

Image Restoration in Digital Photography

Edmund Y. Lam, *Member, IEEE*

Abstract — *This paper introduces some novel image restoration algorithms for digital photography, which has one of the fastest growing consumer electronics markets in recent years. Many attempts have been made to improve the quality of the digital pictures in comparison with photography taken on films. A lot of these methods have their roots in discrete signal and image processing developed over the last two decades, but the ever-increasing computational power of personal computers has made possible new designs and advanced techniques. The algorithms we are presenting here take advantage of the programmability of the pixels and the availability of a compression codec commonly found inside digital cameras, and work in compliance with either the JPEG or the JPEG-2000 image compression standard¹.*

Index Terms — **Image Restoration, Resolution Enhancement, Digital Photography, JPEG, JPEG-2000.**

I. INTRODUCTION

IN traditional film-based photography, a picture is recorded when incident photons hit the silver halide film to create specks of silver that form a latent image of the scene. This is an irreversible process. Various chemical processes are involved, and it is not easy to modify the picture recorded on the film.

Digital photography, on the other hand, is vastly different. Images are recorded on electronic sensors, either in the form of charge coupled devices (CCD) or complementary metal-oxide semiconductor (CMOS). The latter has shown a lot of promise as a viable, low cost solution because they enable the integration of sensing, processing, and memory. In both cases, the images are recorded directly in a digital medium, without the need for further chemical processes. They can be stored permanently without degradation, while photographic film and prints will inevitably fade over time. The electronic sensors can also be reused once the previous image has been saved in memory, and therefore over time it will be more economical than constantly buying new film. The most important feature of digital photography, however, is the programmability of the pixels. It has already been demonstrated that pixel-level analog-to-digital conversion (ADC) is possible [1]. With further transistor size shrinkage in the future, more logic can be put on the pixel level to enhance the function of the cameras. Furthermore, we already have application-specific integrated circuits (ASIC) on the digital cameras to perform operations such as white balancing, color processing, compression, storage, and transmission. These provide the

bases with which we can manipulate the digital images much more easily than with traditional photography.

In this paper, our goal is to explore methods by which we can restore degraded images by taking advantage of the programmability of digital cameras. We devise novel image restoration algorithms to be implemented in the camera alongside the compression of the images. We aim at creating a more versatile digital camera that would make use of the power of its programmability.

II. IMAGE PROCESSING PIPELINE

A prerequisite to designing effective image restoration algorithms for digital photography is a good understanding of the image acquisition process, especially under incoherent illumination which is the typical case for consumer use of digital cameras. Linear system theory provides us a very powerful tool for the analysis of the optical system inside the cameras. In particular, the frequency domain offers very compact relationships between the object and the image, establishing the preeminent role of Fourier analysis in the theory of incoherent imaging. In this section, we explore a simple model of the imaging system, and describe how restoration can be performed as suggested by linear system theory and Fourier analysis.

A. The Imaging System

Many steps are involved in the imaging process for digital cameras, including white balancing, demosaicing, color matrix correction, nonlinear conversion, compression, and storage. A block diagram of the different steps in the image processing pipeline inside a digital camera is shown in Fig. 1. Here we focus on the image acquisition, which describes how images are formed.

In its simplest form, the image acquisition system can be represented as in Fig. 2. The lens in the diagram represents a system of lenses. The object plane is located at a distance z_o from the entrance pupil, while the image plane is located at a distance z_i from the exit pupil. Our goal is to find the image intensity distribution given a certain object at the object plane. We assume the imaging system has the property that a diverging spherical wave, emanating from a point-source object, is converted by the system into a new spherical wave converging towards a point in the image plane. The locations of the ideal image point and the original object are related by a simple scaling factor, which is constant across all points in the image field of interest. Such systems are called diffraction-limited when the object of interest is confined to the region where the above property holds [2]. We further assume that the camera is used under incoherent light. One important property

¹ Edmund Y. Lam is with the Department of Electrical and Electronic Engineering, University of Hong Kong, Hong Kong. (e-mail: elam@eee.hku.hk).

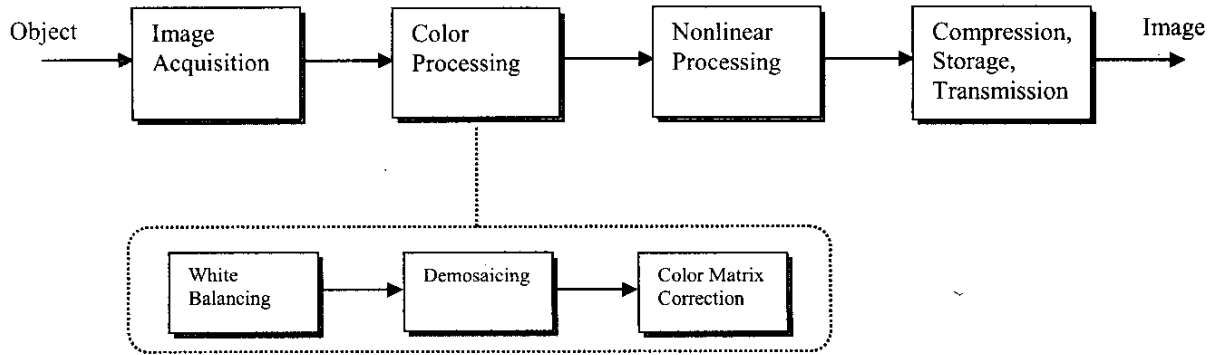


Fig. 1. Simplified block diagram of the image processing pipeline.

of an incoherent imaging system is that it is linear in intensity. Defining $G_g(f_x, f_y)$ as the normalized frequency spectrum of the object and $G_i(f_x, f_y)$ as the normalized frequency spectrum of the image, we have the important relationship [2]

$$G_i(f_x, f_y) = H(f_x, f_y)G_g(f_x, f_y) \quad (1)$$

where $H(f_x, f_y)$ is called the optical transfer function (OTF) of the system. For a diffraction-limited system with incoherent imaging, we can have an analytical formula for the OTF in terms of the pupil function, wavelength, and z_i [2].

In all practical imaging systems, noise is inevitably present. It is common to model it as additive white Gaussian noise at the output, when the dominant source is the random thermal motion of electrons [3]. A uniformly distributed noise assumption is also common for the quantization noise distribution [4], while Poisson noise is typically used for astronomical images, which are taken at low light levels. Now, let us take equation (1) back to the space domain by inverse Fourier transforms, and sample the quantities at regular intervals in both the horizontal and vertical directions. This is useful because computers can only handle discrete images of finite size. We also include the contribution of additive noise to arrive at the following equation:

$$i(x, y) = h(x, y) * g(x, y) + n(x, y), \quad (2)$$

where $g(x, y)$ is the object, $i(x, y)$ is the image, $n(x, y)$ is the noise, and $h(x, y)$ is the point spread function (PSF).

B. Aberrations

In the presence of a point-source object, if the wavefront leaving the exit pupil departs significantly from ideal spherical shape, the imaging system is said to have aberrations [2]. By nature all aberrations are linear. When a particular aberration is space-invariant, its effects can be incorporated by changing the OTF, so that the general imaging equations (1) and (2) are still applicable but with new $H(f_x, f_y)$ and $h(x, y)$.

A simple focusing error is one of the most common space-invariant aberrations. In the literature, the inverse Fourier transform of the OTF is usually modeled as a circular disc with

unit magnitude, the size of the disc being indicative of the amount of focusing error. This is a valid model from a geometric (ray) optics perspective. Its usefulness, however, is limited in cases when diffraction effects govern, for instance, in a lens with small relative aperture. For those situations, Fourier (wave) optics allows us to derive more accurate expressions. Only in the limiting case where the focusing error is severe does the geometric optics expression become a good approximation, because the effects of diffraction become negligible [2]. Readers are referred to [5] for discussions on the analytical expressions of the OTF for various amount of defocus. It suffices to note that as the amount of defocus increases, the magnitude of $H(f_x, f_y)$ for middle to high frequencies is smaller, sometimes even becoming negative. This corresponds to a loss of image details, and visually the effect is that the image is blurred.

The OTF's of many other aberrations can also be derived or measured, such as spherical aberration, coma, astigmatism, curvature of field, and distortion, which are collectively called Seidel aberrations [6]. However, one should be aware that most of these aberrations are space-variant. Therefore, only "local" OTF's can be used. In some cases, if we know the type and amount of the aberration, we can perform a geometric coordinate transformation to turn the space-variant imaging equation into a space-invariant imaging equation similar to equations (1) and (2).

C. Image Restoration

From the perspective of equation (2), the goal of image restoration is to recover the object $g(x, y)$ given the observed image $i(x, y)$. Since the imaging equation represents a convolutional relationship, the restoration problem is often called deconvolution. Generally for image restoration, the PSF $h(x, y)$ is assumed to be known. The restoration problem is termed blind image deconvolution when $h(x, y)$ is not known precisely, and is a subject of active research. See [7] for some recent developments in this area.

When the imaging system has a negligible amount of noise, equation (1) leads to a very direct method for restoration. At

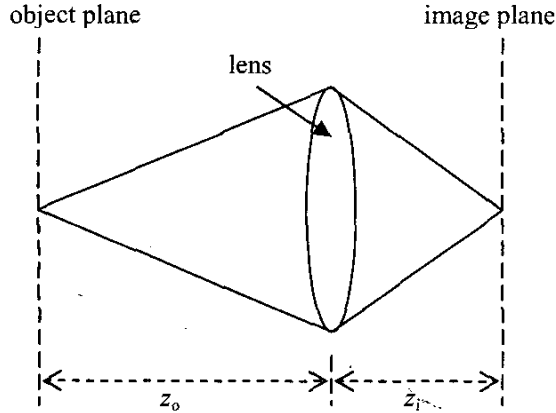


Fig. 2. Generalized model of an imaging system.

each frequency, we define the restored object spectrum $\hat{G}_g(f_x, f_y)$ as

$$\hat{G}_g(f_x, f_y) = \begin{cases} \frac{G_i(f_x, f_y)}{H(f_x, f_y)} & H(f_x, f_y) \neq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

This is called pseudoinverse filtering. Although simple, this method has limited use because of the presence of noise. A better solution is found by using Wiener filtering. This is computed with

$$\hat{G}_g(f_x, f_y) = \frac{H^*(f_x, f_y)}{|H(f_x, f_y)|^2 + \frac{\Phi_n(f_x, f_y)}{\Phi_g(f_x, f_y)}} G_i(f_x, f_y) \quad (4)$$

where $\Phi_g(f_x, f_y)$ and $\Phi_n(f_x, f_y)$ are the power spectra of the object and noise respectively.

III. GENERAL SETTING FOR RESTORATION IN CAMERA

Because of the large quantity of data captured in the camera image, compression is necessary. Each camera is therefore equipped with a compression engine, while the decompression is done after the images are downloaded to a computer host. It is advantageous to see whether it is possible to achieve some restoration akin to equation (4) as part of the compression and decompression process. In this section we will attempt a general formulation, while in sections IV and V we will deal more specifically with JPEG [8] and JPEG 2000 [9] respectively, which are the two most common compression standards.

The method comes in using different encoding and decoding parameters. Consider Fig. 3, which represents the process of blurring, compression, quantization, decompression, and deblurring. Quantization is represented as the addition of quantization noise [4]. Typically, we use a symmetric pair of encoder and decoder, *i.e.*, if $i(x, y)$ is the image before the encoder, then

$$D[E[i(x, y)]] = i(x, y), \quad (5)$$

or to simplify notation,

$$D = E^{-1}. \quad (6)$$

Alternatively, we should rather use another decompressor D' , which incorporates the effect of deblurring as well. This principle has been applied for vector quantization design in the past [10]. Its feasibility can be argued from the following thought experiment: suppose that the quantization effect is small. The blurring causes some loss in high frequency components. If the compression is based on the frequency components, we should design a decompressor that has a boosting effect in the high frequency components with respect to the compressor. The deblurring operation is then "absorbed" into the decompression. In practice, the quantization noise may not be negligible, especially for low bit-rate transmission. Furthermore, the compression is often not in Fourier frequency components; it is not for JPEG and JPEG 2000. The two sections below deal with the specific algorithms possible with the two compression standards. Related works that deal specifically with out-of-focus images have been reported in [5, 11].

IV. IMAGE RESTORATION WITH THE JPEG ALGORITHM

A. Motivation

In JPEG, the first step is to divide an image into 8×8 pixel non-overlapping blocks. Each block is then subjected to a discrete cosine transform (DCT), as follows: let $i(p, q)$ denote a pixel value, $p = 0, \dots, 7$, $q = 0, \dots, 7$, within an 8×8 block. The DCT coefficients are then

$$I(m, n) = \frac{C(m)}{2} \frac{C(n)}{2} \sum_{p=0}^7 \sum_{q=0}^7 i(p, q) \cos\left(\frac{(2p+1)\pi m}{16}\right) \cos\left(\frac{(2q+1)\pi n}{16}\right) \quad (7)$$

with $m = 0, \dots, 7$, $n = 0, \dots, 7$, and

$$C(m), C(n) = \begin{cases} \frac{1}{\sqrt{2}} & \text{for } m, n = 0 \\ 1 & \text{for } m, n > 0. \end{cases} \quad (8)$$

The coefficients are then quantized according to the quantization matrix Q_e , by rounding off the quotients when the DCT coefficients are divided entrywise by the corresponding element in the matrix. They are then entropy-coded, with either modified Huffman coding or arithmetic coding, before transmission. Upon receiving the compressed data, the decoder reverses the process for the entropy coding, dequantizes the coefficients by multiplying entrywise with the matrix Q_d , and performs the inverse DCT [8]. The compression is considered lossy because of the quantization process.

Usually we use the same quantization matrix for both encoding and decoding, *i.e.*, $Q_e = Q_d$. As an example, the JPEG committee uses a matrix that attempts to take into account some human visual system properties. One then needs to transmit the matrix as part of the compressed image




Fig. 3. Image blurring and compression model.

bitstream to the decoder. To adjust the quality of the compression, one usually scales the whole quantization matrix by a constant factor. This is called the quality factor.

However, as argued in the previous section, in the presence of blur or quantization noise using a decompressor that is the inverse of the compressor may not be optimal. In the case for JPEG, that means we should not always set $Q_e = Q_d$. This is explained in [5] and summarized below. Let $I_q(m, n)$ denote the quantized coefficients and $I_n(m, n)$ the quantization noise. To compare the original and the decompressed image, we need to calculate the mean-square error (MSE) in the space domain, which can be done in the DCT domain because of its unitary nature and Parseval's theorem [12]. Therefore,

$$\begin{aligned} \text{MSE} &= \sum_{m=0}^7 \sum_{n=0}^7 (I - \hat{I})^2 \\ &= \sum_{m=0}^7 \sum_{n=0}^7 ((I_q + I_n)Q_e - I_q Q_d)^2 \\ &= \sum_{m=0}^7 \sum_{n=0}^7 ((Q_e - Q_d)I_q + Q_e I_n)^2, \end{aligned} \quad (9)$$

where it is understood that the above quantities all have arguments (m, n) . We can see that when $I_n(m, n)$ is small and $I_q(m, n)$ large, it is reasonable to set $Q_e = Q_d$ to generate a small MSE. However, typically for high frequencies we have larger $I_n(m, n)$ and smaller $I_q(m, n)$, and changing Q_e at those frequencies would provide a tradeoff between the sources of noise and may reduce the overall MSE.

B. Algorithm

We seek to adjust the dequantization matrix by letting it be

$$Q_d(m, n) = a(m, n)Q_e(m, n). \quad (10)$$

In order to calculate $a(m, n)$, we assume that we have at our disposal a collection of images that have both the blurred and restored versions available. Let $I_o(m, n)$ be the vector containing the DCT coefficients of the blurred images, and $I_r(m, n)$ be the corresponding vector for the restored images. We seek to find $a(m, n)$ where $I_o(m, n)a(m, n) = I_r(m, n)$. The best value of $a(m, n)$, in the mean-square sense, can be found by [5]

$$a(m, n) = \frac{\langle I_o(m, n), I_r(m, n) \rangle}{\langle I_o(m, n), I_o(m, n) \rangle}, \quad (11)$$

where $\langle x, y \rangle$ denotes the inner product of x and y .

We can further improve this algorithm if we have some prior estimate of the blur. In this case, we can first process $I_o(m, n)$ with a restoration filter, such as using a Wiener filter with the estimated $h(x, y)$ to produce $\hat{I}_o(m, n)$. Equation (11) is then computed with $\hat{I}_o(m, n)$ replacing $I_o(m, n)$. Because $\hat{I}_o(m, n)$ is now closer to $I_r(m, n)$, but usually not identical due to limitation in the restoration filter and blur estimate, $a(m, n)$ would be more stable and the approximation in $I_o(m, n)a(m, n) = I_r(m, n)$ would be closer to equality. Note however that this would add extra computational burden to the digital camera and may not be feasible for those cameras that need to conserve power as much as possible.

In both cases, since image restoration is an ill-posed deconvolution problem, a minimum MSE solution is known to be highly sensitive to noise, especially at high frequencies. It is very important to incorporate some regularization constraint. A technique to this effect is proposed in [13, 14]. Let $L(m, n)$ be a highpass filter in the DCT domain. We measure the amount of desired high frequencies in all the blocks by

$$D = \sum_{b=1}^B \left(\sum_{m=0}^7 \sum_{n=0}^7 (I_{r,b}(m, n)L(m, n))^2 \right) \quad (12)$$

where b is an index of the block with a total of B blocks. Let R measure the amount of MSE in the restored image, with

$$R = \sum_{b=1}^B \left(\sum_{m=0}^7 \sum_{n=0}^7 (I_{r,n,b}(m, n)Q_d(m, n))^2 \right) \quad (13)$$

where the subscript n again is used to denote the quantization noise. Using a Lagrange multiplier formulation that attempts to balance the MSE while staying close to the amount of desired high frequency components, and simplifying, we have

$$\hat{Q}_d(m, n) = \frac{Q_d(m, n)D}{D + L^2(m, n)R} \quad (14)$$

Combining this with equation (12), therefore,

$$\hat{Q}_d(m, n) = \frac{a(m, n)Q_e(m, n)D}{D + L^2(m, n)R} \quad (15)$$

V. IMAGE RESTORATION WITH THE JPEG 2000 ALGORITHM

A. Motivation

The JPEG image compression is a very successful standard, as evident from its proliferation in many areas of image representation, storage, and transmission. Yet technology is constantly evolving, and new applications are being devised. The emergence and wide adoption of the Internet prompts new developments in image compression, with its demand for progressive transmission in quality and resolution. Electronic commerce also pushes for better image security, giving rise to many schemes for digital watermarking and steganography. Digital photography is also a new application to the JPEG compression standard, and is heavily demanding for a more efficient compression scheme with a relatively low complexity. Progressive coding of images is also a desirable quality because of the limited memory available in the camera, which calls for a need to discard the less significant information of the existing pictures to make room for a new photograph, if necessary.

To address for these needs, much research has shown that wavelet-based compression schemes can produce better image quality than DCT-based compression at the same bitrate, while providing many enhancements to the features desired. That is primarily because it is ideally suitable for multiresolution transmission due to its subband decomposition nature. At the same time, a breakthrough in using bit-plane coding allows for progressive decoding, and when combined with an integer wavelet basis, can provide a range of lossy to lossless transmission of an image. A more elaborate bitstream syntax could also allow for better region-of-interest coding.

As a response, the JPEG 2000 committee was formed to develop a new, advanced standardized image coding system to serve new applications. The JPEG 2000 standard is divided into two parts: part I defines the minimum set of functionality and features, while part II includes advanced techniques and algorithms. Details of the standard can be found in [9]. For our purpose here, we note that although wavelet is much different from DCT, they both fit the blur model depicted in Fig. 3. In the case of JPEG 2000, we try to alter the wavelet transform coefficients given a certain wavelet basis.

B. Algorithm

Similar to the DCT domain image restoration algorithm, our wavelet domain restoration also takes advantage of the flexibility of using different quantization for compression and decompression. We begin with a reference quantization step sizes for the discrete wavelet transform (DWT) coefficients. This step size can be different for different subbands. Using $I_o(m, n)$ and $I_r(m, n)$ from the training data set, we can use equation (11) to best value of $\alpha(m, n)$, in the mean-square sense. Again, if we have some estimate of the blur, we can improve the algorithm with the use of $\hat{I}_o(m, n)$ similar to the previous section.

The dequantization is somewhat different. In the standard, the process is governed by the equation

$$\hat{I}(m, n) = \{I_q(m, n) + \alpha \operatorname{sgn}(I_q(m, n))\}Q_d(m, n), \quad (16)$$

where $\operatorname{sgn}(x)$ denotes the sign of x . It is zero if $x = 0$. α is a design parameter that we can make use of here to further improve the restored image quality. Ordinarily, we can pick $\alpha = 0.5$ in the absence of knowledge about the probability density function of q . However, we know that I , the unquantized DWT coefficient, is usually modeled as Laplacian [15], i.e.,

$$p(I) = \frac{\mu}{2} e^{-\mu|I|}. \quad (17)$$

(Note that this does not mean the quantized coefficient I_q is also Laplacian.) As such, a smaller value of α can reduce the overall distortion.

The question arises as to how to estimate μ . If we have the values of $I(m, n)$ available, we can compute the maximum likelihood (ML) estimation. Let $l(\mu)$ denote the log likelihood. With the probability density function in equation (17), we have

$$l(\mu) = \sum_{m=1}^M \sum_{n=1}^N (\log \mu - \log 2 - \mu |I(m, n)|) \quad (18)$$

Setting $l'(\mu) = 0$, we have

$$\mu = \frac{MN}{\sum_{m=1}^M \sum_{n=1}^N |I(m, n)|} \quad (19)$$

When we only have $I_q(m, n)$ available instead of $I(m, n)$, we can still compute the ML estimation with a probability mass function

$$p(I_q) = \begin{cases} \int_{-1}^1 \frac{\mu}{2} e^{-\mu|I|} dI & I_q = 0 \\ \int_{I_q}^{I_q+1} \frac{\mu}{2} e^{-\mu|I|} dI & I_q > 0 \\ \int_{I_q-1}^{I_q} \frac{\mu}{2} e^{-\mu|I|} dI & I_q < 0 \end{cases} \quad (20)$$

The ML estimation still has an analytical expression, although the expression is somewhat cumbersome. A similar derivation can be found in [16].

Once we have an estimate of μ , we can use it to find the centroid of each partition, which is the optimal decoding according to the Lloyd-Max condition [4]. This is derived from the equation

$$\hat{I} = \frac{\int_{I_q Q_d}^{(I_q+1)Q_d} I p(I) dI}{\int_{I_q Q_d}^{(I_q+1)Q_d} p(I) dI} \quad (21)$$

After some arithmetic manipulation, we have

$$\alpha = 0.5 - \left(\frac{1}{2} \coth \left(\frac{\mu Q_d}{2} \right) - \frac{1}{\mu Q_d} \right) \quad (22)$$

Note that $0 < \alpha \leq 0.5$ [11].

VI. CONSIDERATIONS FOR IMPLEMENTATION AND EXTENSIONS

In implementing the above algorithms, we need to have training images to assist us in adjusting the quantization levels. Depending on applications, this may or may not be totally feasible. One alternative would be to have models of the blur and simulate the adjustments in quantization levels necessary. The advantage that these algorithms bring to us is then an efficient way of performing the restoration without actually computing a filtering operation, such as that in equation (4). We can pre-compute some common blurs, such as defocusing, and store the quantization changes in the camera.

Another point to note about these algorithms is that they provide a basic setting for restoration implemented together with the compression. There are at least two areas where this work can be extended:

1. *Restoration of images with linear space-variant aberrations.* As explained earlier, a number of optical aberrations are space-variant, and some can be converted to space-invariant degradations by suitable geometric coordinate transformations. To incorporate these in our restoration scheme, we need to include both the forward and inverse coordinate transformations before and after the linear space-invariant equivalence of the blur. The challenge is then to devise algorithms that can embed these additional operations in the transform domain used for compression.
2. *Restoration of color images.* As it stands, our algorithm deals with monochrome images, and therefore should only be applied on the luminance channel of a color image. Since almost all consumer level digital cameras produce color pictures, it is important to extend our algorithms to color images. Color image processing requires much more sophistication than processing of the three color planes separately. First of all, one needs to be concerned with how the color image is acquired, for instance, the particular demosaicing algorithm associated with a certain color filter array. Second, one should have a prudent choice of the color space to carry out the restoration. Traditionally, a luminance-chrominance color space such as YUV or YC_bC_r , is preferable to RGB or CMY color spaces. Finally, whatever the color space, the aberrations on the three color planes are related, and this provides extra information for image restoration. How that information can be put to good use is itself a large area for further studies.


VII. CONCLUSION

The digital photography industry is experiencing exponential growth as digital cameras become more viable competitors with film-based cameras. The push for advancement in technology has been extensive and diverse,

with a lot of research focusing on increasing the resolution of the sensor, designing new shapes of the pixels, or improving the readout of the image for faster frame rate. In this paper, we concentrate our efforts to show that digital photography allows complex calculations using the pixel intensities, which enables us to implement image processing algorithms inside a digital camera to restore the quality of degraded images. This would be a definitive advantage over film-based photography.

REFERENCES

- [1] D. Yang, B. Fowler, and A. El Gamal, "A Nyquist rate pixel level ADC for CMOS image sensors," *IEEE J. Solid State Circuits*, vol. 34, no. 7, pp. 348-356, Mar. 1999.
- [2] J. W. Goodman, *Introduction to Fourier Optics*, 2nd ed., McGraw-Hill, New York, 1996.
- [3] K. R. Castleman, *Digital Image Processing*. Prentice Hall, Englewood Cliffs, New Jersey, 1996.
- [4] A. Gersho and R. Gray, *Vector Quantization and Signal Compression*. Kluwer Academic Publishers, Boston, 1992.
- [5] E. Y. Lam and J. W. Goodman, "Discrete cosine transform domain restoration of defocused images," *Appl. Optics*, vol. 37, no. 26, pp. 6213-6218, Sep. 1998.
- [6] H. H. Hopkins, *Wave Theory of Aberrations*. Oxford University Press, Oxford, 1950.
- [7] E. Y. Lam and J. W. Goodman, "An iterative statistical approach to blind image deconvolution," *J. Opt. Soc. of Amer. A*, vol. 17, no. 7, pp. 1177-1184, Jul. 2000.
- [8] W. Pennebaker and J. Mitchell, *JPEG Still Image Data Compression Standard*. Van Nostrand Reinhold, New York, 1992.
- [9] D. Taubman and M. Marcellin, *JPEG 2000: Image Compression Fundamentals, Standards and Practice*. Kluwer Academic Publishers, Boston, 2001.
- [10] D. Sheppard, A. Bilgin, M. Nadar, B. Hunt, and M. Marcellin, "A vector quantizer for image restoration," *IEEE Trans. on Image Proc.*, vol. 7, no. 1, pp. 119-124, Jan. 1998.
- [11] E. Y. Lam, "Digital restoration of defocused images in the wavelet domain," *Appl. Optics*, vol. 41, no. 23, pp. 4806-4811, Aug. 2002.
- [12] K. Rao and P. Yip, *Discrete Cosine Transform: Algorithms, Advantages and Applications*. Academic Press, New York, 1990.
- [13] R. Prost, Y. Ding, and A. Baskurt, "JPEG dequantization array for regularized decompression," *IEEE Trans. on Image Proc.*, vol. 6, no. 6, pp. 883-888, Jun. 1997.
- [14] W. Philips, "Corrections to 'JPEG dequantization array for regularized decompression'," *IEEE Trans. on Image Proc.*, vol. 7, no. 12, pp. 1725, Dec. 1998.
- [15] S. Mallat, "A theory for multiresolution signal decomposition: the wavelet representation," *IEEE Trans. on Patt. Anal. and Mac. Intel.*, vol. 11, no. 7, pp. 674-693, Jul. 1989.
- [16] J. R. Price and M. Rabbani, "Biased reconstruction for JPEG decoding," *Sig. Proc. Lett.*, vol. 6, no. 12, pp. 297-299, Dec. 1999.



Edmund Y. Lam (S'97-M'00) received the B.S. degree in 1995, the M.S. degree in 1996, and the Ph.D. degree in 2000, all in electrical engineering from Stanford University, U.S.A. At Stanford, he was a member of the Information Systems Laboratory, conducting research for the Stanford Programmable Digital Camera project. His focus was on developing image restoration algorithms for digital photography. Outside Stanford, he also consulted for industry in the areas of digital camera systems design and algorithms development. Before returning to academia, he worked in the Reticle and Photomask Inspection Division (RAPID) of KLA-Tencor Corporation in San Jose as a senior engineer. His responsibility was to improve on the core die-to-die and die-to-database inspection algorithms, especially for phase shift masks. He is now an Assistant Professor of Electrical and Electronic Engineering at the University of Hong Kong. His research interests include optics, defect detection, image restoration, and superresolution.