Image noise estimation using color information

Carl Staelin, Hila Nachlieli
HP Laboratories Israel
HPL-2005-2(R.1)
September 4, 2007*

noise estimation, denoising, image processing

Imaging products are continuously increasing their optical resolution, resulting in higher perceived noise levels due to increased sensitivity to noise and graininess in the original image and to the decreased signal-to-noise ratio in the smaller pixel sensors. Denoising images requires accurate noise estimates; too low an estimate and too much noise remains in the image, too high an estimate and too many details are erased from the image. In addition, the noise level is often not constant over natural images, particularly images from grainy negatives. We present a novel noise estimation algorithm for scanned images, based on the statistical prior that the color channel gradients are highly correlated in natural noise-free images, extend it to provide local noise estimates, and then extend a denoising algorithm to utilize the spatially variant noise estimates for improved denoising while retaining details in the original image.

* Internal Accession Date Only

Approved for External Publication

Image noise estimation using color information

Carl Staelin and Hila Nachlieli

September 2, 2007

Abstract

Imaging products are continuously increasing their optical resolution, resulting in higher perceived noise levels due to increased sensitivity to noise and graininess in the original image and to the decreased signal-to-noise ratio in the smaller pixel sensors. Denoising images requires accurate noise estimates; too low an estimate and too much noise remains in the image, too high an estimate and too many details are erased from the image. In addition, the noise level is often not constant over natural images, particularly images from grainy negatives. We present a novel noise estimation algorithm for scanned images, based on the statistical prior that the color channel gradients are highly correlated in natural noise-free images, extend it to provide local noise estimates, and then extend a denoising algorithm to utilize the spatially variant noise estimates for improved denoising while retaining details in the original image.

1 Introduction

denoising algorithms usually assume a-priori knowledge of the noise level in the image to control the amount of denoising. In practice, this knowledge does not exist, and is usually replaced by manual tuning of the noise parameter. Some denoising algorithms, e.g. Donoho [1], tune the denoising parameter based on the input image statistics, but these are usually based on statistics of the luminance channel, and they can easily confuse high frequency texture with noise. Since scanners take three independent samples (usually R, G, and B) at each pixel, and since the color channel derivatives in natural, noise-free images tend to be highly correlated [2], we may utilize the redundant information to estimate the noise in the sampled images. High correlation in the color channel derivatives indicates low noise, while low correlation indicates high noise.

We extend this approach to mosaic images by first pre-processing the mosaic image to obtain a full-color image using simple bilinear interpolation of each color channel independently. We treat the two greens in the mosaic pattern as two separate colors, giving us a four color image

rather than the three color image for scanned inputs. In addition, since the actual distance between same-color pixels in the input mosaic image is two, we use derivatives with a distance of two pixels.

2 Prior work

Noise level estimation is typically the first step in image denoising, as the noise estimates are used to control the activity of the denoising algorithms. Most prior work, described below, concentrates on information obtained from high frequency variations in the luminance channel, whereas we propose using the correlation between color channels to evaluate the noise level. The basic idea behind most prior methods of noise estimation is that noise frequencies are typically higher than feature frequencies. Several of those approaches include analysis of the distribution of the local gradient amplitude [3], and analysis of the power spectrum of the image in order to estimate the variance of additive white noise [4, 5]. High frequencies are characteristic to noise, but also to texture, which increases the error of the noise estimator. One way of avoiding confusing texture with noise is to look only at smooth image regions [6, 7]. Similar approaches include estimating the noise using the standard deviation of unsharp mask values in pixels with low local gradients [8], or only in regions with the smallest gradient values [9]. Additional approaches to luminance based noise estimation, using wavelets or SVD, are described in [10, 11, 12].

3 Correlation versus scale

In natural, noise-free images, the derivatives of the color channels tend to be highly correlated [2], so in noisy images with color-channel independent noise, one would expect this correlation to decay as noise levels increase. We have attempted to verify this observation, both with scanned images with naturally occurring noise and with artificially noised images.

Fig. 1(a) contains the contour maps of the density of the color correlations for three natural images. These images were captured using a high quality digital camera with a 5Mpixel sensor. The images were then downscaled by a factor of four using pixel averaging to reduce the noise still further and to reduce the demosaicing effects on the image statistics. It is readily apparent that the color channel derivatives are highly correlated.

For scanned images with naturally occurring noise, we do not have an ideal, noise-free version of the image. However, by downscaling the image using pixel averaging we can see how the correlation changes as the assumed noise levels decrease. (Assuming the noise is independent, and each downscaled image is the result of the averaging of the NxN high resolution pixels, then the average noise per-pixel should decrease by a factor of N.)

The naturally noisy images were scanned by a Nikon Super Coolscan 5000 ED negative and slide scanner at 8bit and 4000dpi with no noise reduction option enabled. Color reconstruc-

(a) Color derivatives of noise-free natural images

image (a) image (b) image (c)

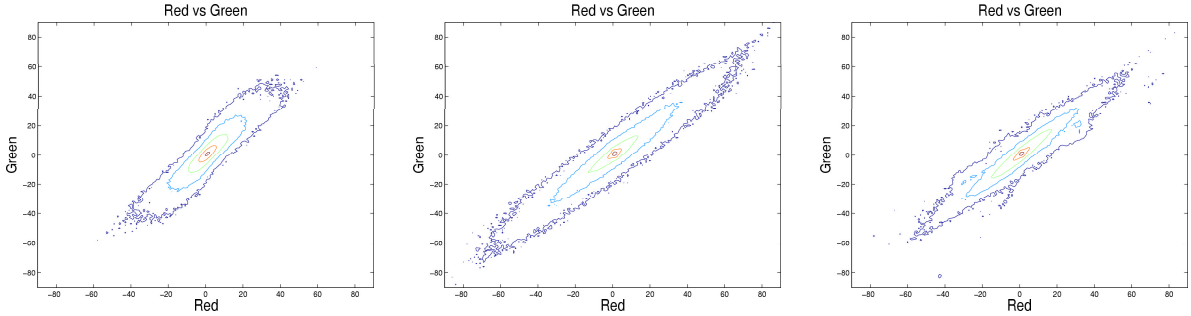

(b) Red-green color derivatives of a natural noisy image

original ×4 downscale ×16 downscale

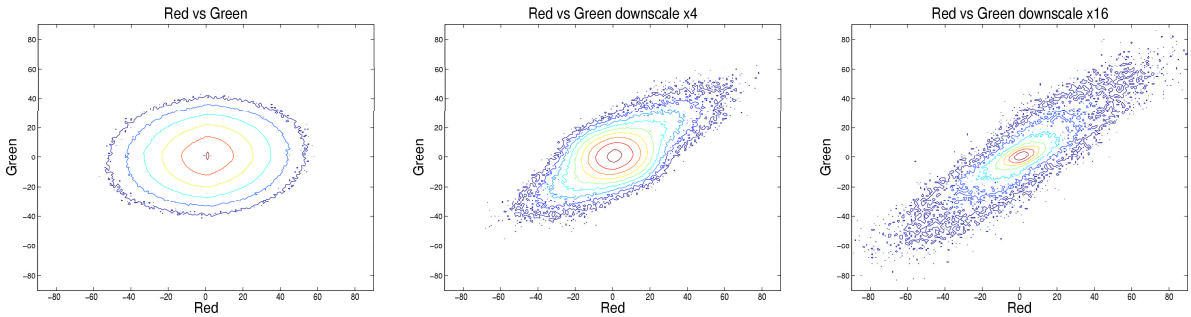

(c) Color derivatives of an image corrupted by ten gray levels of Gaussian noise

original ×4 downscale ×16downscale

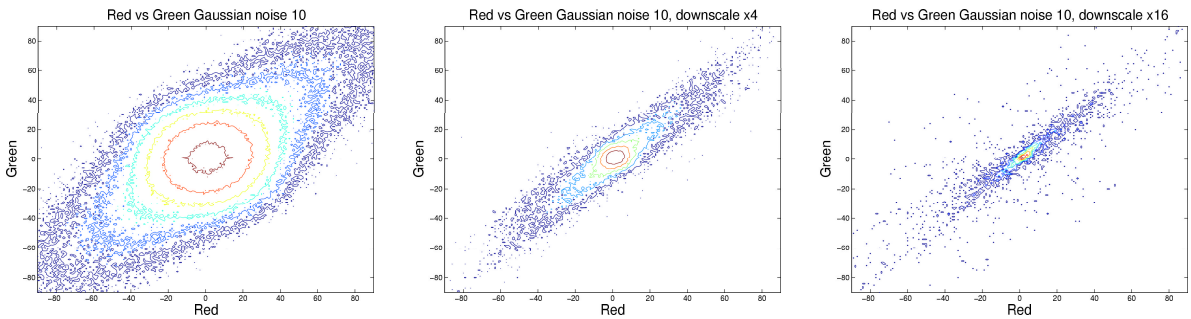

Figure 1: Correlation between color derivatives of images in various scales

tion and scratch reduction were enabled, and the scanned images were stored in a lossless TIFF format. The scanned images were about 20MPixels each, for roughly 60MB images. The negatives are about twenty years old using a standard Kodak film. Most of the noise in the image is actually grain in the negative, as we have tried a number of experiments where we scan the negative up to sixteen times each and use the average result, with little change in the apparent image noise. In addition, numerous works have shown that film grain may be modeled as luminance-dependent independent identically distributed (i.i.d.) Gaussian noise [13, 14, 15, 16].

Fig. 1(b) shows a sequence of contour maps of the density of the color channel derivatives for a natural scanned image. In Fig. 1(b), there appears to be little or no correlation between the color channel derivatives at the original resolution, while as we reduce the image size via pixel averaging, the correlation increases.

For an artificially noised image, if one looks at the color plane derivative correlations versus the downscale factor, one may see that the correlations increase as the scale (and noise) are reduced. Fig. 1(c) shows a sequence of contour maps of the density of the color channel derivatives. A diagonal line on a forty five degree slope would indicate perfect correlation, while a ball centered around zero indicates no correlation. In point of fact, the color channel derivatives in noise-free natural images are not perfectly correlated, else we would have a grayscale image rather than a color image.

As one may notice, the artificially noised images behave in a fashion similar to the naturally noised image, with much lower correlations at higher resolution, and increasing correlation as the noise is gradually reduced via pixel averaging.

4 A model for using color information for noise estimation

An overly simplified image model is that the color channel derivatives are perfectly correlated in natural images, and that noise is i.i.d. Gaussian. Under this model the color derivatives in 3-space would lie on the identity line. The experimental results from the previous section show that as noise is added to an image, the color derivatives tend to drift further from the identity line. Conversely, the average amount of drift away from the identity line should give us a measure for the amount of noise in the image. More formally, we can estimate the noise at a pixel as the distance between the deviation of the local derivatives $\delta\vec{u} = (\delta u_r, \delta u_g, \delta u_b)^T$ from the identity line $\vec{e} = (1,1,1)^T$ as:

$$D_{CCA} = ||\delta\vec{u} - < \delta\vec{u}, \hat{e} > \hat{e}|| \tag{1}$$

where $<>$ is the inner product and $\hat{e} = \frac{\vec{e}}{||\vec{e}||}$. We call this measure the *CCA distance*. Notice that a similar expression is obtained when calculating the (biased) standard deviation of the

three differences:

$$s = \sqrt{\frac{\sum_{i\in\{r,g,b\}} \left(\delta u_i - \overline{\delta u}\right)^2}{3}} \tag{2}$$

where $\overline{\delta u}$ is the mean of δu_i.

The per-pixel mean CCA distance is measured by the average of Equation 1 over the eight directional derivatives with the pixel's eight nearest neighbors. In some cases, discussed in the next section, we prefer the minimum of the pixel's eight nearest D_{CCA} values, instead of the average. We refer to this value as D^{min}_{CCA}. Also, while Equations 1 and 2 use the first derivative, we can also use the second derivative since if the first derivatives are correlated, then the second derivatives should also be correlated. The directional first derivative at x_i is $\partial x_i = x_i - x_{i-1}$, while the directional second derivative is $\partial\partial x_i = x_{i-1} - 2x_i + x_{i+1}$. We refer to D_{CCA} calculations using the second derivative as D_{CCA2} and D^{min}_{CCA2} respectively, depending on whether *mean* or *min* is used for the per-pixel values.

Let us denote a 3D volume element by dv, and by r the difference between two independent noisy pixels, both with an i.i.d. Gaussian noise with standard deviation σ. The distribution of r is Gaussian with standard deviation $\sigma_r = \sqrt{2}\sigma$. The expectation value of the per pixel mean CCA distances, $E\left\{D_{CCA}\left(\sigma_r\right)\right\}$, is then given by:

$$\begin{aligned} & E\left\{D_{CCA}\left(\sigma_r\right)\right\} \\ = & \int_V P_{r,\theta,\varphi}\left(\sigma_r\right)\left(r\sin\theta\right)dv \\ = & \frac{\frac{1}{4\pi}\frac{1}{\sigma_r\sqrt{2\pi}}\int_{r=0}^{\infty} r^3 e^{-\frac{r^2}{2\sigma_r^2}}dr\int_{\theta=0}^{\pi} sin^2\theta\, d\theta\int_{\varphi=0}^{2\pi} d\varphi}{\frac{1}{4\pi}\frac{1}{\sigma_r\sqrt{2\pi}}\int_{r=0}^{\infty} r^2 e^{-\frac{r^2}{2\sigma_r^2}}dr\int_{\theta=0}^{\pi} sin\theta\, d\theta\int_{\varphi=0}^{2\pi} d\varphi} \\ = & \sqrt{\frac{\pi}{2}}\sigma_r \sim \frac{\sigma}{0.564} \end{aligned} \tag{3}$$

D_{CCA} is the norm of a vector in 3-space (Equation 1), which has 3 components, one for each color channel. Given a three dimensional Gaussian i.i.d. noise,

$$\frac{1}{(2\pi)^{d/2}\sqrt{det\overline{\overline{Q}}}}\ \exp\left(-\frac{1}{2}\left(\delta\vec{u}-\vec{\mu}\right)^T\overline{\overline{Q}}^{-1}\left(\delta\vec{u}-\vec{\mu}\right)\right) \tag{4}$$

Where $\overline{\overline{Q}}$ is the covariance tensor:

$$\overline{\overline{Q}} = E\left\{\left(\delta\vec{u}-\vec{\mu}\right)\left(\delta\vec{u}-\vec{\mu}\right)^T\right\} \tag{5}$$

and $\vec{\mu} = (0,0,0)^T$ for zero mean noise. One can calculate the projections of the CCA distance on each of the color channels, denoted by D^r_{CCA}, D^g_{CCA}, D^b_{CCA}, for the red, green and blue colors, respectively, and extract the three eigenvalues of the covariance matrix.

In the isotropic case, substituting Equation 4 in the Expecation value in Equation 6 leads to:

$$\begin{aligned} E_x &= \int_V P_{r,\theta,\varphi}(\sigma_r) \left| \sqrt{\frac{2}{3}} r \sin(\theta - \theta_0) \cos(\varphi - \varphi_0) \right| dv \qquad (6) \\ &= \sqrt{\pi}\sigma \frac{4}{2\pi} \sqrt{\frac{2}{3}} = \sqrt{\frac{8}{3\pi}}\sigma \sim \frac{\sigma}{1.0854} \end{aligned}$$

with $x \in \{r, g, b\}$ for the red, green and blue axes respectively, where E_x is $E\{D^x_{CCA}(\sigma_r)\}$, dv is a 3D volume element, and $\left|\sqrt{\frac{2}{3}} r \sin(\theta - \theta_0) \cos(\varphi - \varphi_0)\right|$ is the value of the projection of the distance from the measured point $\delta\vec{u}$ to the unity line[1].

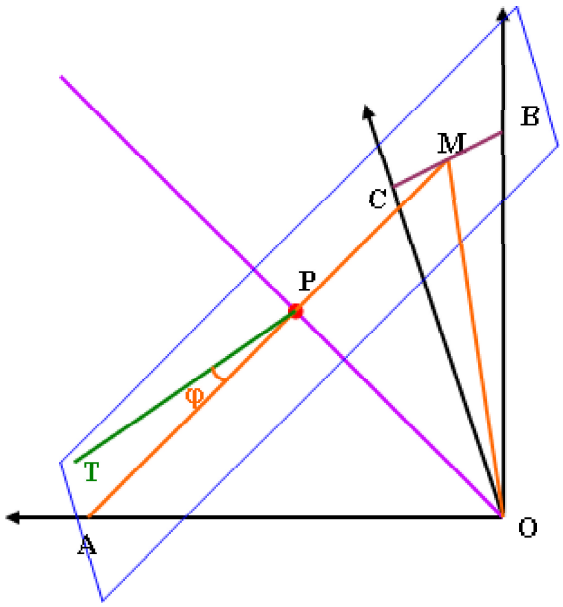

Figure 2: 3D angles

5 Results

We tested the use of the CCA distance for both global noise estimation and local noise estimation. The experimental results presented in this section are based on the data obtained by the following procedure: We used 145 high quality digital camera images, downscaled

[1]The projection is calculated as follows: in Fig. 2 the line OP is the three dimensional identity line. All lines which are orthogonal to the identity line lie on a surface, illustrated by the rectangle. This surface intersects the three axes at the points A,B and C. The line $PT = r \sin\theta$ is the distance from the measured point $T \equiv \delta\vec{u}$ to the identity line. We find the projection of PT on the OA axes by moving to the rotated ($\cos\theta_0 = 1/\sqrt{3}$, $\phi_0 = 0$) set of coordinates: $x' = PA = (2, -1, -1)^T/\sqrt{6}$, $y' = (0, 1, -1)^T/\sqrt{2}$ and $z' = OP = (1, 1, 1)^T/\sqrt{3}$: Angle POA equals acos($1/\sqrt{3}$), which is the inner product between $\hat{x}'$ and $\hat{x}$, and hence the cosine of angle PAO is $\sqrt{\frac{2}{3}}$, and the projection of PT on the axis OA is $\sqrt{\frac{2}{3}} r \sin(\theta - \theta_0) \cos(\varphi - \varphi_0)$.

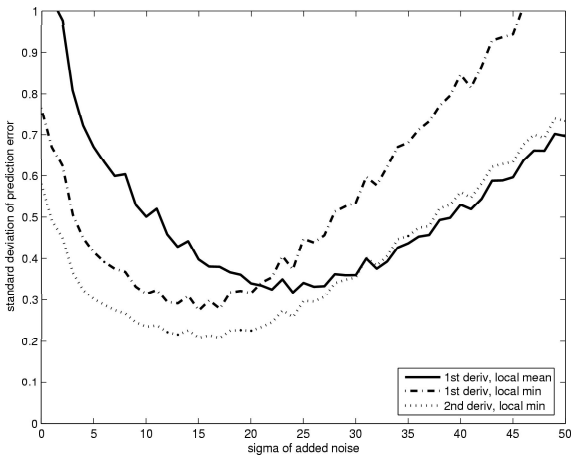

Figure 3: Image noise estimation errors

them by a factor of four to essentially eliminate noise, JPEG artifacts and de-mosaicing affects, added i.i.d. Gaussian noise with known standard deviation σ to each image, clip to $[0 \ldots 255]$, and calculated the per-pixel CCA distance for all pixels in the image.

Testing the use of the per-color CCA distance noise estimator (Equations 3, 6), averaging over 1 out of each 10 pixels in the image, we empirically derived the equations:

$$\widetilde{\sigma} = (0.570 \pm 0.008)\, D_{CCA} - (0.771 \pm 0.537) \tag{7}$$

and

$$\begin{aligned} \widetilde{\sigma}^r &= (1.099 \pm 0.015) D^r_{CCA} - 0.910 \pm 0.582 \\ \widetilde{\sigma}^g &= (1.094 \pm 0.011) D^g_{CCA} - 0.525 \pm 0.433 \\ \widetilde{\sigma}^b &= (1.099 \pm 0.013) D^b_{CCA} - 0.880 \pm 0.599 \end{aligned} \tag{8}$$

While the slope of the calculated fit agrees with the predicted slope, the entire line is shifted. This shift, common to all noise level, is probably hidden in the supposedly noise free images, which either contain a tiny amount of noise or contain colored edges which are interpreted as noise by the CCA distance. We return to this discussion later in this section.

In addition, each pixel has eight directional first derivatives and four directional second derivatives. We have at least two options for choosing the "representative" value of these per-pixel values: mean and min. Equation 9 shows the relevant polynomials used to convert the mean of the per-pixel CCA values from Equation 2 to the image noise estimate. Note the slope differs from Equation 7 due to the $\frac{1}{\sqrt{3}}$ in Equation 2.

$$\begin{aligned} \widetilde{\sigma} &= 1.00\, D_{CCA} - 0.94 \\ &= 2.53\, D^{min}_{CCA} - 0.86 \\ &= 0.96\, D^{min}_{CCA2} - 0.58 \end{aligned} \tag{9}$$

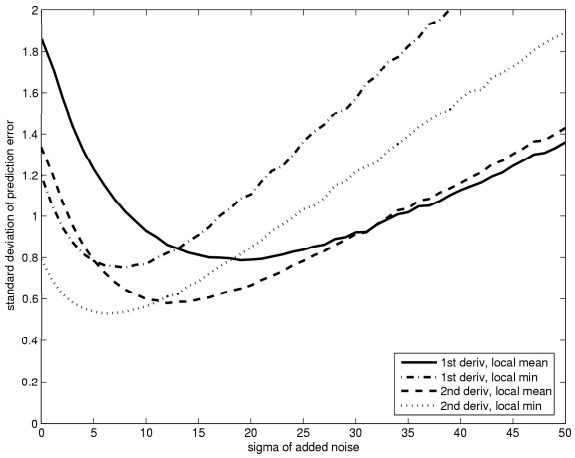

Figure 4: Spatially-variant noise estimation errors for 13x13 pixel neighborhoods

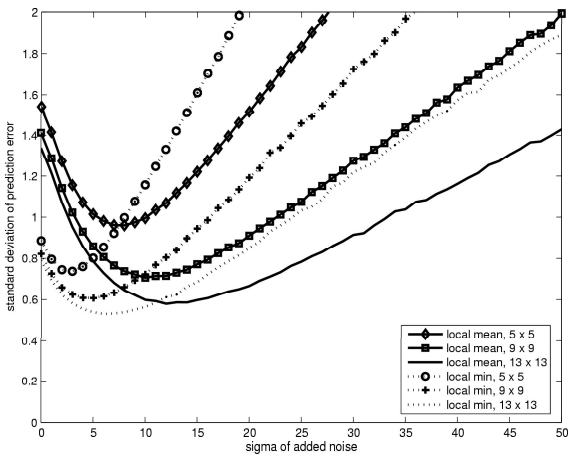

Figure 5: Spatially-variant noise estimation errors for 2^{nd} derivative

Fig. 3 shows the standard deviation of the prediction errors as a function of the injected σ. The prediction errors are computed using leave-one-out evaluation of the linear fitting process. In other words, when computing the prediction error for a test image, the curve fitting is done using the CCA distance measurements from the other images. It is clear that using the second derivative gives more accurate results, and that for smaller noise levels it is better to use the minimum rather than the mean of the various directional derivatives at each pixel.

Fig. 4 shows the standard deviation of the prediction errors for 13x13 windows for various configurations similar to Fig. 3, except that in this case the noise prediction and resulting error calculations are done per window rather than per image. As expected, the errors are generally higher than those of Fig. 3, but they are still small. However, using two prediction methods, local minimum for small noise levels and local mean for larger noise levels, appears to be more useful as the difference between them is much larger than in the per-image noise estimates.

Fig. 5 shows how the standard deviation of the prediction error shrinks as the neighborhood size increases. In practice, a 9x9 window should provide sufficient accuracy.

The reason for the larger variances at smaller noise levels is probably that natural images do generally have some color and hence some small, natural variation in the color derivatives. It is hard to tell whether the original images, which we assumed to be noise free, still contained some small level of noise, or if our noise estimation method confuses some of the colored texture and edges in the image with noise.

6 Localized Correlation between noise and color-noise estimation

A photograph captured on a chemical film may have different amounts of noise in different regions due the way the film reacts to different colors and light exposures. Fig. 6 is a 4000 dpi scan of a negative photograph captured on a chemical film (the printed images are, of course, down sampled). Fig. 7 is the CCA-based noise estimate according to Equation 1 convolved with an 11×11 averaging kernel. To make the result more visible, the grey levels are $15\times$ the actual noise level estimate. Especially near the top left of Fig. 7 there are some visible thin dark lines. These are a result of the dust and scratch reduction algorithm used in the scanner driver which identified some mold on the negative and filled in those regions, which correspond to the dark lines.

For comparison purposes, we also show in Fig. 8 the sigma map for a "noise-free" version of the same image, shown in Fig. 9, obtained by denoising and then downscaling by a factor of eight in each dimension by pixel averaging. This image supports the assumption stated above that our noise estimator may mistake color edges for noise. One way to overcome this

Figure 6: Sample scanned image

Figure 7: Sample noise estimate

Figure 8: Noise estimate for clean image

Figure 9: Clean image

Figure 10: Sample noise to signal estimate

Figure 11: Noise to signal estimate for clean image

issue is to estimate the noise to signal ratio (Equation 10) instead of the noise:

$$N2S_{CCA} = \frac{D_{CCA}}{\max\left(\frac{1}{N_{channels}}, \overline{\delta u}\right)} \tag{10}$$

Fig. 6 shows noise to signal ratio values in the noisy image. The noise to signal ratio in the sky is larger than in the skin tone regions, and edges are filtered out. Fig. 11 shows the noise to signal in the clean image, multiplied by 32 in order to make the result more visible. The residual noise to signal in the sky is bigger than the residual noise to signal on the skin, and the edges are not visible. Research on using the information obtained from the local derivatives ($\overline{\delta u}$) to improve per-image noise estimation is on-going.

7 Noise estimation in mosaic images

We extended the proposed noise estimation for mosaic images by calculating the second derivatives of each of the colors in each of the four directions. The distance between neighboring pixels in each of the directions and colors is doubled, relative to the distance between neighboring pixels in the same direction and color in the scanned image, i.e., $\partial\partial x_i = x_{i-2} - 2x_i + x_{i+2}$. We treat the two greens in the mosaic pattern as two separate colors, giving four colors rather than three. As in the scanned images, we calculate the CCA distance in each direction and choose the direction that gives the minimal D_{CCA2} value, denoted by D^{min}_{CCA2}. Fig. 12 shows the per-image noise estimation error in mosaic images, which where created by mosaicing our 145 test images. As one can see, the best method

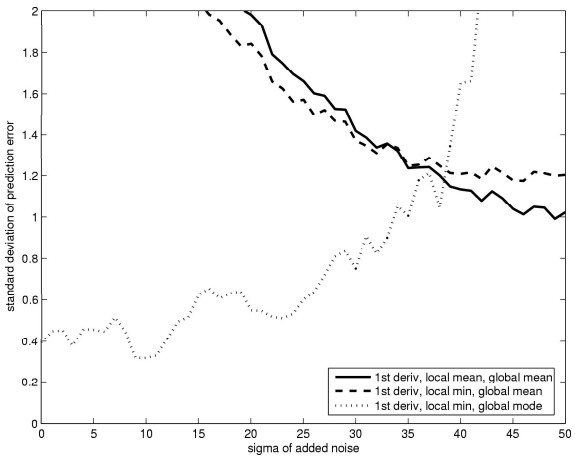

Figure 12: Estimation error for per-image noise estimation in mosaic images

for calculating the per-image noise estimation out of the per-pixel D^{min}_{CCA2} is by using the a variant of the mode or most common value: the D^{min}_{CCA2} histogram was prepared using 0.25 spaced bins, the most common value (mode) was found, and the median of the values of D^{min}_{CCA2} whose density was at least 95% of the density of the mode was calculated. The reason for using the median of the values whose density is at least 95% of the density of the mode is that the mode for large noise levels is somewhat unstable due to random fluctuations in the observed density, whereas the median of the densest portions of the histogram is much more stable.

8 Application to spatially variant denoising

We previously developed [17] a bi-selective filter based on over complete DCT domain processing which can selectively smooth and sharpen in a single pass. The heart of the algorithm is a function that transforms the DCT coefficients: small coefficients are squashed, large coefficients are (slightly) enhanced, and there is a smooth transition between the two regimes. This function is governed by three parameters: L, a lower threshold below which all coefficients are set to zero, S, a sharpening factor, and H, an upper threshold above which the coefficient is increased by the sharpening factor. Between the lower and upper thresholds the function uses linear interpolation. In our previous work [17], the three parameters were set globally for the whole image, usually setting L to about two times the global noise estimate, $S = L$, and $H = 3 \cdot L$ respectively generally gives good results.

Given the spatially variant noise estimates, however, we are able to extend this algorithm to adaptively filter the image depending on the local noise estimates. To do this we take the noise estimate at the central pixel, σ_{ij} in the DCT block and scale the three transform coefficients by the local noise estimate. Experimentally we found that setting $L = 2.5 \cdot \sigma_{ij}$ and S and H as before generally gives good results.

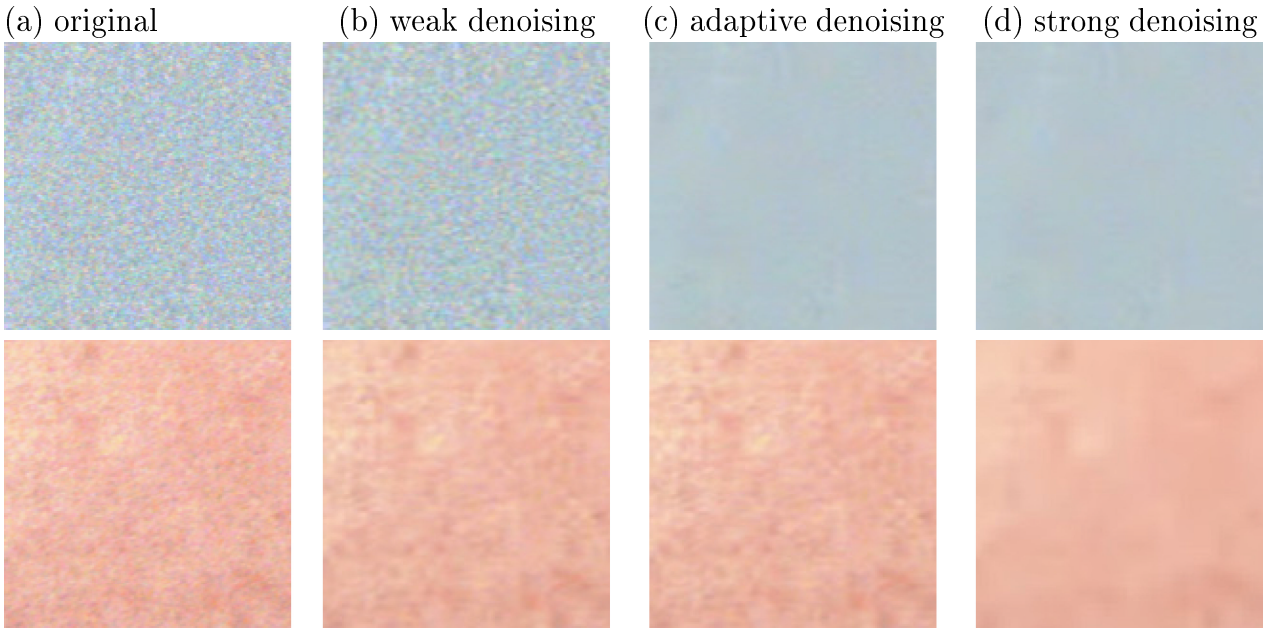

Figure 13: Sample image and noise estimate

Fig. 13 demonstrates the value of spatially variant, or adaptive filtering based on the spatially variant noise estimates. It shows two fragments of the image in Fig. 6, one from the sky just above the girl's head (top), and one from her forehead (down). Fig. 13(a) contains the original fragments, while Fig. 13(b) was processed with uniformly weak denoising to preserve the texture. Fig. 13(c) were processed adaptively using the noise estimates from Fig. 7, while Fig. 13(d) was processed with uniformly strong denoising to clear the noisy sky. It is apparent that the adaptive filtering is able to preserve the texture in the relatively clean and textured face while clearing the noise in the relatively noisy and flat sky.

9 Conclusions

We have demonstrated a simple method for estimating the noise in color images, assuming that each color channel is fully sampled and that the noise is Gaussian and independent. Perhaps surprisingly, using artificial data this method can estimate noise levels within a single gray value across a whole image for many noise levels. In addition, a spatially-variant form of the measure is able to estimate noise levels within a few gray levels for many noise levels. We demonstrated a sample application of the spatially variant noise estimate by extending an existing denoising method to control the amount of denoising at each part of the image by the estimated amount of noise in the image, and demonstrated that it can potentially perform much better than a denoising method with a single image-wide parameter setting. We then extended the noise estimation method to mosaic images, which have only one of R, G, or B samples at each pixel, and using a modified algorithm were able to determine the

noise level within a single gray value for most noise levels.

Acknowledgment

We gratefully acknowledge the advice and feedback of Yaacov Hel-Or, Doron Shaked, Renato Keshet, and Mani Fisher.

10 References

[1] D. L. Donoho, "De-noising by soft-thresholding," *IEEE Transactions on Information Theory*, vol. 41, no. 3, pp. 613–627, May 1995.

[2] Y. Hel-Or, "The canonical correlations of color images and their use for demosaicing," Hewlett-Packard Laboratories, Tech. Rep. HPL-2003-164(R.1), February 23 2004. [Online]. Available: http://lib.hpl.hp.com/techpubs/2003/HPL-2003-164R1.pdf

[3] V. Kayargadde and J. Martens, "An objective measure for perceived noise," *SIGNAL PROCESSING*, vol. 49, no. 3, pp. 187–206, 1996.

[4] N. Nill and B. Bouzas, "Objective image quality measure derived from digital image power spectra," *OPTICAL ENGINEERING*, vol. 31, no. 4, pp. 813–825, 1992.

[5] E. P. Simoncelli and E. H. Adelson, "Noise removal via Bayesian wavelet coring," in *Proceedings 3rd IEEE International Conference on Image Processing*, vol. 1, 1996, pp. 379–382.

[6] D. Shaked and H. Nachlieli, "System and method for estimating image noise," USPTO patent application 20050244075, patent pending, November 3 2005.

[7] F. Murtagh and J.-L. Starck, "Image processing through multiscale analysis and measurement noise modeling," *Journal Statistics and Computing*, vol. 10, no. 2, pp. 95–103, 2000.

[8] M. P. Keyes and K. A. Hoff, "Sharpening system adjusted for measured noise of photofinishing images," US patent no. 6,118,906, September 2000.

[9] J. Katajamaki and H. Saarelma, "Objective quality potential measures of natural color images," *Journal of Imaging Science and Technology*, vol. 42, no. 3, pp. 250–263, 1998.

[10] S. Olsen, "Noise variance estimation in images," 1993. [Online]. Available: citeseer.ist.psu.edu/olsen93noise.html

[11] A. B. M. Jansen, "Multiple wavelet threshold estimation by generalized crossvalidation for images with correlated noise," *Image Processing, IEEE Transactions on Publication*, vol. 8, no. 7, pp. 947–953, July 1999.

[12] G. S. Y. Konstantinos Konstantinides, Balas K. Natarajan, "Noise estimation and filtering using block-based singular value decomposition," *IEEE Transactions on Image Processing*, vol. 6, no. 3, pp. 479–483, 1997.

[13] O. K. Al-SHaykh and R. M. Mersereau, "Restoration of lossy compressed noisy images," *IEEE Transactions on Image Processing*, vol. 8, no. 10, pp. 1348–1360, October 1999.

[14] J. C. K. Yan and D. Hatzinakos, "Signal-dependent film grain noise removal and generation based on higher-order statistics," in *Proceedings of the IEEE Signal Processing Workshop on Higher-Order Statistics*, Banff, Alberta, Canada, July 21–23 1997, pp. 77–81.

[15] D. G. Falconer, "Image enhancement and film-grain noise," *Optica Acta*, vol. 17, no. 9, pp. 693–705, 1970.

[16] H. H. Arsenault, C. Gendron, and M. Denis, "Transformation of film-grain noise into signal-independent additive Gaussian noise," *Journal of the Optical Society of America*, vol. 71, no. 1, pp. 91–94, January 1981.

[17] H. Nachlieli, C. Staelin, M. Fischer, D. Shaked, R. Keshet, and P. Kisilev, "Transform domain robust filtering," Hewlett-Packard Laboratories, Tech. Rep. HPL-2005-185, October 2005, http://lib.hpl.hp.com/techpubs/2005/HPL-2005-185.pdf.