

A Method for Modeling Noise in Medical Images

Pierre Gravel*, Gilles Beaudoin, and Jacques A. De Guise

Abstract—We have developed a method to study the statistical properties of the noise found in various medical images. The method is specifically designed for types of noise with uncorrelated fluctuations. Such signal fluctuations generally originate in the physical processes of imaging rather than in the tissue textures. Various types of noise (e.g., photon, electronics, and quantization) often contribute to degrade medical images; the overall noise is generally assumed to be additive with a zero-mean, constant-variance Gaussian distribution. However, statistical analysis suggests that the noise variance could be better modeled by a nonlinear function of the image intensity depending on external parameters related to the image acquisition protocol. We present a method to extract the relationship between an image intensity and the noise variance and to evaluate the corresponding parameters. The method was applied successfully to magnetic resonance images with different acquisition sequences and to several types of X-ray images.

Index Terms—Image processing, magnetic resonance imaging, noise measurement, robustness, X-rays.

I. INTRODUCTION

IMAGE noise is a common problem in most image processing applications as evident in the extensive literature on the ways to reduce or circumvent it. Beginners in image processing soon discover that a denoising step is often required before any relevant information can be extracted from an image. The plethora of denoising functions included in popular commercial software is also evidence of the importance of extracting a signal from a noisy image. We adopt a different approach by using noise statistics to extract new information from standard images. Our method relies on the measurement of the relationship between the image intensity I and the noise variance σ_N^2 . This relationship is of the form

$$\sigma_N^2 = f(I, \alpha_1, \alpha_2, \alpha_3, \dots) \quad (1)$$

and depends on the noise model whereas the values of the parameters $\alpha_1, \alpha_2, \alpha_3, \dots$ are determined by the image acquisition protocol. The relationship (1) should apply to every pixel in the

image. The method is also designed to study uncorrelated noise, i.e., for which the correlation length is zero.

Several methods for estimating the variance of white additive noise in images have been proposed [1] but they cannot be used on images where the relationship (1) applies since the noise variance on them is nonuniform. The method presented here is more general and measures the noise variance in smooth, preferentially homogeneous regions of an image. This is particularly useful when only one image is available for analysis. When the same object can be imaged several times, the local mean intensity and noise variance can be estimated for each pixel [2], [3]. This provides a large number of (I, σ_N^2) measurements from which the parameters in (1) can be determined. However, as it is often the case in practice, only a single image might be available and the local mean intensity must be estimated using a smoothed version of the image. As a result, the local mean intensity and noise variance are less precise and must be combined with measurements at other pixels to increase accuracy. This is the approach we elected to use.

A common misconception in image processing is to assume noise to be additive with a zero-mean, constant-variance Gaussian distribution or to be Poisson distributed. This assumption simplifies image filtering and deblurring, but the poor quality of the results generally indicates that a better understanding of the noise properties is required. Hence, the noise on magnetic resonance (MR) images was found to have a Rician probability density function (PDF) instead of a Gaussian one [4] whereas the noise on computed tomography (CT) images was found to be Gaussian instead of Poisson distributed [3], [5]. The modeling of image noise is not new and has led to investigations in various fields. Several authors have characterized the noise on aerial images [6]–[9] because texture analysis (based on image variance) is a fundamental tool of remote sensing for terrain classification. Similar techniques were also used in medical imaging for tissue classification and segmentation [10], [11]. In a different type of application, the modeling of noise during MR image acquisition and Shannon's theory of information content were used to derive an optimum in the trade-off problem between image resolution and contrast-to-noise ratio [12]. Given the relevance of noise modeling in the previous applications, we present and validate a method to do so which is more robust and reliable than previously published ones.

In Section II, we describe the statistical properties of several types of noise found in medical images. The method to extract the relationship (1) from a given image as well as the various imaging protocols used in this work are described in Section III. The method is then applied to several types of MR and X-ray images with the results being analyzed in Section IV. The conclusion follows in Section V.

Manuscript received February 2, 2004; revised May 24, 2004. This work was supported by Valorization-Recherche Québec, Canada. The Associate Editor responsible for coordinating the review of this paper and recommending its publication was J. Liang. Asterisk indicates corresponding author.

*P. Gravel is with the Laboratoire de recherche en imagerie et orthopédie (LIO) and the Département du génie de la production automatisée, École de technologie supérieure (ÉTS), Montréal, QC H3C 1K3, Canada (e-mail: pierre.gravel@etsmtl.ca).

G. Beaudoin is with the Centre hospitalier de l'Université de Montréal (CHUM), Montréal, QC H2L 4M1, Canada (e-mail: gilles.beaudoin@att-global).

J. A. de Guise is with the Laboratoire de recherche en imagerie et orthopédie (LIO) and the Département du génie de la production automatisée, École de Technologie Supérieure (ÉTS) (e-mail: jacques.deguise@etsmtl.ca).

Digital Object Identifier 10.1109/TMI.2004.832656

II. NOISE MODELS

In this section, we discuss the statistical properties of three common types of noise found in medical imaging (Gaussian, Poisson, and Rician) and derive the relationship (1) for each of them. Whenever an image contains different types of uncorrelated noise, the overall noise variance σ_N^2 can be expressed by summing up the various noise contributions

$$\sigma_N^2 = \sigma_1^2 + \sigma_2^2 + \dots + \sigma_n^2 \quad (2)$$

For example, the images from a charge—coupled device (CCD) camera are free of grain noise but are degraded by Poisson and read-out noises [7]. The noise variance on such images is either constant or linear dependent on the signal intensity. Unless otherwise mentioned, the images analyzed in this work were such that the noise contributions from secondary sources were negligible.

A. Gaussian Noise

This most common type of noise results from the contributions of many independent signals. This is a consequence of the central limit theorem which states that the sum of many random variables with various PDFs results in a signal with a Gaussian PDF. For example, the reading noise from a CCD detector is generated by the thermal fluctuations in many interconnected electronics components and, thus, has a Gaussian PDF. Gaussian noise is such that σ_N^2 is a constant.

B. Poisson Noise

Poisson noise prevails in situations where an image is created by the accumulation of photons over a detector. Typical examples are found in standard X-ray films, CCD cameras, and infrared photometers. We focus our attention on images saved with linear or logarithmic intensity scalings.

1) *Linear Intensity Scaling*: The following analysis assumes a pixel intensity corresponding to the number of monochromatic photons captured in a given amount of time. Real X-ray beams are not monochromatic and have energy spectra showing strong characteristic emission lines superimposed over a Bremsstrahlung radiation background [13]. The energy deposited at a pixel location (the image intensity) does not correspond exactly to the number of captured monochromatic photons since the X-rays follow a compound Poisson noise process. Whereas the number of X-rays follows a Poisson noise distribution, the X-ray energy converted counts follow a compound Poisson noise distribution due to the wide spectrum of the energy. Our analysis based on monochromatic photons remains nevertheless a good approximation of reality as will be shown in Section IV.

For a Poisson process of mean X (where X represents the number of captured photons), the expectation value for the variance is

$$\sigma_X^2 = X. \quad (3)$$

Because a recorded image is usually linearly rescaled to accommodate a given range in gray scales, the relation between

the intensity X reaching the detector and the recorded intensity I is

$$I = \gamma X + \delta \quad (4)$$

where γ and δ are constants. The corresponding variance σ_I^2 varies linearly with the intensity I

$$\sigma_I^2 = \gamma I - \delta\gamma. \quad (5)$$

2) *Logarithmic Intensity Scaling*: The noise statistics are altered on photographic plates, such as those used for X-rays, where the stored information is the optical density. An X-ray film being a *negative* recorder, an increase in light exposure causes the developed film to become darker. The degree of darkness of the film is quantified by the optical density D which is measured in this work by a scanner. The relation between the intensity X and the optical density depends on the film, the developing process used and the scanner. Within the operational intensity range of a given film, the relation can be expressed as

$$D = D_m - \log\left(\frac{X}{X_M}\right) \quad (6)$$

where X/X_M is the transmittance of the film [13] and corresponds to the fraction of the light that reaches the scanner detector (X_M is the maximum intensity). The minimum optical density D_m (or maximum transmittance) is found in the bright bone regions and the maximum optical density (or minimum transmittance) is found in the dark background regions. Taking into account image rescaling, the relation between the optical density and the scanned intensity, I , becomes

$$I = \gamma D + \delta \quad (7)$$

and the corresponding variance σ_I^2 varies exponentially with the intensity I

$$\ln(\sigma_I^2) = \left[\ln\left(\frac{\alpha^2 \gamma^2}{X_M}\right) - \frac{I_m}{\alpha\gamma} \right] + \frac{I}{\alpha\gamma} \quad (8)$$

with

$$\alpha = \ln_{10}(e). \quad (9)$$

The scanned image is inverted such that a large value of I corresponds to a large value of optical density D . Both α and γ constants are, thus, positive and the optical density variance increases with the optical density.

The analysis assumed again monochromatic photons but (8) still remains an excellent approximation of reality as will be shown in Section IV.

3) *Rician Noise*: The noise in MR images has a Rician PDF [4]. For these tests, we have used a standard volume coil (bird cage) which has uniform efficiency throughout the volume of interest. The signals are acquired in quadrature. Each signal produces an image X that is degraded by a zero-mean Gaussian noise of standard deviation σ_0 (which we define as the noise level). The two images are then combined into a magnitude

image I and the Gaussian noise PDF is transformed into a Rician noise PDF. The expectation values for the mean magnitude and the variance are [14]

$$I = \sigma_0 \sqrt{\frac{\pi}{2}} \exp\left(-\frac{X^2}{4\sigma_0^2}\right) \cdot \left[\left(1 + \frac{X^2}{2\sigma_0^2}\right) I_0\left(\frac{X^2}{4\sigma_0^2}\right) + \frac{X^2}{2\sigma_0^2} I_1\left(\frac{X^2}{4\sigma_0^2}\right) \right], \quad (10)$$

$$\sigma_I^2 = X^2 + 2\sigma_0^2 - \frac{\pi\sigma_0^2}{2} \exp\left(-\frac{X^2}{4\sigma_0^2}\right) \cdot \left[\left(1 + \frac{X^2}{2\sigma_0^2}\right) I_0\left(\frac{X^2}{4\sigma_0^2}\right) + \frac{X^2}{2\sigma_0^2} I_1\left(\frac{X^2}{4\sigma_0^2}\right) \right]^2 \quad (11)$$

where I_0 and I_1 are modified Bessel functions of the first kind. There was no MR image rescaling found in the data we analyzed and, thus, no scaling relation like (4) and (7) was used here.

When $X \gg \sigma_0$, the Rician PDF approaches a Gaussian PDF [14] with $I \cong X$ and $\sigma_I \cong \sigma_0$. When $X \ll \sigma_0$, the Rician PDF approaches the Rayleigh PDF which does not depend on X [14]. The expectation values for the mean magnitude and the variance of a Rayleigh PDF are

$$I = \sigma_0 \sqrt{\frac{\pi}{2}} \quad (12)$$

$$\sigma_I^2 = \frac{4 - \pi}{2} \sigma_0^2. \quad (13)$$

Equations (12) and (13) can be used to estimate the value of σ_0^2 in an image from measurements in background regions where the magnitude is almost zero.

Before leaving this section, a last comment must be made about the optimum in the trade-off problem between MR image resolution and contrast-to-noise ratio [12]. The derivation assumed the noise to be gaussian and additive in order to estimate Shannon's information content of an image based on its power spectrum. These two assumptions are valid when the noise level is small ($X \gg \sigma_0$). However, this is not true for very noisy images ($X \ll \sigma_0$) for which the trade-off problem becomes particularly crucial. Special care must be taken in those cases.

III. MATERIALS AND METHODS

The first part of this section describes a method for measuring the relationship (1) between the intensity and the noise variance in an image. A noise image is first generated as the difference between the original image and a smoothed version of it. A mask image is then created to identify the pixels on the image plateaus. The noise variance on these image plateaus is then evaluated using robust estimators. The second part describes the imaging protocol for the acquisition of MR and X-ray images.

A. Noise Characterization

1) *Image Smoothing*: To a first level of approximation, each noise type studied can be modeled as a zero-mean Gaussian noise with a locally varying standard deviation. We assume that the difference between an image and a smoothed version of it

produces a zero-mean noise signal with a standard deviation of $\sigma_N \approx \sigma_I$. A noise image N is obtained by subtracting a smoothed version $\bar{I}$ from the original image I

$$N_{i,j} = I_{i,j} - \bar{I}_{i,j} \quad (14)$$

where i and j refer to pixel row and column indexes.

The image smoothing is performed by convolving the original image with a boxcar low-pass filter of size D .

The smoothing method works well for pixels located on plateaus where the intensity gradients are small. Near the edges, where the intensity gradients are large, the image smoothing does not reproduce the local mean intensity well and the noise signal has a nonzero mean.

The filter size D depends on image resolution and is found by trial and error. If D is too small, the smoothed image tends to follow the original image too closely and the noise variance is underestimated. If D is too large, the intrinsic variations in the image are smoothed out and the noise variance is overestimated or may not be of the form (1). All the images in this study were processed with $D = 5$ pixels and the limited size of the filter was taken into account by multiplying the noise variance with a correction factor (see Section III).

2) *Binary Mask Generation*: Edge pixels are discarded in the analysis and are masked out using a binary mask based on the edge information. The mask is first created by applying a threshold τ to a gradient image computed using Sobel filters. The threshold value is found using a method described in Section IV.

The binary mask is eroded by half the size of the smoothing filter to remove the pixels where the computed local mean intensity is imprecise due to the proximity to the image edges. The image boundaries are also eroded to the same depth to remove boundary effects due to filtering. Finally, the binary mask is cleaned from binary noise using standard morphological operators (opening and closing). This method eliminates highly textured regions from the noise analysis.

3) *Noise Variance Estimation*: To compute the noise variance σ_N^2 for a given mean intensity $\bar{I}$, we first need to identify all the unmasked pixels in the smoothed image that have this intensity. A noise sample (14) is available for each such pixel and the noise variance can be computed from them. We use an equidistant-bin histogram of the local mean intensity $\bar{I}$ of the unmasked pixels; the bin width being small enough that σ_N can be considered constant within each of them. For each bin k , we compute the mean intensity $\bar{I}_k$, the noise standard deviation $\sigma_{N,k}$, and its associated error $\sigma_{NN,k}$. In what follows we will use the bold-face notation to indicate vector quantities.

The noise distribution changes from bin to bin and the Gaussian estimator of the variance cannot be used in general. We use instead a robust estimator of the noise standard deviation in the k th bin that is based on the median absolute deviation [2]

$$\hat{\sigma}_{N,k} = 1.4826 \text{ med} \left| \mathbf{N}^k - \text{med}(\mathbf{N}^k) \right| \quad (15)$$

where the $\mathbf{N}^k$ is the vector of noise samples in the bin. The error associated to the last estimator was found using Monte Carlo simulations

$$\hat{\sigma}_{\mathbf{NN}} \approx 1.14 \frac{\hat{\sigma}_{\mathbf{N}}}{\sqrt{\mathbf{n}}} \quad (16)$$

where $\mathbf{n} \gg 1$ is a vector that holds the number of noise samples in each bin.

The last two results must be corrected for the limited size of the boxcar low-pass filter

$$\hat{\sigma}_{\mathbf{N}} = \sqrt{\frac{n_p}{n_p - 1}} \hat{\sigma}_{\mathbf{N}} \quad (17)$$

$$\hat{\sigma}_{\mathbf{NN}} = \sqrt{\frac{n_p}{n_p - 1}} \hat{\sigma}_{\mathbf{NN}} \quad (18)$$

where n_p is the number of pixels in the filter kernel.

4) *Data Fitting*: The intensity histogram has L bins. Using the L values of the mean intensity $\bar{I}$, the noise standard deviation $\sigma_{\mathbf{N}}$, and its associated error $\sigma_{\mathbf{NN}}$, we can now evaluate the M free parameters $\alpha_1, \dots, \alpha_M$ in the noise model (1) by minimizing the weighted mean square error (WMSE) defined as

$$\Omega(\tau) = \frac{1}{L - M} \sum_{k=1}^L w_k \left(\frac{f(\bar{I}_k, \boldsymbol{\alpha}) - \sigma_{\mathbf{N},k}}{\sigma_{\mathbf{NN},k}} \right)^2. \quad (19)$$

where k is the bin index, w_k is the corresponding weight, $\boldsymbol{\alpha}$ is the vector of parameters, and τ is the threshold value used in the binary mask generation. We use the Nelder-Mead simplex (direct search) method [15] with an adjustment of the weights.

Outliers have a large influence on a least squares fitting because squaring the residuals magnifies the effects of these extreme data points. To minimize this problem, we fit the data using a robust regression scheme based on Tukey's bisquare weights [16] where the weight given to each data point depends on how far the point is from the fitted curve. Points near the curve get full weight. Points farther from the curve get reduced weight. Points that are farther from the curve than would be expected by random chance get zero weight [17], [18]. The weights are computed as follows.

1) Set $\mathbf{w} = 1$.

2) Compute the residuals $\mathbf{r}$ from the weighted least squares fit that minimizes the WMSE (19)

$$\mathbf{r} = \frac{f(\bar{I}, \boldsymbol{\alpha}) - \sigma_{\mathbf{N}}}{\sigma_{\mathbf{NN}}}. \quad (20)$$

3) Compute the standard deviation of those residuals using a robust estimator based on the median absolute deviation

$$\sigma_{\text{MAD}} = 1.4826 \text{med}|\mathbf{r} - \text{med}(\mathbf{r})|. \quad (21)$$

4) Compute the normalized residuals

$$\mathbf{u} = \frac{\mathbf{r}}{6\sigma_{\text{MAD}}}. \quad (22)$$

5) Compute the bisquare weights $\mathbf{w}$

$$\mathbf{w} = \begin{cases} (1 - \mathbf{u}^2)^2, & |\mathbf{u}| < 1 \\ 0, & |\mathbf{u}| \geq 1 \end{cases}. \quad (23)$$

6) Go to Step 2 and continue the iterations until convergence is obtained.

The measured values of the mean intensity $\bar{I}$, the noise standard deviation $\sigma_{\mathbf{N}}$, and its associated error $\sigma_{\mathbf{NN}}$, depend implicitly on the threshold value τ (which corresponds to an intensity gradient). Too small values of τ result in a limited number of intensity measurements and too large values produce distortions in the shape of (1). The optimum threshold value is found as follows. For a given value of τ , we generate a binary mask, compute $\bar{I}$, $\sigma_{\mathbf{N}}$, and $\sigma_{\mathbf{NN}}$ from the unmasked pixels, and find the values of the free parameters $\alpha_1, \dots, \alpha_M$ that minimize $\Omega(\tau)$. We repeat the same series of steps for different threshold values and then select the value of τ that produces the smallest value of $\Omega(\tau)$.

The parameter-estimation procedure is slightly different for MR images with Rician noise since (10) and (11) are parametric equations. The second step (20) requires using a look-up table of (I, σ) values generated using (10) and (11) to compute $f(\bar{I}, \sigma_0)$.

We recommend using relative rather than absolute threshold values. The absolute value of an intensity gradient τ changes whether an image has a bit depth of say 8, 12, or 16 bits. However, when the relative threshold value $\tau\%$ is defined as the fraction of the pixels having an intensity gradient less or equal to τ then its value does not depend on the image bit depth. To ease writing, we will also use the symbol τ rather than $\tau\%$ with the understanding that the threshold is expressed as a ratio.

B. Image Acquisition Protocol

We apply our noise characterization method to five types of medical images.

1) *Magnetic Resonance Images*: We used the image acquisition protocol for the clinical assessment of knee-joint osteoarthritis at the University of Montreal Hospital Research Centre (CRCHUM). Magnetic resonance imaging (MRI) is performed with a Siemens, 1.5 T horizontal-bore magnet using a three-dimensional (3-D) FISP sequence with fat saturation or a 3-D double-echo in steady-state (DESS) sequence. Both sequences provide adequate contrast for the bone-cartilage and the cartilage-synovium interfaces.

The frequency-encoding direction is from head to foot and the phase-encoding direction, from anterior to posterior. Each 12-bit image is 512×480 (zero filled to 512×512) with a square field of view of 160 mm, giving an effective voxel size of $0.31 \text{ mm} \times 0.39 \text{ mm} \times 1 \text{ mm}$ (FISP) or $0.31 \text{ mm} \times 0.39 \text{ mm} \times 2 \text{ mm}$ (DESS). The imaging parameters are $\text{TR}/\text{TE} = 42/7$ with a flip angle of 20° (FISP) or $\text{TR}/\text{TE} = 25/9$ with a flip angle of 40° (DESS). To increase the number of data points and to maximize the dynamic range, we use five images taken from each of the FISP (110 images) and the DESS (56 images) data blocks. The five images sample uniformly the data blocks.

2) *Radiography From a Micro Strip Gas Chamber With an Amplifier Grid* [19]: We used a MICROMEAS testing bed (Biospace Instruments, France) installed at the École de technologie supérieure in Montreal. This system has a fan-beam geometry which allows the direct detection of X-rays using a 6-bar Xenon gas chamber. The beam is $254 \mu\text{m}$ thick and is recorded on a 1764 pixel grid. The 16-bit image (1 mAs) was acquired line by line as the source (70 kV_p) and the detector scanned an object while moving at a uniform speed. The effective pixel size is $254 \mu\text{m} \times 254 \mu\text{m}$ giving a field of view of $0.40 \times 0.45 \text{ m}^2$.

3) *Computed Radiography*: We used the image acquisition protocol for the clinical assessment of scoliosis used at Montreal's Sainte-Justine Hospital. The radiography was acquired on a FUJI FCR 7501 using a 90-kV_p source. The square pixel size is 400 μm . The 10-bit image is 1281 $\times$ 769 and has a field of view of 0.51 $\times$ 0.31 m². The image intensity was compressed logarithmically to facilitate clinical assessment, the original, uncompressed image was not available to us.

We also used the General Electrics Revolution XQ/i installed at Montreal's Hôtel-Dieu. A series of 94 images (10 mAs) was acquired using a 70-kV_p source without a anti-diffusion grid. Each 14-bit image is 2021 $\times$ 1751 and has a field of view of 0.40 $\times$ 0.35 m², giving a square pixel size of 200 μm .

4) *Radiography on a Screen Film*: Two X-ray films were scanned simultaneously with a Vidar VXR Twain DS (32 bit) and Microsoft Photo Editor scan software. The multi-film-support scanner has a moving stage and stationary sensor, optics and illumination. The 12-bit knee image is 2100 $\times$ 1524 and was scanned at 300 dpi giving a square pixel size of 85 μm for a field of view of 0.18 $\times$ 0.13 m². The 12-bit test pattern image (AGFA) has a field of view of 0.43 $\times$ 0.36 m² but we only use a 504 $\times$ 3060 image section for a field of view of 4.2 $\times$ 25.9 cm².

5) *Computerized Tomography*: We used the Pickers 5000 installed at CRCHUM. A series of 101 images (275 mAs) was acquired using a 100 kV_p source. Each 16-bit image is 512 $\times$ 512 and has a field of view of 0.48 $\times$ 0.48 m², giving a square pixel size of 936 μm . The image slices were 5 mm thick and a standard filter for thoracic imaging was used.

IV. RESULTS AND DISCUSSION

This section compares the measurements of noise standard deviation (or variance) in five types of medical images with their theoretical noise models. Our method is much less sensitive to the tuning parameters (the filter size and the threshold) usually associated to measurements of local noise variance. This should lead to more efficient automated image analysis.

Since the images are recorded with a limited precision (e.g., 10-bit images) the rescaling and the truncation of the intensity values always modify their noise statistics. We take into account the rescaling effects in our analysis but the truncation artefacts cannot be removed from the data. Such artefacts are negligible for most of the images used in this work. Image clipping due to pixel saturation by the noise generally affects the high-intensity end of the images where the noise variance is the greatest. This saturation effect was negligible in this work since the results that follow reproduce quite well the theoretical models for most of the intensity ranges of the images.

A. Magnetic Resonance Images

Rician noise differs greatly from Gaussian and Poisson noises. The variance of a Gaussian noise is constant whereas the variance of a Poisson noise is proportional to the noise mean (5). Rician noise is such that the noise variance depends non linearly on the noise mean. This is effectively what we observe on MR images.

Fig. 1(a) and (c) shows the same knee observed under the FISP and the DESS imaging sequences. The white circles are

cross sections through a horizontal cylinder filled with a liquid providing a strong signal. For the DESS sequence, two images are acquired and averaged together. As a result

$$\sigma_{\text{DESS}}(I) = \frac{\sigma(I)}{\sqrt{2}} \quad (24)$$

with $\sigma(I)$ given by (11). The noise standard deviation σ_{DESS} must be multiplied by $\sqrt{2}$ for comparison with the predictions of a Rician noise model.

Fig. 1(b) and (d) shows how the noise standard deviation depends on the magnitude for each sequence. The underlying curve on Fig. 1(b) is a fit to the theoretical relationships (10) and (11) and reproduces well the overall distribution of the data except for a few points in the bend of the curve. A better fit is found on Fig. 1(d) where the Rician noise model takes into account the $\sqrt{2}$ factor. In the analysis of Fig. 1(a) and (c), 59% and 93% of the image pixels were used.

The first data points in Fig. 1(b) and (d) are superfluous and result from our choice of the filter size and threshold values. Using a larger filter size and a smaller threshold eliminates the spurious points but also reduces the magnitude data range and prevents an accurate estimation of σ_0 .

Good estimates of σ_0 can usually be found by measuring the mean magnitude $\bar{I}$ or the standard deviation σ_I in a background region of an image and by using (12) or (13) to obtain σ_0 . This technique was applied to Fig. 1(a) and (c) as a validation step for our method. The background region is a rectangular section in the lower left corner of both figures. The results are listed in Table I. Comparison of the predictions of (12) and (13) with the least squares fit parameters are in good agreement. It also suggests that (13) is a more accurate estimator of σ_0 than is (12).

The remaining part of this section presents a detailed example of the application of the method to the DESS image [Fig. 1(c)]. Fig. 2 shows the effect of the threshold τ on masking edge pixels. Fig. 2(a)–(c) shows binary masks generated from the original image using threshold values $\tau_a < \tau_b < \tau_c$. The unmasked pixels appear in white and represent respectively 27%, 93.5% and 97% of the total number. Fig. 2(c)–(e) shows the corresponding noise standard deviation results computed from the unmasked pixels. The smallest threshold value results in too few pixels being selected [Fig. 2(a)] and, thus, in poor statistics and a limited intensity range [Fig. 2(d)]. The pixels with the largest noise intensity gradients are also eliminated from the analysis, even in flat intensity plateaus, which results in an underestimation of the noise variance. Fig. 2(c) represents the other extreme where a large threshold value results in too many selected pixels. The intensity range on Fig. 2(f) is the largest one but nonlinearities are also observed in the data; the shape of the curve in the bend is not well reproduced. The middle two panels provide the best results where the data and the theoretical model agree and the masked pixels are located along the most significant image edges.

Fig. 3(a) shows the WMSE for a range of threshold values. For each value of τ corresponds: 1) a binary mask; 2) the resulting $(\bar{I}, \sigma_N, \sigma_{NN})$ data set; 3) a nonlinear fit to the data set; and 4) the value of the WMSE for that fit (19). A minimum is observed at $\tau \cong 92.5\%$ which was the value used to generate




Fig. 1. Examples of typical MRI sagittal slices of a knee observed under (a) the FISP and (c) the DESS imaging sequences. (b) Relation between the noise standard deviation and the magnitude in (a). The curve is a robust least squares fits to a Rician noise model of parameter $\sigma_0 = 14.47 \pm 0.12$ (d). Relation between the noise standard deviation and the magnitude in the image (c). The standard deviation measurements were multiplied by $\sqrt{2}$ for comparison with a Rician noise model of parameter $\sigma_0 = 20.93 \pm 0.14$. The error bars correspond to ± 2 rms values of the standard deviation.

TABLE I

RICE NOISE LEVEL ESTIMATION BASED ON IMAGE BACKGROUND STATISTICS. THE FIRST COLUMN LISTS THE MR ACQUISITION SEQUENCES. THE SECOND COLUMN LISTS THE MEAN AND THE STANDARD DEVIATION OF THE MAGNITUDE DATA AS MEASURED IN A BACKGROUND REGION OF EACH IMAGE. THE THIRD COLUMN LISTS THE PREDICTED NOISE LEVELS ESTIMATED USING (12) AND (13) AND THE DATA IN THE SECOND COLUMN. THE PARAMETER VALUES ESTIMATED USING OUR METHOD ARE LISTED IN THE LAST COLUMN. THE ERROR BARS CORRESPOND TO ± 2 ROOT MEAN SQUARED (RMS) VALUES OF THE STANDARD DEVIATION. THE ERRORS ARE ESTIMATED USING THE BOOTSTRAP METHOD [15]

Acquisition Sequence	Measurements	Theory	Least Squares Fit
FISP	$\mu = 17.23 \pm 0.20$ $\sigma = 9.30 \pm 0.14$	$\sigma_0 = 13.75 \pm 0.16$ $\sigma_0 = 14.20 \pm 0.22$	$\sigma_0 = 14.47 \pm 0.12$
DESS	$\mu = 24.84 \pm 0.16$ $\sigma = 9.59 \pm 0.11$	$\sigma_0 = 19.82 \pm 0.13$ $\sigma_0^* = 20.70 \pm 0.24$	$\sigma_0 = 20.93 \pm 0.14$

* After taking into account the $\sqrt{2}$ factor

Fig. 1(d) and Fig. 2(b) and (e). This represents the best compromise between the intensity range, the scattering of the data points and the agreement with the theoretical model.

Fig. 3(b) shows that the estimated σ_0 parameter in each fit (the noise level) does not display such minimum/maximum behavior and tends to vary monotonously with the threshold value. That




Fig. 2. (a)–(c) Binary masks generated after applying the thresholds $\tau_a < \tau_b < \tau_c$ respectively to the intensity gradients in the DESS knee image. (d)–(f) Noise standard deviation results for each mask above. The standard deviation measurements were multiplied by $\sqrt{2}$ for comparison with Rician noise models (underlying curves). The error bars are not shown to avoid cluttering the figures.

remains true for all the other parameters estimated in this work. This observation strengthens the importance of using the WMSE statistics rather than the parameter values or a *Chi-by-eye* criterion to estimate the optimal threshold value.

B. X-Ray Images

1) *Linear Intensity Scaling*: Fig. 4(a) is a radiography of the torso of a humanoid phantom (Sectional Phantom, The Phantom Laboratory, Salem, NY). The resin matrix of the phantom encloses real bones and uses air-filled cavities to simulate the lungs. The data was obtained using a micro-strip gas chamber with an amplifier grid [19]. The MICROMEGAS system was designed to minimize the radiation exposure of a patient. Diagnostic radiation is known to be a causative factor of breast cancer in scoliotic patients [20].

Because the X-ray detector records a signal proportional to the X-ray flux, Fig. 4(b) shows that the noise variance measured on the image is proportional to the mean local intensity according to (5). There is a good agreement between the data and the predictions of the simple model assuming a mono chromatic X-ray beam.

There is gap in the intensity data (from 5000 to 60 000) that corresponds to the pixels forming the edge of the torso. These pixels are eliminated from the analysis because the corresponding intensity gradients are very large. The information for intensity data beyond 60 000 represents only 8% of the unmasked pixels in the bright background and is not shown here (the data points lay near the straight line). Forty-two percent of the image pixels were used in the noise analysis.

2) Logarithmic Intensity Scaling:

a) *Computed Radiography*: Computed radiography (conventional X-ray source with imaging phosphor plates) can also

reduce the required dose of radiation in comparison to standard radiology. When an X-ray beam strikes a photostimulable phosphor detector, most of the absorbed X-ray energy is trapped in the imaging plate and can be read out later using a laser beam [13]. The laser light stimulates the emission of trapped energy in the plate and visible light is released from the plate, collected by a fiber-optic light guide and sent to a photomultiplier tube. The reemitted flux is proportional to the X-ray beam intensity.

In the following example, we use a computed radiography that was saved on disk using a logarithmic intensity scaling for which the transfer function is unknown to us. Thus, the image gray levels correspond to rescaled optical densities rather than the X-ray beam intensities. The variance of the noise in the image should, thus, depend on the mean local optical density according to (8).

As described in Section II-B, the X-rays follow a compound Poisson noise process where their numbers follow a Poisson noise distribution and their energy converted counts follow a compound Poisson noise distribution due to the wide energy spectrum. During the last stage of image acquisition, the energy/intensity distribution is modified by the logarithmic intensity scaling that simulates an ideal film-conversion characteristics.

Fig. 5(a) shows the chest radiography of a patient with scoliosis. Fig. 5(b) reveals a good agreement between the data and the underlying logarithmic fit to (8) based again on a simple model assuming a mono chromatic X-ray beam. 41% of the image pixels were used in the noise analysis.

b) *Radiography on a Screen Film*: A screen film is characterized by a Hurter and Driffield (H&D) curve which is a plot of a film's optical density as a function of the log of exposure [13]. The curve is a straight line within the operational intensity range of the film. Outside that range, the optical density levels




Fig. 3. (a) Relation between the weighted mean square error and the threshold value. (b) Relation between the estimated noise level σ_0 and the threshold value. A smoothing spline (a) and a cubic curve (b) were used to enhance the visual appearance of the data on both panels.

off at low and high exposures due to under exposition and saturation of the film. The film must first be scanned before noise can be studied on it and the scanning process introduces photon, read-out and quantization noise in addition to grain noise already present in the film. The impact of the scanner-induced noise on the total noise balance must be evaluated to differentiate it from the film grain noise.

We first measured the noise variance on a screen-film test pattern (AGFA) which is a synthetic image without grain noise [Fig. 6(a)]. The analysis was applied to the horizontal band identified with an arrow. 72% of the band pixels were used in the analysis. Photon noise from the scanner illumination is the dominant noise source and its variance should obey (8). This is true for the upper half of the intensity range as Fig. 7 shows. There is gradual levelling off at relative optical densities $D \leq 30 \times 10^3$. According to Fig. 7, this noise has a standard deviation of $\sigma_E \cong 32$ which, for a 16-bit data range, gives the scanner an effective range of 11 bits. (The 12 bit-images were automatically rescaled to fit a 16-bit data range.) The levelling off is due to quantization and read-out noises. Whereas some noise was produced by inhomogeneities in the ink distribution, most of it was generated during the scanning process.




Fig. 4. (a) Chest radiography of a phantom torso (positive image) observed with a micro-strip chamber using a 70-kV_p beam. (b) Relation between the noise variance and the pixel intensity in the image. The line is a robust least squares fit. The error bars correspond to ± 2 rms values of the standard deviation.

Indeed, the test pattern film appears smoother visually than on the scanned image.

Fig. 6(b) shows the radiography of a patient who had knee surgery. Fig. 7 shows the corresponding noise variance to obey (8) with a breakdown of the logarithmic behavior at relative optical densities $D \leq 40 \times 10^3$ where the low X-ray fluxes are not properly recorded by the film. The noise variance at low optical density is about 15 times larger on the radiography than on the test pattern (both images were scanned simultaneously and, thus, have similar read-out and quantization noises). The levelling off on the radiography is not due to scanning noise only. Most of the noise in the base and fog section of the H&D characteristic curve arises from the development of some unexposed grains [13]. The grain rough texture on small scales generates the grainy appearance of radiographies (mostly visible in the bone regions) and is an important source of noise in our analysis. The toe section of the same curve is found in the interval $35 \times 10^3 \leq D \leq 40 \times 10^3$ where the film has a slow response to light. There is no image saturation on the film. 71% of the image pixels were used in the analysis.

The results presented here indicate that the scanner quantization noise dominates over all other types of noise in the radiography. The photon noise that was present during the X-ray exposure and that was recorded in the grain distribution is negligible when compared to the scanning noise. This is shown by the two roughly identical distributions of data points along the straight




Fig. 5. (a) Chest radiography of a patient with scoliosis (negative image) obtained with the imaging-plate system. (b) Relation between the noise variance and the relative optical density in the image (a). The bone and the background regions have respectively low and high optical densities. The line is a robust least squares fit. The error bars correspond to ± 2 rms values of the variance.

line on Fig. 7. There is no visible vertical offset that would have revealed the presence of an another significant noise source beside scanning noise.

C. Limits of the Method

We discuss in this section about the main limitations of the noise analysis method and the type of images where it can be used successfully.

1) *Statistical Homogeneity*: The statistical homogeneity of noise is the most fundamental assumption upon which our method is based. It assumes that the noise properties are the same across an image and the relationship (1) applies everywhere. The hypothesis can be tested by imaging several times the same object, by computing the mean and the variance images, and by plotting a scatterplot of the $(\bar{I}, \sigma_N^2)$ data. A coherent data distribution should be visible if a global relationship (1) exists. We tested the hypothesis using a series of 101 CT scans of the phantom torso obtained on the Pickers 5000




Fig. 6. (a) Screen-film test pattern (used as a negative image). The arrow points to the image band used for noise analysis. (b) Radiography of a patient who had knee surgery (negative image). The white objects are screws anchoring the ligaments between the joints.

and a series of 94 computed radiographies of the same object obtained on the General Electrics Revolution XQ/i.

Fig. 8(a) and (b) shows the mean and the variance images of the 101 CT scans. To get the larger possible intensity range, the phantom torso was positioned as if the patient were sitting in the scanner which explains the unconventional thoracic view. The CT data was recorded in Hounsfield units [13]. The data does not correspond to X-ray intensity or film optical density but corresponds instead to a linear attenuation coefficient. No strongly coherent data distribution appears in the $(\bar{I}, \sigma_N^2)$ scatterplot [Fig. 8(c)]. The analysis of series of synthetic CT images also produces very noisy nonlinear relationships between pixel mean and variance that are image-dependent. Similar observations were made by Lu *et al.* [3] using 900 CT images of a physical phantom. The very noisy shape of the scatterplot indicates that the relationship (1) does not apply everywhere; such images can not be analyzed using our method. Indeed, Fig. 8(b) shows that the noise variance depends on the local intensity and the radial distance from the image center. The centrosymmetric




Fig. 7. Relations between the noise variance and the relative optical density for the radiography and the test pattern images. The bone and the background regions in the radiography have respectively low and high optical densities. The line is a robust least squares fit and the error bars are not shown to avoid cluttering the figure.

ring structure is generated by the filtering step in the image reconstruction algorithm.

We observed Gaussian rather than Poisson distributions of local pixel intensities on real and synthetic CT scans series thereby reproducing another observation of Lu *et al.* [3]. The image reconstruction algorithms involve weighted linear combinations of intensities that transform the noise PDF from Poisson to Gaussian according to the central limit theorem. The image smoothing also introduces noise correlation that can not be handled with our method.

These results were predicted by Lei and Sewchand [5] in their statistical description of X-ray CT imaging, from the projection data to the reconstructed image. The noise PDF on such an image has an asymptotic Gaussian distribution. Moreover, neighboring pixels are statistically dependent when their spatial separation is smaller than a threshold depending on the physical imaging system and image reconstruction algorithm.

Fig. 8(d) and (e) presents the mean and the variance images of the 94 computed radiographies. The $(\bar{I}, \sigma_N^2)$ scatterplot [Fig. 8(f)] shows a coherent data distribution where the noise variance increases linearly with the image intensity. This is the expected behavior for an image with Poisson noise. The two clusters of points above the intensity of 800 correspond to the bright background level which is slightly different below and above the shoulders of the phantom torso.

The shape of the data distribution on Fig. 8(f) indicates that the relationship (1) applies to every image pixel and the noise is statistically homogenous. Our noise-analysis method, designed for a single image, was applied to one of the 94 computed radiographies and the results are shown as data points on Fig. 9. The underlying curve is a linear fit computed using all 94 images by first binning the intensity data on Fig. 8(f) and then by computing the mean intensity and noise variance in each bin. There is a good agreement between the results of the two methods even though our noise-analysis method tends to slightly overestimate the noise variance of the low intensity pixels. The middle third

of the intensity range is missing due to the pixels with the large intensity gradients being masked out. They correspond to the pixels along the torso boundary.

2) *Image Smoothness*: Better performances are achieved if numerous intensity plateaus are present in an image but such plateaus are not necessarily required to obtain good results. Synthetic images show that the method works also well on smooth image regions and any sharp discontinuity is eliminated from the analysis.

When most of the high-intensity pixels in an image lie along the boundaries between smooth portions of the image, the pixels are masked out by the edge detection and thresholding procedure. Thus, the resulting $(\bar{I}, \sigma_N^2)$ data set does not contain any high-intensity information. However, the relationship (1) is often monotonous and can be extrapolated to larger values of intensity when a theoretical model [e.g., (5), (8), (10), and (11)] is available. Mammography is a good example. The parenchyme tissue distribution creates a smooth overall signal in mammograms upon which are superposed bright microcalcifications. Masking those small-scale discontinuities narrows the intensity range of the $(\bar{I}, \sigma_N^2)$ data but the relationship (1) can nevertheless be extracted from the much smoother surrounding signal since the relationship is generally linear (5) or exponential (8). This is not possible when an unknown nonlinear operation was initially applied to the images. Extrapolation can only be used if the noise model is known; the relevant parameters can, thus, be estimated from the data.

3) *Linear vs Nonlinear Noise Models*: The precision of the parameter estimation depends on the shape of the relationship (1). A graph of the weighted mean squares error (19) often displays a well defined minimum [Fig. 3(a)] when the relationship (1) is nonlinear. This is not always the case when (1) is linear; the WMSE may simply increase monotonically with the threshold value. For example, a scatterplot of three $(\bar{I}, \sigma_N^2)$ measurements may result in a smaller WMSE than 100 of them slightly scattered along a line and covering a large intensity range. A more robust definition of the WMSE is, thus, required. Similar problems may be encountered if the intensity range of the measurements does not fully cover the non linear domain of the relationship (1).

4) *Image Textures and Correlated Noise*: Strong image textures are eliminated from the analysis due to their large intensity gradients. Smoother image textures can be a problem however when their correlation lengths are smaller than the size of the smoothing filter since they contribute to the local image variance. For such images with large signal-to-noise ratios, the texture-related variance is greater than the noise variance and no accurate relationship (1) may be extracted from the data. The image textures were not a problem in this work given the agreement between the $(\bar{I}, \sigma_N)$ measurements and the theoretical models.

V. CONCLUSION

We developed a method to study various types of image noise for which a relationship of the form (1) exists. Using local measurements of the mean intensity $\bar{I}$, the noise standard deviation σ_N , and its associated error σ_{NN} , we can determine the type of




Fig. 8. (a) Mean and (b) variance images of a series of 101 CT scans of the phantom torso. (c) Scatterplot of the local mean and variance data. (d) Mean and (e) variance images of a series of 94 computed radiographies of the same object. (f) Scatterplot of the local mean and variance data. The straight line is a robust least squares fit through the data.




Fig. 9. Relation between the noise variance and the pixel intensity on the GE Revolution XQ/i system. The straight line is a robust least squares fit computed using the local mean and the variance of a series of 94 images of the phantom torso. The data points were measured for a single image using the method presented in this paper. The error bars correspond to ± 2 rms values of the variance.

noise (Poisson, Rician, Gaussian) and compute the parameters $\alpha_1, \alpha_2, \alpha_3, \dots$ in the relationship (1). The method relies on robust statistical estimators and is less dependent on tuning parameters. It was applied successfully to six medical images acquired on five different detecting systems. The noise variance results for MR images and computed radiographies were in agreement with the results obtained using different methods.

Ideally, no assumption about the shape of the relationship (1) should be made. This can be easily done when a series of noisy images of a same object are available [2], [3] or when a special test pattern such as a grey level wedge can be used. The problem is much more difficult when a single image is available. The shape of the experimental relationship (1) often depends on several tuning parameters that must be properly selected. It is unwise to base our judgement on *chi-by-eye* least squares fits including only short, well-behaved sections of a $(\bar{I}, \sigma_N)$ scatter plot or on the basis of the fit parameters $\alpha_1, \alpha_2, \alpha_3, \dots$ values. Special care must also be taken before using an experimental relationship (1) as a basis to further analysis, unless it is properly determined.

The image noise must be uncorrelated and statistically homogeneous in order for the method to be useful. It should, thus, not be applied to CT images which require a different noise analysis approach.

The method performs best on raw or unprocessed (and uncompressed) images. Such images are usually modified by an operator to enhance their contrast and their edges, to smooth out the noise or to deblur the images. Unfortunately, these operations also alter the very noise statistics we intend to study. The raw images should be used as such, i.e., cosmetic-free, which fortunately requires the least amount of work.

The method presented in this study has several possible applications and future work will focus on: 1) image restoration; 2) image calibration; 3) nonlinear image denoising; and 4) improving the performances of two-dimensional/3-D segmenta-

tion algorithms for the reconstruction of various 3-D tissue distributions [21].

ACKNOWLEDGMENT

The authors would like to thank A. Bleau and S. Rajagopalan for their critical reading of the manuscript and to P. Després, F. Dinelle, and A. Dagenais for their help during data acquisition.

REFERENCES

- [1] S. I. Olsen, "Noise variance estimation in images," presented at the 8th Scandinavian Conference on Image Analysis, Tromsø, Norway, May 25–28, 1993.
- [2] B. Waegli, "Investigations into the noise characteristics of digitized aerial images, Proc. ISPRS Commission II Symp., Cambridge, U.K., July 13–17, 1998," *Int. Arch. Photogrammetry and Remote Sensing*, vol. XXXII/2, pp. 341–348.
- [3] H. Lu, X. Li, I. T. Hsiao, and Z. Liang, "Analytical noise treatment for low-dose CT projection data by penalized weighted least-squares smoothing in the K-L domain," *Proc. SPIE, Medical Imaging 2002*, vol. 4682, pp. 146–152.
- [4] H. Gudbjartsson and S. Patz, "The rician distribution of noisy MRI data," *Magn. Reson. Med.*, vol. 34, pp. 910–914, Dec. 1995.
- [5] T. Lei and W. Sewchand, "Statistical approach to X-ray CT imaging and its applications in image analysis—Part I: Statistical analysis of X-ray CT imaging," *IEEE Trans. Med. Imag.*, vol. 11, Apr. 1992.
- [6] J. C. Dainty and R. Shaw, *Image Science*. New York: Academic, 1974.
- [7] R. Brügelmann and W. Förstner, "Noise estimation for color edge extraction," in *Robust Computer Vision*, W. Förstner and S. Ruwiedel, Eds. Karlsruhe, Germany: Wichmann, 1992, pp. 90–107.
- [8] J. S. Lee and K. Hoppel, "Noise modeling and estimation of remotely-sensed images," *Proc. 1989 Int. Geoscience and Remote Sensing Conf.*, vol. 2, pp. 1005–1008.
- [9] R. A. Schowengerdt, "Noise Models," in *Remote Sensing. Models and Methods for Image Processing*, 2nd ed. New York: Academic, 1997, ch. 4, pp. 127–131.
- [10] I. N. Bankman, T. S. Spisz, and S. Pavlopoulos, "Two-dimensional shape and texture quantification," in *Handbook of Medical Imaging, Processing and Analysis*. New York: Academic, 2000, pp. 215–230.
- [11] A. Wismüller, F. Vietze, and D. R. Dersch, "Segmentation with neural networks," in *Handbook of Medical Imaging, Processing and Analysis*. New York: Academic, 2000, pp. 107–126.
- [12] M. Fuderer, "The information content of MR images," *IEEE Trans. Med. Imag.*, vol. 7, Oct. 1988.
- [13] J. T. Bushberg, J. A. Seibert, E. M. Leidholdt Jr., and J. M. Boone, *The Essential Physics of Medical Imaging*, 2nd ed. Baltimore, MD: Lippincott Williams & Wilkins, 2002.
- [14] J. Sijbers, "Signal and Noise Estimation from Magnetic Resonance Images," Ph.D. thesis, Univ. Antwerp, Antwerp, Belgium, 1999.
- [15] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C, The Art of Scientific Computing*, 2nd ed. Cambridge, MA: Cambridge Univ. Press, 1992.
- [16] D. C. Hoaglin, F. Mosteller, and J. W. Tukey, Eds., *Understanding Robust and Exploratory Data Analysis*. New York: Wiley, 1983.
- [17] W. C. Cleveland, *Visualizing Data*. Summit, NJ: Hobart Press, 1993, pp. 110–119.
- [18] *Curve Fitting Toolbox User's Guide*, The MathWorks Inc., Natick, MA, 2002.
- [19] Y. Giomataris, P. Rebourgeard, J. P. Robert, and G. Charpak, "MICROMEGAS: A high-granularity position-sensitive gaseous detector for high particle-flux environments," *Nucl. Instrum. Meth. A*, vol. 376, pp. 29–35, 1996.
- [20] M. M. Doody, J. E. Lonstein, M. Stovall, D. G. Hacker, N. Luckyanov, and C. E. Land, *Spine*, vol. 25, no. 2052, 2000.
- [21] C. Kauffmann, P. Gravel, B. Godbout, A. Gravel, G. Beaudoin, J. P. Raynauld, J. Martel-Pelletier, J. P. Pelletier, and J. A. De Guise, "Validation of a computer-aided method for quantification of thickness and volume changes in human knee cartilage using MRI," *IEEE Trans. Biomed. Eng.*, vol. 50, pp. 978–988, Aug. 2003.