Digital Image Forensics Through the Use of Noise Reference Patterns

K.L.Bouman* N.Khanna, E.J.Delp
School of Electrical and Computer Engineering
Purdue University
West Lafayette, IN 47907, USA

Abstract
As result of the growing use of digital cameras, methods to identify an image's origin quickly, reliably, and inexpensively are needed. By designing and implementing methods for mitigating or "attacking" forensic features in the images, the robustness of various forensic identification techniques can be analyzed. The objective of this research was to see if drastically modified images could still be identified using a camera's defects in the form of a reference noise pattern. Experimental data has been gathered using several digital cameras and has been placed in an image database to organize and manage the gathered experimental data. Various types of filters, including blurring, weighted blurring, sharpening, histogram equalization, and pseudo-random noise filters, have been implemented and used to process various images. These images were then used to see if the image's source camera is still identifiable after attacking the image. A reference pattern is found by averaging the noise pattern found in multiple images obtained from the same camera. This intrinsic signature of the camera can be used to correlate with a noise pattern of an image with an unknown origin. If the patterns are similar, and the correlation is above a certain threshold, the camera containing that particular reference pattern is the source camera. All images' source camera were identified correctly, and except for a few exceptions, most correlation values for attacked images were lower, and therefore had less similar noise patterns than the original images'.

Introduction and Methods

Digital image quality has become so high in the last few years that not only has the general public been rapidly replacing classical analog cameras with digital cameras, but law enforcement agencies are doing so as well. However, with the availability of powerful editing programs, it is very easy for an amateur to modify digital media and create realistic looking forgeries.

In film photography, methods for camera identification have been perfected. Small scratches on negatives are one of the many ways to identify an image captured on an analog camera. However, because of the growing use of digital cameras, methods to identify an image's origin quickly, reliably, and inexpensively are needed. Forensic tools that help establish the origin and authenticity of digital images are essential to a forensic examiner. These tools can prove to be vital whenever questions of digital image integrity are raised [1]. Although an image is originally accompanied with a large amount of data about the image in the EXIF header, this header may not be available if the image was saved as a format other than JPEG or recompressed [2].

A noise pattern can be used as a watermark for images. Like analog cameras, defects in the image are used to determine which camera produced the image. Each camera has non-uniformities ranging from dust specks on the optics, to dark current. This noise is relatively stable over the camera's life span and can be used to determine the source of an image.

One possible way to identify the source of a digital image is by using the pattern of hot or dead pixels, which give luminance values of white or black pixels respectively regardless of the images content. However, if the camera did not contain any defective pixels, or they were eliminated by processing after the image was taken, the source can not be identified.

Although alone hot and dead pixels are not reliable for identifying images, they can be used in conjunction with other defects to create a more reliable "watermark". Hot and dead pixels fall into the range of noise caused by array defects, which include point defects, hot point defects, pixel traps, column defects, and cluster defects, all of which alter individual pixel values dramatically.

Pattern noise can also be mapped for cameras. Any spatial pattern that does not vary from frame to frame can be used as a "watermark" for an image [1]. Such patterns include dark currents (dark meaning the current was formed despite no exposure to light), and photo response non-uniformity noise (PRNU). Dark current arises from an excess of electrons that are captured by the sensors for various reasons and calculated as part of the total signal. These brighter areas (since there are more electrons captured at those particular points), defined as the fixed pattern noise (FPN), are due to variations in detector size, and foreign matter trapped during fabrication of the sensor. The FPN can be detected when the senor is not illuminated. Dark current increases with temperature, since electrons become

* Corresponding author: klbouman@ecn.purdue.edu
Proceedings of the 2006 SURF Confrence

more active with heat [1,3]. Dark current also increases with exposure time [1]. To suppress this noise some camera manufacturers subtract a dark frame from every image taken, and therefore can be unreliable [3]. PRNU on the other hand is always present and is the variation in pixel responsively that can be detected when the sensor is illuminated [1]. This noise is related to detector size, spectral response, thickness in coatings and other small imperfections made when the camera was manufactured [1,4].

Let $y_{ij} = f_{ij}(x_{ij} + \eta_{ij}) + c_{ij} + \varepsilon_{ij}$ define the digitized output of the sensor (y_{ij}), where $i = 1...m$, $j = 1...n$, and $m \times n$ is the image resolution (m pixels vertically, n pixels horizontally). This is found by multiplying the raw data (x_{ij}) and the random shot noise (η_{ij}) by the PRNU noise factor (f_{ij}) and adding the additive random noise (ε_{ij}) and dark current (c_{ij}). The factors (f_{ij}) are different for all cameras, however, since PRNU is image dependent, some areas of noise may be largely suppressed due to a pixel in the image (x_{ij}) being approximately 0 [1].

The noise component of an image is subtracted from the original image by the use of a wavelet based denoising filter (F_σ). The value of σ determines the strength of the filter. The filter extracts Gaussian noise with a variance of σ^2. The filter parameter (σ) found in [1] to have the best performance is 5(Figure 1.3), since the gathered noise is not found to be too highly image dependent (Figure 1.2 with σ=15), or so small that finding correlation between reference patterns becomes hard to determine (Figure 1.4 with σ=2).

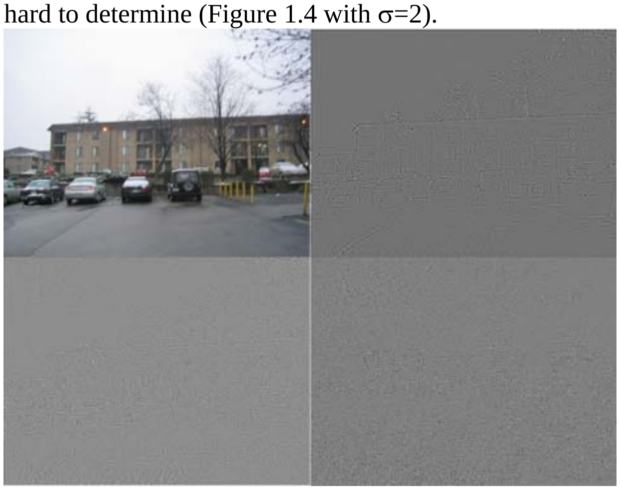

Fig 1.1 (top left) Original Image
Fig 1.2 (top right) Noise Pattern with σ of 15
Fig 1.3 (bottom left) Noise Pattern with σ of 5
Fig 1.4 (bottom right) Noise Pattern with σ of 2 Images 1.2-1.4 have been contrast enhanced by 100% Images 1.2-1.4 From Nitin Khanna

On the silicon chips that capture and record the light that enters a digital camera, many light sensitive areas called photo-sites are laid out in an array to capture the charge of accumulated electrons that are produced when light hits the sensor. However, since these photo-sites are just capturing how many electrons are released, or the intensity of light, they are insensitive to any differences in wavelength [2,3,5,6]. Therefore, most sensors need a method implemented to obtain the value of multiple colors' intensity at each pixel. Rather than being sensitive to all the different wavelengths of light, a sensor only needs to be monochromatic and sensitive to each of the primary colors of light, red, green, and blue, or sometimes green and the three secondary colors of light, cyan, magenta, yellow. With these colors, it is possible to create any color by simply overlapping the three or four values [3,5,6]

One of the most practical and common methods used for creating colored images in digital cameras is by using a color filter array (CFA). In this type of sensor, a photo-site is broken up into red, blue, and green filters or cyan, magenta, yellow, and green filters [3,5,6]. After the array of data is read by the sensor, the data is run through algorithms in the camera's software to merge three or four intensity values from each pixel into one color value. The most common type of color filter is the bayer filter pattern [3,5,6]. This pattern uses the primary colors of light, red (R), green (G), and blue (B), as filters for each photo-site [5]. The bayer filter involves alternating an odd number of rows of red and green filters with an even number of rows of blue and green filters [3]. There are 25% more green pixels than there are red or blue, since human visual system are more sensitive to small changes in green wavelengths than red or blue. This extra green filter helps to estimate the green's luminance data appropriately for the human visual system [5,6].

The method used to identify a camera's noise pattern is determined by subtracting a denoised version of an image, by performing a wavelet based denoising filter, from the original image. Since there are 25% more green channel luminance values in a bayer filter, there is more noise in the green channel than in either the red or blue channels. For this reason, the noise pattern is found in the images' green channel. The reference pattern is then found by averaging the noise pattern found in multiple pictures obtained from the same camera. This intrinsic signature of the camera can be used to correlate with a noise pattern of an image with an unknown origin. If the patterns are similar, and the correlation is above a certain threshold, the camera containing that particular reference pattern is the source camera.

When subtracting the denoised image from the original image there are small differences in pixel value. In order to find the similarity or correlation value

Noise Pattern X

.101	-.034	.256
.004	-.073	.419
-.066	.152	.001

Noise Pattern Y

-.082	.314	-.072
.093	-.598	-.722
.001	-.291	.099

Noise Pattern Z

.101	-.034	.256
.004	-.073	.419
-.066	.152	.001

Finding Correlation Values

$$x:y=\frac{(.101\times -.082)+(-.034\times .314)+(.256\times -.072)}{\sqrt{.101^2+(-.034)^2+.256^2}\times\sqrt{(-.082)^2+.314^2+(-.072)^2}}=$$

$$\frac{-.008282-.010676-.018432}{\sqrt{.010201+.001156+.065536}\times\sqrt{.006724+.098596+.005184}}=$$

$$\frac{-.03739}{\sqrt{.076893}\times\sqrt{.110504}}=\frac{-.03739}{.277296\times .332421}=\frac{-.03739}{.092179}=-.4056$$

$$x:z=\frac{(.101\times .101)+(-.034\times -.034)+(.256\times .256)}{\sqrt{.101^2+(-.034)^2+.256^2}\times\sqrt{.101^2+(-.034)^2+.256^2}}=$$

$$\frac{.010201+.001156+.065536}{\sqrt{.010201+.001156+.065536}\times\sqrt{.010201+.001156+.065536}}=$$

$$\frac{.076893}{\sqrt{.076893}\times\sqrt{.076893}}=\frac{.076893}{.277296\times .277296}=\frac{.076893}{.076893}=1.000$$

Fig 2: Demonstration on How to Find a Correlation Value Between Two Noise Patterns

between these noise patterns, each pixel in one noise image is multiplied by the corresponding pixel value in another noise image. The values found are then added together and divided by the square root of the sum of the squares of each pixel's noise value in the first image times the square root of the sum of the squares of each pixel's noise value in the second image. Refer to figure 2. Due to PRNU noise being multiplicative and image dependant, in darker areas, where denoising is less effective, the correlation will be lower although the image has not been tampered with.

Since the preferred method of capturing colored images is through the use of a single sensor in conjunction with a color filter array, only a third of the final image is captured by the camera. The other two thirds is interpolated. Thus, it is possible to identify an image's origin can be determined through estimation of the color interpolation parameters used by the source camera. However, this method cannot be used on images that have been compressed, such as in JPEGs, since compression artifacts suppress the correlation between the pixels created by the camera's interpolation [4].

Since a sensor only measures intensity, pixels are stored in a computer as a value that describes how bright the pixel is. In binary images, the pixel value is a 1-bit number (0 or 1). Since there are only two possible one bit numbers, these images are often displayed as black and white. A grayscale image on the other hand contains eight bits (one byte) of information per pixel. Since every bit can contain up to two values, the number of possible values in a byte is 2^8 or 256. Since a grayscale images only needs one value to represent the brightness of the pixel, each pixel contains a number from 0 to 255 (256 possible values) where 0 represents black, 255 represents white, and all values in the middle are shades of gray. A color image, on the other hand, is made up of three different values, red, green, and blue, each having a range from 0 to 255. When these three values are placed together in an image they make a colored image [7].

Filters are used to attack pixel values usually by assigning them new values based on adjacent pixels. When attacking a colored imaged these three values have to be determined separately. Four different examples of filters are described below.

A blurring filter is a three by three array of weighting coefficients where each coefficient is assigned a weight of one. Refer to figure 3.1. The weight for each pixel is multiplied by its intensity value (in this case since they all have a coefficient of one, the values would stay the same), and then divided by nine (when the pixel is on an image's edge the value changes) to find the average of the pixels contained in the filter. This new average value is then assigned to the middle pixel (shown here as bold). This is done for every pixel in the image. The resulting image is smooth when compared to the input image. A weighted blurring filter is very similar to a blurring filter except the coefficients are different. More emphasis is placed on the pixel value that is being changed and the adjacent pixel values rather than the pixel values diagonal from the pixel being changed. The weighting coefficients used for this filter in experimentation can be seen in Figure 3.2. Instead of dividing by nine, after the pixel's values are multiplied by their coefficients and added together they are divided by the sum of the coefficients, in this particular filter sixteen. A weighted blurring filter, like a blurring filter, makes an image smoother, however more detail is preserved.

1	1	1
1	**1**	1
1	1	1

1	2	1
2	**4**	2
1	2	1

Fig 3.1 (left): Weighting coefficients for a blurring filter
Fig 3.2 (right): Weighting coefficients for a weighted blurring filter

A sharpening filter makes use of a weighted blurring filter. The sharpened valued is found by the equation

$$I(m,n)+\lambda(I(m,n)-\left(\sum_{j=-1}^{1}\sum_{k=-1}^{1}h(j,k)\times I(m-j,n-k)\right)$$

where m identifies the row the pixel is in, n identifies the column, j and k equal [-1, 0, 1] to represent a position in the filter, $\left(\sum_{j=-1}^{1}\sum_{k=-1}^{1}h(j,k)\times I(m-j,n-k)\right)$ is the corresponding pixel in the weighted blurred image, and λ is any value the programmer decides on, usually around 1, however the value of 2 was used in order to change the pixel values more drastically [8].

A histogram of an image graphs pixel intensity (0 to 255) by the number of pixels that are each value. A histogram equalizing filter changes pixel intensities so that the histogram of the resulting image is more uniform. This is done by first finding the number of pixels in an image. The number of pixels at each intensity value is then divided by the total number of pixels in the image. This is done for every intensity value and then is placed into an array containing 256 values. A new array of 256 values is created by adding the fractions in the original array until reaching the intensity value that is being recorded in the new array. For example, for the newer array's 5th position, the values in the 0th, 1st, 2nd, 3rd, 4th, and 5th positions of the original array are added together and assigned to the newer array's 5th position. This is done for every pixel value (0 to 255) until the array's 256th term, which is assigned 1. These new values are multiplied by 255, and the new values of each pixel intensity are assigned to the pixels containing the original intensity [9].

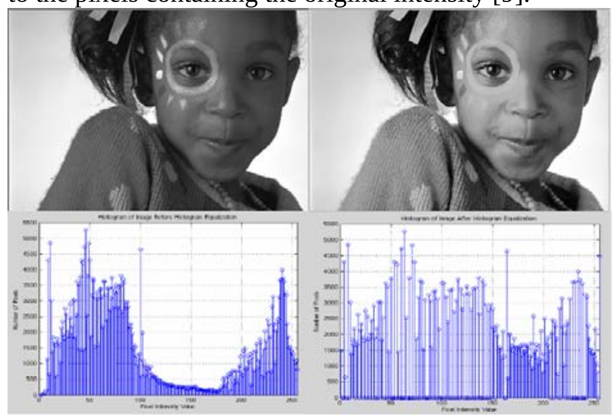

Fig 4.1 (top left): Original Image [8]
Fig 4.2 (bottom left): Histogram of Original Image
Fig 4.3 (bottom left): Image 4.1 after histogram equalization
Fig 4.4 (bottom right): Histogram of Image 4.3

Noise was added to an image's green channel through the use of a pseudo-random generator, causing the noise to be randomly distributed throughout the image.

These five filters were applied to 40 images acquired from four different cameras. These 240 resultant images were placed in the denoising filter and then correlation values were found between each image and their camera reference pattern.

Specific Objectives

The objective of this research was to see if drastically modified images could still be identified using a camera's defects in the form of a reference noise pattern. A higher correlation value was expected to be resultant from between an image and its source camera's pattern than between other camera patterns, yet lower correlation was expected to be achieved for a filtered image than the original image for authentication purposes. In addition, the impact a filter has on an image when it comes to identifying the source camera and discovering a method to possibly identify what modifications were made to an image were explored.

Results and Discussion

The correlation value method described above was used to find the similarity between 240 images, 200 of those having been modified, and four different possible noise patterns from four different cameras. Each camera had a noise pattern that was determined by taking the average noise patterns of 240 images from that particular camera.

Fig 5.1:The Averaged Camera Correlation Values, Total Averages of Correlation Values, Standard Deviations and Statistical Significances of Each Method of Filtered Attack						
	Original Image	Blurring	Weighted Blurring	Histogram Equalization	Sharpening	Pseudo-random Noise
Camera1	0.1353814	0.047144	0.1283843	0.0114812	0.1390119	0.079597
Camera 3	0.0788258	0.0322505	0.0681803	0.0107104	0.0726947	0.0486451
Camera 5	0.1247873	0.0506303	0.1185903	0.0149163	0.1139722	0.0762584
Camera 7	0.2016535	0.075928	0.1964297	0.0099474	0.2157401	0.1205999
Average	0.135162	0.0514882	0.12789615	0.01176383	0.13535473	0.0812751
Standard Deviation	0.04858858	0.02139085	0.05225711	0.00445059	0.05741372	0.03110314
Ratio of Standard Deviation to Average	41.9246565%	47.7484918%	53.108978%	38.6168234%	54.7276819%	39.0385009%

Fig 5.2:The Averaged Camera Correlation Values, Total Averages of Correlation Values, Standard Deviations and Statistical Significances of Each Method of Filtered Attack After the Division of Each Correlation Value by the Correlation Value of the Original Image						
	Original Image	Blurring	Weighted Blurring	Histogram Equalization	Sharpening	Pseudo-random Noise
Camera1	1	0.34921249	0.94456087	0.08546658	1.02773684	0.58905234
Camera 3	1	0.34921249	0.94456087	0.08546658	1.02773684	0.58905234
Camera 5	1	0.41082469	0.94974442	0.11818338	0.90876352	0.61944589
Camera 7	1	0.37135714	0.971526	0.05031676	1.07062071	0.59385376
Average	1	0.38247865	0.93115751	0.09787313	0.98099923	0.60270173
Standard Deviation	0	0.1004961	0.09075903	0.05152614	0.104143	0.10419966
Ratio of Standard Deviation to Average	0%	26.2749584%	9.74690456%	52.645845%	10.6160128%	17.2887604%

An image's source camera was determined by finding the camera reference pattern that had the highest correlation value with that image. All images', including the modified images, source camera was correctly identified. As described above, the closer a correlation value is to 1, the more similar it is to the camera's noise pattern. Therefore, the highest correlation values for each of the images were sorted by method of modification and then averaged. The arithmetic mean for each camera and the total arithmetic mean of the image's correlation values for the original, blurred, weighted blurred, histogram equalized, sharpened, and pseudo-random noise images, are found in figure 5.1. The standard deviation, which is a measure of how widely values are dispersed from the arithmetic mean, is found by the equation

$$\sqrt{\frac{n\sum x^2-\left(\sum x\right)^2}{n(n-1)}}$$

where n is the number of items in the data set, and x is a data set's deviation from the mean. Since the correlation values of the original images have a standard deviation of .0486, there is a lot of variance in the data. Therefore, it is easier to analyze the data when every value is divided by the value obtained when comparing the original image with its corresponding reference pattern, such as in figure 5.2. This assumes that every original image's noise pattern is a facsimile to the averaged camera noise pattern. As can be seen in both tables, the correlation values of the original, weighted blurred images, and the sharpened images are closer to 1 than any of the other filtered images' values. The original image supposedly should have the highest correlation value, considering there has been no altering of pixel values. However, as can be seen in the averages for the sharpening filter in figure 5.2, some sharpened images had a more similar noise pattern to the reference pattern than the original image did. This may have been caused because the sharpening filter makes the noise more prominent, by subtracting the original image's "blurred" or unwanted image. In addition, both the weighted blurring and sharpening filter, since the sharpening filter makes use of a weighted blurring filter, have placed a larger weight on the pixel value that was changed rather than the surrounding pixel values. This in turn means that each pixel intensity value, in each location, although having been changed, has been changed to a value that places most of its emphasis on the original value of the pixel at that location. Therefore the noise pattern values will be more similar than if they were given an equal or lesser weight than the other values that were being used to calculate the new value of the pixel.

Fig 2.3: The Average and Standard Deviation of Each Camera's Images When Compared to Four Different Reference Patterns					
		Camera 1 Reference Pattern	Camera 3 Reference Pattern	Camera 5 Reference Pattern	Camera 7 Reference Pattern
Camera 1	Average	0.09016663	0.00663237	0.00309788	0.0046133
	Standard Deviation	0.05103388	0.00493683	0.00207572	0.00448122
Camera 3	Average	0.00379927	0.05188447	0.00247983	0.00147783
	Standard Deviation	0.00297967	0.02818761	0.00178452	0.00151231
Camera 5	Average	0.00218258	0.00374062	0.08319247	0.00172957
	Standard Deviation	0.00117863	0.00248705	0.04331861	0.00124896
Camera 7	Average	0.00639413	0.0063168	0.00470943	0.13671643
	Standard Deviation	0.00606546	0.00481828	0.00280582	0.08023762

Although the average correlation values for different methods of attack seem quite a bit different, the variance of the data is high, making it impossible to be able to specifically assign a correlation value with a method of attack without more collected data.

As can be seen in figure 5.3 the averages of the correlation values for a set of images when compared to the reference patterns of camera 1, 3, 5, and 7 that were created by averaging 240 noise patterns, are shown. The values in yellow are the values corresponding to the correct reference pattern for each set of images. The averages in yellow are much larger than those shown in white, meaning they are more similar to the correct reference pattern than the incorrect reference patterns for that particular set of images. In fact, in every case, the set of values that make up the yellow values were found to be significantly larger than those that make up the white values, with a negligible probability that the results and resultant means were found just by coincidence (ranging from 1.26×10^{-18} to 1.71×10^{-21}).

There are a few possible variables that could have affected the resultant data. Firstly, since the denoising filter only detects noise in the image's green channel, it cannot detect any modification in the image through either the red and blue channels. If significant modifications are made solely in the green channel, it may be more difficult to identify the source camera. In the filter used to add pseudo-random noise to an image, noise was only added to the image through the green channel. Therefore, this filter did not affect either the red or blue channels, and if the denoising filter found noise in either of those channels the correlation value would be exactly the same for that image and the original image, although they are in fact very different.

In order to create a camera reference pattern, many images from that camera need to be denoised and then averaged into one noise pattern. In this research 240 images were used to create each reference pattern. Therefore, in order to have accurate data to compare images to, the source camera must be in working order and many images must be obtained from it before any correlation can be determined. In addition, it would be difficult to determine if an image is from a camera if that is the only camera that its noise pattern is being compared to. Instead, an image must be compared to multiple reference patterns in order to be sure it has been identified correctly.

Although more data should be collected at this point from different cameras to verify the results found. The gathered results suggest that correlation values will always be larger between an image noise pattern and it's source camera's noise reference pattern than a reference pattern from a camera other than it's source camera no matter what filter has been implemented on the image.

The data collected on correctly identifying source digital cameras is being used currently to extrapolate this method to identification of digital scanners.

Summary and Conclusions

A method of correctly identifying the source camera of an image by the use of the image's source pattern was implemented. Five filters, blurring, weighted blurring, sharpening, histogram equalization and pseudo-random noise filters were implemented in C, and applied to 40 images, coming from four different source cameras. The noise patterns of these resultant 240 images were found by the use of a denoising filter with a filter parameter (σ) of 5. These noise patterns were compared to 4 camera reference noise patterns, which were created by taking an average of the noise patterns of 240 images coming from each camera. A correlation value was found for each comparison.

It was found that every image's source camera was identified correctly. The correlation value of image and its source camera's reference pattern was found to be significantly greater than the correlation value of image and a reference pattern other than from its source camera. This strongly suggests that this method of camera identification can be used to identify most images, even after drastic modification.

References

[1] J. Lukas, J. Fridrich, M. Goljan, "Determining Digital Image Origin Using Sensor Imperfections." 2005.

[2] M.R.S, Alken, "How Digital Cameras Work." http://www.alkenmrs.com/digital-photography/how-digital-cameras-work.html. 2006.

[3] T. Hogan, "How Digital Cameras Work". http://www.bythom.com/ccds.htm. 2004.

[4] N. Khanna, "A Survey of Forensic Characterization Methods for Physical Devices." 2006.

[5] J.Adams, K. Parulski, K. Spaulding, "Color Processing in Digital Cameras". IEEE. 1998.

[6] M. Kudenov, "Charged Coupled Devices (CCD's)." http://ffden-2.phys.uaf.edu/212_fall2003.web.dir/Mike_Kudenov%20/CCD.htm. 2003b.

[7] Fisher, Bob, et al, "Pixel Values". http://www.cee.hw.ac.uk/hipr/html/value.html. 1994.

[8] EE637: Digital Signal Processing I. "Digital Image Processing Laboratories." http://dynamo.ecn.purdue.edu/~bouman/grad-labs/. 2005.

[9] L.C. Mai, "Histogram Equalization." http://www.netnam.vn/unescocourse/computervision/22.htm. 2000.