Digital Camera Sensor Noise Estimation from Different Illuminations of Identical Subject Matter

Corey Manders
University of Toronto
Dept. of Electrical and Computer Engineering
10 King's College Rd.
Toronto, Canada.
manders@ieee.org

Steve Mann
University of Toronto
Dept. of Electrical and Computer Engineering
10 King's College Rd.
Toronto, Canada
mann@eecg.utoronto.ca

***Abstract*— A noise model that is generated from pictures taken by differently illuminating the same subject matter is presented. The superposigram, a three dimensional structure based on three different illuminations of the subject matter is used to directly calculate the noise present in the images which is attributable to the noise originating in the sensor array. Though a similar method has been used to determine the dynamic range compression present in a particular camera, this paper focuses on using data in which no dynamic range compression has taken place (raw files from a Nikon D2h were used). The paper shows through a mathematical proof that the noise in the camera is well estimated by the superposigram and then shows empirically that this is true. Finally, the method is applied to four commercially available cameras, to evaluate the noise in each.**

***Keywords*—lightspace, superposimetric imaging, digital camera noise analysis.**

I. INTRODUCTION: A SIMPLE CAMERA MODEL AND THE SUPERPOSIGRAM

Several papers demonstrate that the dynamic range compression is present in virtually all cameras, for example [1] [2][3]. The fact is covered in great detail by Poynton in [4]. All papers use various methods to recover the response function from images taken with a given camera. However, [1] uses a three-dimensional "superposigram" structure to recover the response function using a singular-value decomposition. The paper shows several examples of superposigrams including superposigrams of images after undoing the range compression. Such a plot exhibits a plane defined by $x+y-z=0$ in three-dimensional space. A better example of this is demonstrated in [5] where superposigrams were used to show that images composed of 12-bit data from the raw data of a camera were in fact linear.

This paper presents a new method for determining the noise present in the sensor by analysing the perpendicular distances of points in the superposigram to the plane $x+y-z=0$. The paper will show through a mathematical proof the the variance of the noise in the camera is the same as the variance of the distances of points in the superposimetric plane as the number of points approaches infinity. The method is compared to another traditional method of finding the variance of the noise in which multiple images of the same subject matter with identical illumination is used. The new work is largely based on the concept of lightspace presented in [6].

Fig. 1: One of the many datasets used in the computation of the superposigram. Leftmost: Picture of Deconism Gallery with only the upper lights turned on. Middle: Picture with only the lower lights turned on. Rightmost: Picture with both the upper and lower lights turned on.

Research in this area may be broadly divided into two areas, comparametrics[7] and superposimetrics[1]. Both areas play upon this linearity of light, and exploit the following two properties of linearity:

1) Homogeneity: A function is said to exhibit homogeneity if and only if $f(ax)=af(x)$, for all scalar a.
2) Superposition: A function is said to exhibit superposition if and only if $f(x+y)=f(x)+f(y)$.

Comparametrics use the homogeneity property of linearly, and superposimetrics uses the property of superposition in which the principle tool is the superposigram (as opposed to the comparagram in comparametrics).

The construction of a superposigram is quite simple given a camera and two lights. An image is taken with the first light on, a second image is taken with the second light on, and finally a third image is taken with both lights on. For all images, it is essential that the camera remains in exactly the same position. For added accuracy, the lighting should be $d.c.$ lighting or three-phase lighting rather than the usual single-phase $a.c.$ lighting. Single-phase lighting will diminish the quality of the result by nature of the inherent flicker created by the sinusoidal nature of the power. Three-phase lighting (a set of three lights with each light on a different phase), will greatly reduce any flicker or unevenness in the lighting. As one light dims corresponding to its phase, the other lights will supplement the intensity, analogous to the trigonometric identity $\sin^2(t)+\cos^2(t)=1$. Thus the illumination will be constant as a function of time, and artifacts caused by uneven illumination will be attenuated. Such images are shown in figure 1, where three phase lighting was used. In the case of this paper, the camera was securely fastened that no movement occurred from one image to the next. The only change which is allowed to occur in the different images is the amount of light falling on each element of the sensor array due to the structured changes in illumination.

The Nikon D2h camera was set to raw data mode. The raw data from each camera was confirmed to be linear (this is evident later from figure 4 where the superposigrams are noisy planes rather than noisy curved manifolds). Note that the raw files may be decoded using dcraw, Dave Coffin's raw file decoder which is freely available from his website at www.cybercom.net/∼dcoffin/dcraw, or neftoppm, available at www.eyetap.org/∼corey/code.html. In the case of using dcraw, which was used for Canon raw files, the Bayer interpolation was disabled and each red, green, or blue pixel was separated into individual twelve-bit ppms.

The superposigram (and the method) plays upon the linearity of light and the light-linear[4] raw file format available from the cameras used. The superposigram structure is constructed using a three-dimensional array. Each array element is initially zero. Each pixel location is considered individually in the three images yielding an array location. For example, if we were to consider the first pixel location, and the values of the pixels were 300,400, and 705 respectively (12-bit images), the array location (300,400,705) would be incremented by one. The superposigram plays upon the linearity of light to characterize properties of the output of the camera.

As shown in previous work[1], this information may be used to obtain the camera response function (or range compression) for a given camera. However, if we consider the output of raw data from a camera, where no range compression is present, this paper will show that the superposigram may be used to determine the noise present in the sensor array. Note that as 12-bit raw data is used, there is no noise from data compression such as JPEG compression, and quantization noise is reduced by using 12-bit data rather than the 8-bit data usually found in JPEG images or other file formats found as output from digital cameras.

After the method of noise estimation has been presented, the method will be applied to estimate the noise in a Nikon D2h, a Nikon D70, a Canon 1D mark II, and a Cannon 10D.

II. ARTIFICIALLY GENERATED IMAGES AND SUPERPOSIGRAMS

To begin to understand the effect of sensor noise on a superposigram, images may be constructed by using gradient patterns. For simplicity, two images were created using horizontal and vertical gradients. A third image was constructed by adding these two gradient images together. Then zero mean Gaussian noise was added to all three images with a standard deviation of 0.02. These images are depicted in figure 2.

Fig. 2: Three of the images used to test the effect of Gaussian noise on a superposigram. Each Gradient image was perturbed by zero-centered Gaussian Noise with a standard deviation of 0.02.

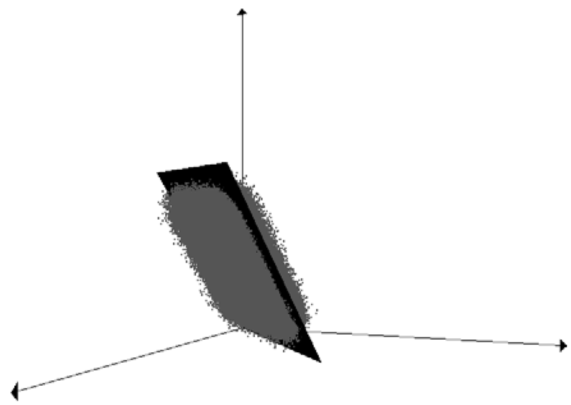

Fig. 3: The superposigram resulting from the images shown in figure 2. Each Gradient image was perturbed by zero-centered Gaussian Noise with a standard deviation of 0.02. The plane $x + y - z$ has been plotted in black to show the perturbation from the plane.

After generating the gradient images and adding Gaussian noise, a superposigram was constructed from the images. Note that noise from the first image will only perturb the superposigram along its first axis, the second image along the second axis, and the third along the third axis. Furthermore, the sensor used for each of the images is the same, therefore the noise for each of the images will follow the same noise model which we assume to have zero mean. Thus it is reasonable to assume that the noise in each of the three images will have largely identical properties (for example identical standard deviations). For this reason, the one-dimensional noise in each of the three images will result in three-dimensional noise and show up as a perturbation from the plane $x + y = z$ (or equivalently $x+y-z = 0$). Conveniently, the noise will show up as deviations from the plane. We will show that measuring the stand deviations of distances from the plane amounts to measuring the standard deviation of noise in the camera.

The superposigram resulting from the images in figure 2 is shown in figure 3.

In our artificially created example, as expected, the distances of the superposimetric points from the plane $x + y - z = 0$ are distributed with a mean of 0 and a standard deviation of 0.02.

III. QUANTIFYING NOISE IN DIGITAL CAMERAS USING THE SUPERPOSIGRAM

At first, it may not seem completely apparent how the distribution of distances from the plane are approximately (and very closely) related to the distribution of noise in the camera, consider that the noise in the camera will perturb a single data point in the superposigram in the x, y, and z direction with the same noise distribution.

The superposimetric distance theorem. *The variance of noise in the camera is equal to the variance of perpendicular distances of points to the superposimetric plane as* $n \to \infty$.

Proof: *Consider any point in the superposigram. Without noise purterbations, the point lies on the plane* $x+y-z=0$. *With independent perturbations in the* $x, y,$ *and* z *directions any point in the superposigram may be written as:* $(x+\Delta x, y+\Delta y, z+\Delta z)$. *Let the variance of the perturbations be* $\sigma(\eta)$. *Given the standard measure of a point* $X=(x_0, y_0, z_0)$ *from a plane* $ax+by+cz+d=0$ *using the perpendicular normal may be stated as:*

$$D = \frac{ax_0+by_0+cz_0+d}{\sqrt{a^2+b^2+c^2}}. \tag{1}$$

From the fact that our plane is defined by the equation $x+y-z=0$, *it follows that any point in the superposigram may be written as a distance from the plane:*

$$d_\rho = \frac{1}{\sqrt{3}}\left((x_\rho+\Delta x_\rho)+(y_\rho+\Delta y_\rho)-(z_\rho+\Delta z_\rho)\right) \tag{2}$$

Let the variance of the distances from the plane be $\sigma(d)$. *If we calculate the variance of the distances we have:*

$$\sigma(d) = \sum_{i=1}^{n} \frac{(d_i-\bar{d})^2}{n-1}. \tag{3}$$

However, we expect the noise to have mean = 0, thus $\bar{d}=0$. *Also, given that the point unperturbed by the noise is on the plane* $x_1+x_2-x_3=0$, *the variance of the distances may be written as:*

$$\sigma(d) = \sum_{i=0}^{n} \frac{(\frac{1}{\sqrt{3}}\Delta x_i+\Delta y+\Delta z_i)^2}{n-1} \tag{4}$$

$$= \frac{1}{3}\sum_{i=1}^{n} \frac{(\Delta x_i)^2+(\Delta y_i)^2+(\Delta z_i)^2}{n-1} + + \frac{1}{3}\sum_{i=1}^{n} \frac{\Delta x_i\Delta y_i-\Delta x_i\Delta z_i-\Delta y_i\Delta z_i}{n-1} \tag{5}$$

$$= \sigma(\eta) + \frac{1}{3}\sum_{i=1}^{n} \frac{\Delta x_i\Delta y_i-\Delta x_i\Delta z_i-\Delta y_i\Delta z_i}{n-1} \tag{6}$$

$$= \sigma(\eta) + \frac{1}{3}\sum_{i=1}^{n} \frac{\Delta x_i\Delta y_i}{n-1} - \frac{1}{3}\sum_{i=1}^{n} \frac{\Delta x_i\Delta z_i}{n-1} - \frac{1}{3}\sum_{i=1}^{n} \frac{\Delta y_i\Delta z_i}{n-1} \tag{7}$$

$$. \approx \sigma(\eta) + \frac{1}{3}E(\Delta x\Delta y) - \frac{1}{3}E(\Delta x\Delta z) - \tag{8}$$

$$\frac{1}{3}E(\Delta y\Delta z). \tag{9}$$

Note that the line 6 assumes that the variances of the perturbations along each axis are equal. We also know that the expectations of the terms $\Delta x, \Delta y,$*and* $\Delta z = 0$. *Thus, by the property of expectation,* $E(X)E(Y)=E(XY)$, *the expectations of* $\Delta x\Delta y$, $\Delta x\Delta z$, *and* $\Delta y\Delta z$ *are all zero. Finally, note that as* $n \to \infty$, *the error due to the approximation on line 8 tends to 0. Thus, the variance of the distances,* $\sigma(d)$, *tend to the variance of the noise,* $\sigma(\eta)$, *as* n *increases. This implies the variance of the noise in the camera is equal to the variance of perpendicular distances of to the superposimetric plane as* $n \to \infty$. □

Knowing that it is now possible to quantify noise using the superposigram and distances from the plane $x+y-z=0$, we may use the method to quantify the noise in each of the colour channels which make up the sensor.

The method is first used with a Nikon D2H. The superposigrams from each of the colour channels of a Nikon D2H are shown in figure 4, and the histograms corresponding to the resulting distances from the plane $x+y-z=0$ are shown in figure 5.

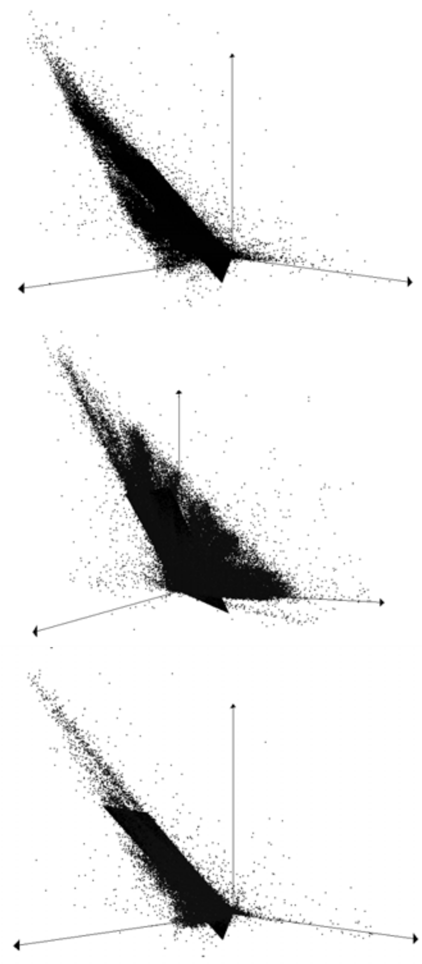

Fig. 4: Superposigrams from the red, green, and blue channels of a Nikon D2H. The superposigram of the red channel is depicted in the top superposigram, the green channel in the middle, and the blue channel on the bottom. The plane $x+y-z=0$ has also been drawn for additional clarity.

IV. VERIFYING RESULTS

To verify the results, the method was compared to a more traditional means of finding the variance of the noise. Specifically, multiple images (approximately 1000), were taken of

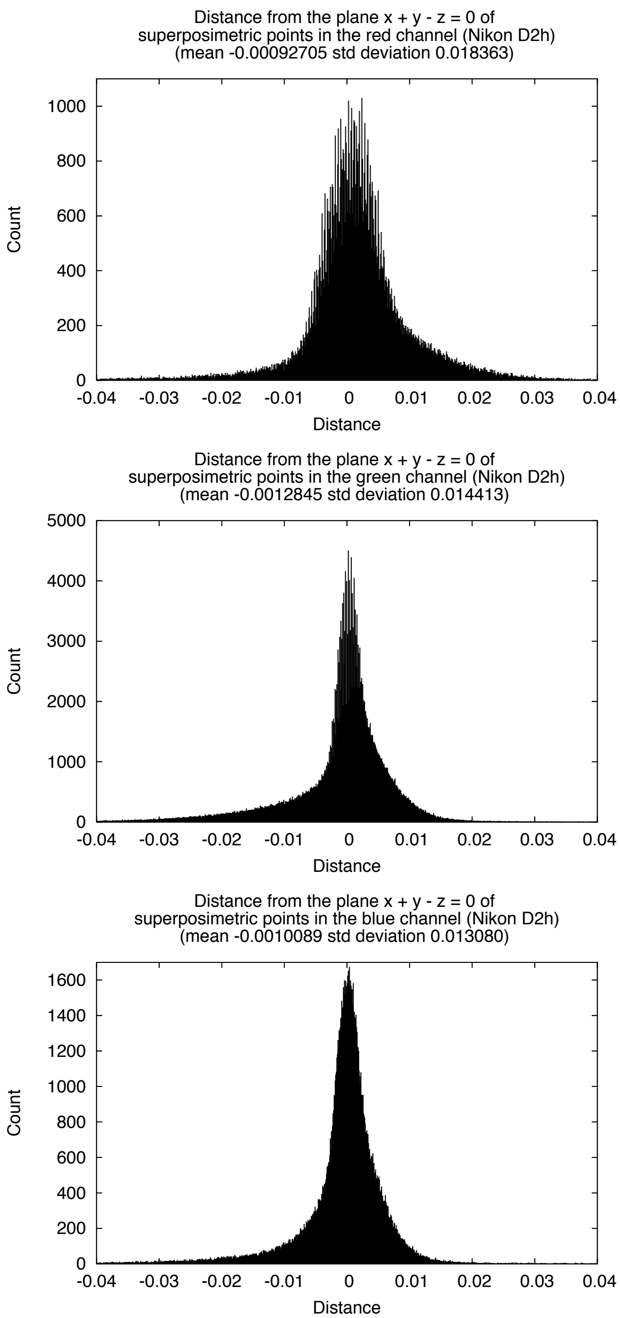

Fig. 5: Histograms of the distances from the plane of the red, green, and blue channels respectively.

the same subject matter. For each spatial location in the image, the variance was computed among each image at that spatial location. The mean was also computed and the variance among like means was averaged. There are difficulties in comparing the two results, partly because of the two tests being taken with different subject matter. Though ideally this should make no differences (the noise in the camera's sensor is no correlated to the subject matter), in practice there is a difference. The variance in low pixels (dark current noise) will generally be lower than noise in the midtones and highlights. Thus, different

	superposigram method (high)	multiple images(high)
Red	5.3×10^{-3}	5.0×10^{-3}
Green	4.7×10^{-3}	4.3×10^{-3}
Blue	3.9×10^{-3}	3.2×10^{-3}
	superposigram method (medium)	multiple images(medium)
Red	3.6×10^{-3}	3.2×10^{-3}
Green	3.1×10^{-3}	2.9×10^{-3}
Blue	2.1×10^{-3}	2.2×10^{-3}
	superposigram method (low)	multiple images(low)
Red	1.3×10^{-3}	1.0×10^{-3}
Green	9.5×10^{-4}	9.3×10^{-4}
Blue	8.6×10^{-4}	8.6×10^{-4}

TABLE I: Comparison of the standard deviations recovered using two methods. The left column shows the results using the method proposed in this paper. The right column shows results using multiple images of the same subject matter under identical illumination.

distributions of light an dark pixels will change the results. This can be overcome to some extent by looking at noise in specific regions (for example low, medium and high pixel values). Low pixels ere considered to be in the bottom third of the range, middle pixels in the central third, and high pixels in the top third.

The results from the two methods are shown in table I.

V. Measuring Noise in Various Cameras

Four digital SLR cameras were compared, two from Canon and two from Nikon. The first was a Nikon D2H camera which is a 4.1 megapixel camera which employs a proprietary JFET LBCAST sensor. The sensor's dimensions are 23.3mm $\times$ 15.5mm and is a three transistor RGB filter design. The lens used was a Nikon Nikkor 28mm lens at ISO 200 (the highest quality setting). The fstop setting used was 2.8 (the most open aperture available). Several exposures ranging from 4 to 250 milliseconds were used and combined into a single superposigram. The second camera used was a Nikon D70 digital SLR camera which is a 6 megapixel camera using a CCD sensor array. The dimensions of the sensor array are 23.7mm $\times$ 15.6mm. The lens used was the same lens used on the D2H, a 28mm Nikon Nikkor lens. Again, the ISO setting was 200, the highest quality setting. The first Canon camera used was a D Mark II digital SLR camera. This camera uses an 8 megapixel CMOS sensor array, the dimensions of which are 28.7mm $\times$ 19.1mm. The sensor uses a 4 transistor design as opposed to the Nikon 3 transistor design. The lens used was a 50mm Canon lens using an aperture of 1.4 (the most open aperture available) and an ISO setting of 100 (the highest quality ISO). Finally, the last camera tested was a Canon 10D. The sensor is a 6.3 megapixel CMOS sensor which measures 22.7mm $\times$ 15.1mm. The same lens (50mm Canon) was used as in the 1D Mark II testing. Again, the ISO was set to 100 (highest quality) and an aperture of 1.4 was used.

For each camera, a superposigram was constructed and distances to the plane of every superposimetric point were calculated. From the distances to the plane, a histogram was derived representing the noise distributions as well as the means and standard deviations of the distances. The standard

channel	Nikon D2H	Nikon D70	Canon 1D Mark II	Canon 10D
Red	5.3×10^{-3}	4.5×10^{-3}	1.8×10^{-3}	4.3×10^{-3}
Green	4.7×10^{-3}	2.4×10^{-3}	1.2×10^{-3}	2.2×10^{-3}
Blue	3.9×10^{-3}	3.0×10^{-3}	1.1×10^{-3}	2.4×10^{-3}

TABLE II: Standard deviations recovered from various commercially available cameras.

deviations from each of the cameras are reported in table II.

VI. Conclusion

Though other relatively simple methods exist for estimating the noise in a sensor array (for example identical exposures of the same image), the superposimetric method can find the noise distribution in a simple and elegant manner.

The method proposed is a simple and direct procedure for measuring the noise in a sensor array for a typical camera as opposed to various sensor- specific methods such as [8] and [9]. Though multiple exposures were used in composing the superposigrams used to analyze each sensor array this was simply utilized to improve the accuracy of the results and to ensure that the entire dynamic range of the camera had many samples. Probability distributions may be found for a given sensor array using as little as three images given the assumptions that the camera does not move and the sensor values for each image are well distributed and cover the dynamic range of the camera.

The method was proven to be accurate by means of using artificially generated images with known values and then adding white Gaussian noise of a known standard deviation and zero centered mean. We showed mathematically this shown be true using the superposimetric distances theorem, and verified the results using a well-known method of finding the sensors variance given a range of pixel values.

References

[1] C. Manders, C. Aimone, and S. Mann, "Camera response recovery from different illuminations of identical subject matter," in *Proceedings of the IEEE International Conference on Image Processing*, Singapore, Oct. 24-27 2004, pp. 2965–2968.

[2] Khurram Shafique and Mubarak Shah, "Estimation of the radiometric response functions of a color camera from differently illuminated images," in *Proceedings of the IEEE International Conference on Image Processing*, Singapore, Oct. 24-27 2004, pp. 2339–2968.

[3] P. E. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," *SIGGRAPH*, 1997.

[4] Charles Poynton, *A technical introduction to digital video*, John Wiley & Sons, 1996.

[5] C. Manders and S. Mann, "Determining camera response functions from comparagrams of images with their raw datafile counterparts," in *Proceedings of the IEEE first International Symposium on Intelligent Multimedia*, Hong Kong, Oct. 20-22 2004, pp. 418–421.

[6] S. Mann and R. W. Picard, "Virtual bellows: constructing high-quality images from video," in *Proceedings of the IEEE first international conference on image processing*, Austin, Texas, Nov. 13-16 1994, pp. 363–367.

[7] S. Mann, "Comparametric equations with practical applications in quantigraphic image processing," *IEEE Trans. Image Proc.*, vol. 9, no. 8, pp. 1389–1406, August 2000, ISSN 1057-7149.

[8] Abbas El Gamal Hui Tian, Boyd Fowler, "Analysis of temporal Noise in CMOS Photodiode Active Pixel Sensor," *Journal of Solid-State Circuits*, vol. 36, pp. 92–101, January 2001.

[9] D. Yang, B. Fowler, A. El Gamal, H. Min, M. Beiley, and K.Cham, "Test structures for characterization and comparative analysis of CMOS image sensors," *Proc. SPIE*, vol. 2950, pp. 8–17, October 1996.