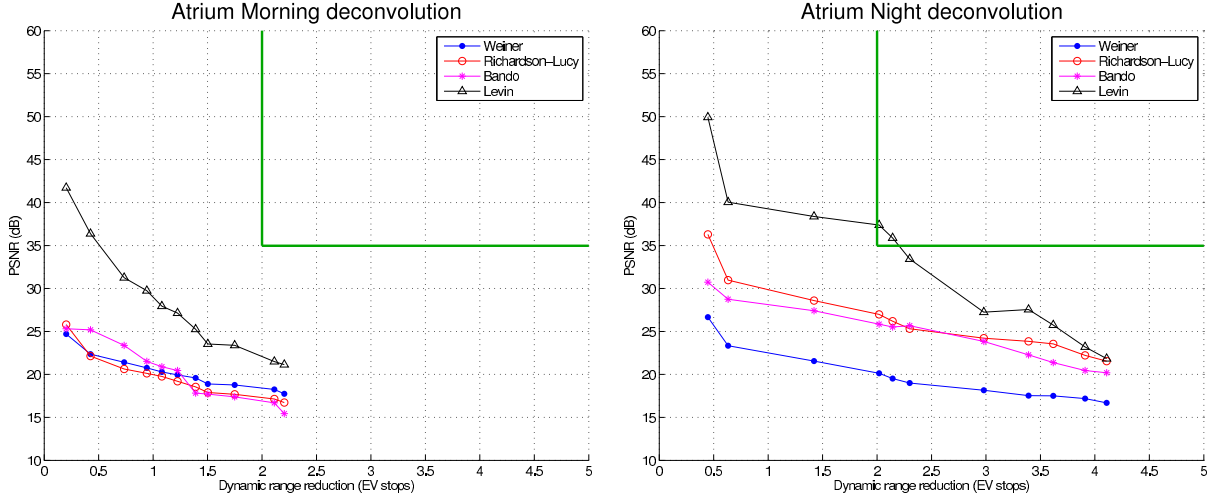


Figure 4. Comparison of all four deconvolution algorithms on the Atrium test scenes without any noise added. All results were computed using the aperture filter proposed by Zhou and Nayar.²⁰ Values above and to the right of the green bars pass our acceptance criteria.

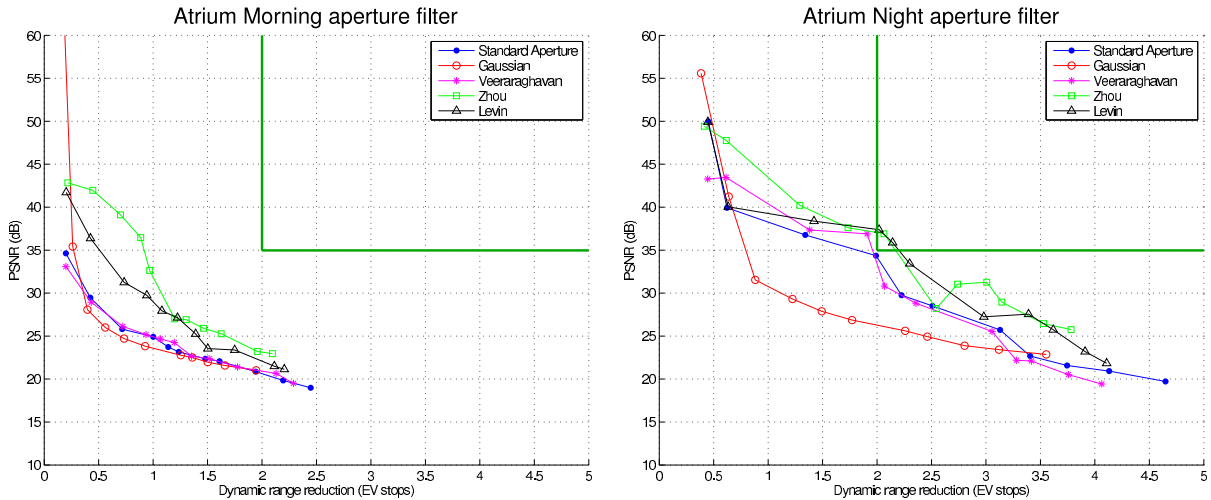


Figure 5. Comparison of all aperture filters on the Atrium test scenes without any noise added. All results were computed using the deconvolution algorithm proposed by Levin et al.¹⁰ Values above and to the right of the green bars pass our acceptance criteria.

compared to compute the amount of dynamic range reduction for that size filter. Next, different amounts of Gaussian noise were added to simulate the random nature of the image acquisition process for different sensor sizes. Then, all of the combinations of filter, radius and noise were deconvolved by each of the deconvolution algorithms. Finally, all results were compared to the original image to compute the PSNR. Figures 4 and 5 summarize our results across the combinations of aperture filter and deconvolution method for images without any noise added, while Figures 6 and 7 summarize our results for images with Gaussian noise $\sigma = 1$ added.

The deconvolution algorithm by Levin et al.¹⁰ was able to produce acceptable results for images with small bright areas, such as *Atrium Night*, when paired with one of the filters by Levin et al.,¹⁰ Veeraraghavan et al.,¹¹ or Zhou and Nayar.²⁰ However, it was only able to do so at noise levels below those found in existing cameras.

Additionally, the Levin et al.¹⁰ method performed worse than Richardson-Lucy for images with a few bright points like *Atrium Night* at realistic noise levels. While it is able to reconstruct very fine details, the method of Levin et al.¹⁰ tends to introduce high frequency ringing. The amount of ringing associated with a given feature is insignificant relative to its overall magnitude, but if that feature is bright enough, the ringing will destroy detail

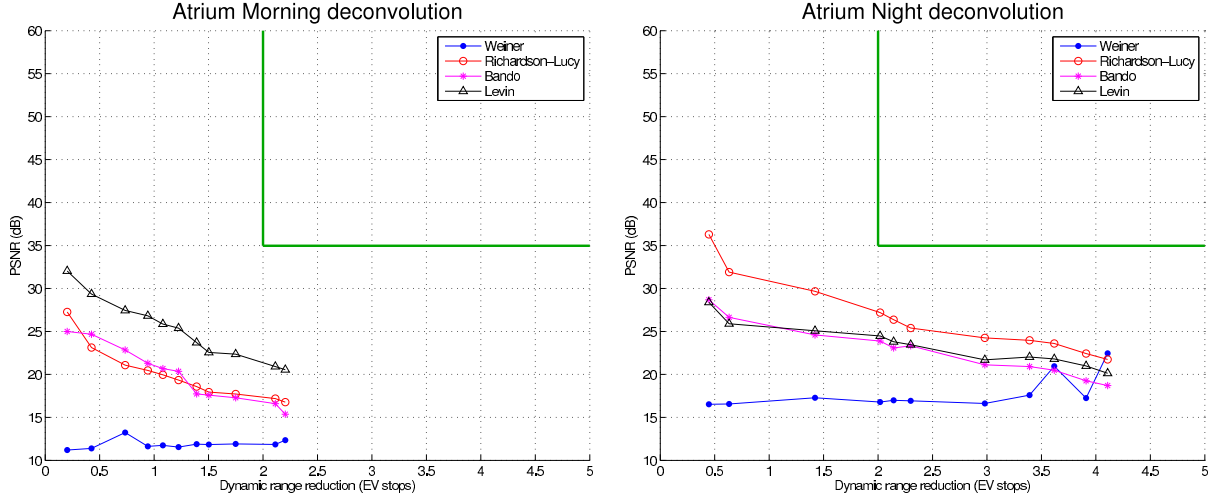


Figure 6. Comparison of all four deconvolution algorithms on the Atrium test scenes with additive Gaussian noise of $\sigma = 1$. All results were computed using the aperture filter proposed by Zhou and Nayar.²⁰ Values above and to the right of the green bars pass our acceptance criteria.

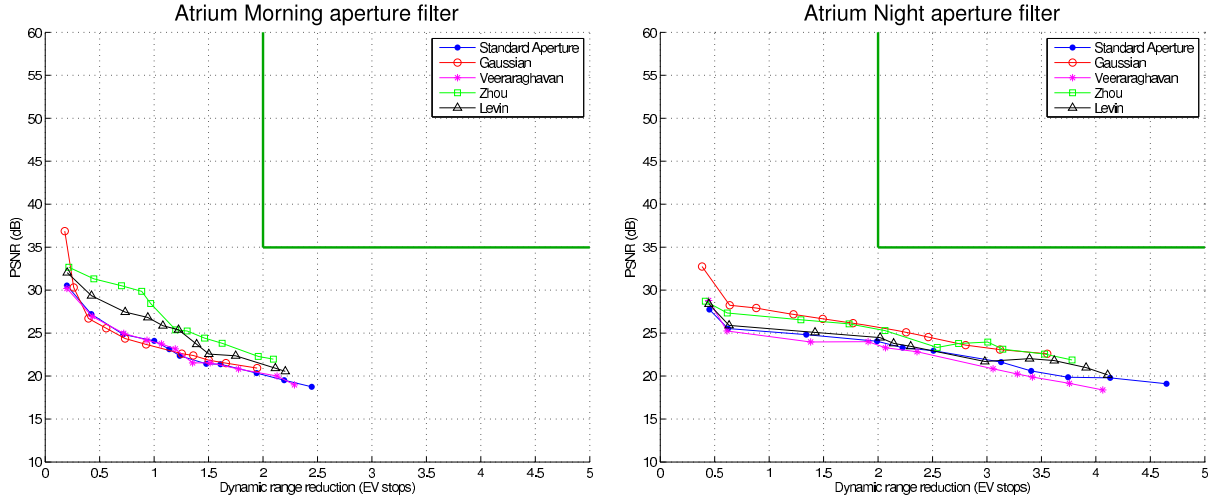


Figure 7. Comparison of all aperture filters on the Atrium test scenes with additive Gaussian noise of $\sigma = 1$. All results were computed using the deconvolution algorithm proposed by Levin et al.¹⁰ Values above and to the right of the green bars pass our acceptance criteria.

in the dark regions of the image. Increasing the smoothing parameter of the algorithm might have produced less noise in the dark regions but would not have recovered any additional detail.

The deconvolution algorithm by Bando and Nishita¹⁷ performed worse than expected, given its success on conventional 8-bit images. This could be caused by the selection of its normalization parameters, which we have not been able to optimize for the given images. However, in all tests on conventional (8-bit) images, it performed at best equal with that of Levin et al.,¹⁰ and we are confident that it could not do significantly better here.

The filters by Levin et al.,¹⁰ Veeraraghavan et al.¹¹ and Zhou and Nayar²⁰ outperformed the other filters for images with low noise levels. As the amount of noise increased, the difference between the capability of all the different filters decreased, and eventually all the filters had the same performance as the conventional aperture. This confirms the observation of Zhou and Nayar²⁰ that the optimized aperture filters they designed became simpler and more like the conventional aperture as the noise level increased.

Overall, the results show that no current combination of aperture filter and deconvolution algorithm can deliver an acceptable combination of dynamic range reduction and image quality for images with large bright

areas like *Atrium Morning* at any noise level, while the desired performance was only possible for images with small bright areas with unrealistically low amounts of noise. Additionally, the quality of the final result depended more on the deconvolution method than on the choice of aperture filter for realistic noise levels.

8. CONCLUSIONS

None of the possible combinations of aperture filter and deconvolution algorithm were able to consistently reduce the dynamic range of the captured image without excessively degrading image quality. Techniques worked for specific scenes under very controlled conditions, but without advances to either aperture filtering or image reconstruction, it is not applicable to general circumstances.

The efficiency of this defocus imaging approach is scene-dependent. The method is good for small over-exposed regions that are just above the maximum photosite capacity of a sensor but performs worse on large overposed areas or in recovering exceedingly bright regions. The more complex deconvolution algorithms performed better than traditional methods but at significant computational cost. These algorithms took on the order of several minutes to produce results for megapixel images. The marginal improvement in dynamic range at acceptable image quality does not justify the amount of computation required by the method.

We could also consider different means of evaluating image quality. High PSNR is rarely the goal of deconvolution. Images may look acceptable, or even better, if they are sharper despite the fact that the PSNR score would be lower than a different version. We could obtain a more accurate estimate of the image quality using different metrics or subjective user preference means of evaluation.

In terms of future work, there are avenues to consider for both improving the aperture filter and the deconvolution algorithm. The aperture filters proposed have not considered maximizing transmission and have not been optimized for this particular application. It would be interesting to see if a different filter could be developed to improve results. Similarly, deconvolution algorithms are not well-tuned for high dynamic range images. The human visual system is sensitive to relative change and can detect small changes in dark regions. However, deconvolution algorithms are linear and minimize absolute error without accounting for how salient that error is to the final observer. An interesting line of research would be attempting to devise a deconvolution algorithm that gives preference to relative error and correctly weights the importance of dark regions.

ACKNOWLEDGMENTS

This work was supported by Dolby Canada under the Dolby Research Chair in Computer Science at the University of British Columbia.

REFERENCES

- [1] Yang, D., Gamal, A., Fowler, B., and Tian, H., "A 640512 CMOS image sensor with ultra wide dynamic range floating-point pixel-level ADC," *IEEE Journal of Solid-State Circuits* **34**, 1821–1834 (1999).
- [2] Debevec, P. and Malik, J., "Recovering high dynamic range radiance maps from photographs," *Proceedings of ACM SIGGRAPH* (Aug 1997).
- [3] Robertson, M., Borman, S., and Stevenson, R., "Dynamic range improvements through multiple exposures," in *Proceedings of International Conference on Image Processing (ICIP) 1999*, 159–163 (1999).
- [4] Nayar, S. and Mitsunaga, T., "High dynamic range imaging: Spatially varying pixel exposures," *International Conference on Computer Vision* **1** (2000).
- [5] IMS Chips, "HDRC VGAX," (2004). <http://www.hdrc.com>.
- [6] Acosta-Serafini, P., Masaki, I., and Sodini, C., "Single-chip imager system with programmable dynamic range." U.S. Patent 6,977,685 (2005).
- [7] Dicke, R., "Scatter-hole cameras for x-rays and gamma rays," *Astrophysical Journal* **153** (1968).
- [8] Fenimore, E. E. and Cannon, T. M., "Coded aperture imaging with uniformly redundant arrays," *Applied Optics* **17**(3), 337–347 (1978).
- [9] Gottesman, S. and Fenimore, E., "New family of binary arrays for coded aperture imaging," *Applied Optics* **28**, 4344–4352 (Oct 1989).

- [10] Levin, A., Fergus, R., Durand, F., and Freeman, W., “Image and depth from a conventional camera with a coded aperture,” *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)* **26** (Aug 2007).
- [11] Veeraraghavan, A., Raskar, R., Agrawal, A., Mohan, A., and Tumblin, J., “Dappled photography: mask enhanced cameras for heterodyned light fields and coded aperture refocusing,” *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)* (Aug 2007).
- [12] Raskar, R., Agrawal, A., and Tumblin, J., “Coded exposure photography: motion deblurring using fluttered shutter,” *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)* **25**, 795–804 (Jul 2006).
- [13] Wiener, N., [*Extrapolation, Interpolation, and Smoothing of Stationary Time Series*], MIT Press, Cambridge, Mass (1964).
- [14] Richardson, W., “Bayesian-based iterative method of image restoration,” *Journal of the Optical Society of America* **62**(1) (1972).
- [15] Lucy, L., “An iterative technique for the rectification of observed distributions,” *The Astronomical Journal* **79**(6), 745–754 (1974).
- [16] Fergus, R., Singh, B., Hertzmann, A., Roweis, S., and Freeman, W., “Removing camera shake from a single photograph,” *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH)* (Jul 2006).
- [17] Bando, Y. and Nishita, T., “Towards digital refocusing from a single photograph,” in [*Proceedings of Pacific Conference on Computer Graphics and Applications*], 363–372 (Sep 2007).
- [18] Figueiredo, M. and Nowak, R., “An EM algorithm for wavelet-based image restoration,” *IEEE Transactions on Image Processing* **12**(8), 906–916 (2003).
- [19] Bioucas-Dias, J., “Bayesian wavelet-based image deconvolution: a GEM algorithm exploiting a class of heavy-tailed priors,” *IEEE Transactions on Image Processing* **15**(4), 937–951 (2006).
- [20] Zhou, C. and Nayar, S., “What are good apertures for defocus deblurring?,” *IEEE Computational Photography* (2009).
- [21] Thomos, N., Boulgouris, N. V., and Strintzis, M. G., “Optimized transmission of JPEG2000 streams over wireless channels,” *IEEE Transactions on image processing* **15** (January 2006).
- [22] Xiangjun, L. and Jianfei, C., “Robust transmission of JPEG2000 encoded images over packet loss channels,” in [*IEEE International Conference on Multimedia*], 947–950, School of Computer Engineering, Nanyang Technological University (2007).