

MR noise measurements and signal-quantization effects at very low noise levels

Olaf Dietrich, José G. Raya, Maximilian F. Reiser

Josef Lissner Laboratory for Biomedical Imaging, Department of Clinical Radiology – Grosshadern, LMU Ludwig Maximilian University of Munich, Munich, Germany

ELECTRONIC PREPRINT VERSION:

Not for commercial sale or for any systematic external distribution by a third party

Final version: *Magn Reson Med* 2008; 60(6): 1477–1487. <URL:<http://dx.doi.org/10.1002/mrm.21784>>

Abstract

The well-known noise distributions of MRI data (Rayleigh, Rician, or non-central chi-distribution) describe the probability density of real-valued (i.e. floating-point) signal intensities. MR image data, however, is typically quantized to integers before visualization or archiving. Depending on the scaling factors applied before the quantization and the signal-to-noise ratio (SNR), very low noise levels with substantial artifacts due to the quantization process can occur. The purpose of this study was to analyze the consequences of the signal quantization, to determine the theoretical absolute lower limit for noise measurements in discrete data, and to evaluate an improved method for noise and SNR measurements in the presence of very low noise levels. Image data were simulated with original noise levels between 0.02 and 2.00. Noise measurements were performed based on the properties of background and foreground data using the conventional approach, which exploits the standard deviation or mean value of the signal, and a maximum-likelihood approach based on the relative frequencies of the observed discrete signal intensities. Substantial deviations were found for the conventionally determined noise levels, while noise levels comparable to or lower than the

quantization error can be accurately estimated with the proposed maximum-likelihood approach.

Keywords:

Magnetic resonance imaging, Signal quantization, Statistical noise distribution, Signal-to-noise ratio

Corresponding author:

Olaf Dietrich, PhD
Josef Lissner Laboratory for Biomedical Imaging
Department of Clinical Radiology – Grosshadern
LMU Ludwig Maximilian University of Munich
Marchioninstr. 15
81377 Munich
Germany
Phone: +49 89 7095-3623
Fax: +49 89 7095-4627
E-mail: od@dtrx.net

(Parts of this study were presented on the ISMRM annual meeting 2008, abstract #1529.)

Introduction

The properties of the MR signal intensity in the presence of noise, such as the statistical distribution, mean value, and standard deviation, have been extensively discussed in several publications (1–11). It is commonly assumed that noise in MRI raw data is normally distributed in each receiver channel. The statistical signal distribution in the final image depends on the reconstruction and channel-combination technique. Signal statistics are described, e.g., by the Rayleigh or Rician distribution for single-channel magnitude data after conventional Fourier-transform reconstruction (1–8) and by the non-central χ -distribution in the case of a root-sum-of-squares (RSS) reconstruction of multi-channel data (9,12,13). These statistical distributions play an important role for the measurement of the signal-to-noise ratio (SNR) or the contrast-to-noise ratio (CNR) in MR images, both of which are based on the exact determination of the image noise level. Several methods for noise measurements have been proposed depending on the noise properties. If a spatially uniform noise distribution can be assumed, then the noise level can be determined by analyzing the background noise (i.e. the MR signal in air) of an MR image (2,14,15). If the noise level is variable over the image (e.g., as a consequence of parallel-imaging reconstruction), then the noise level should be determined at the same location of the image foreground as the signal, for instance by analyzing the signal in a difference image (10,15–17) or by pixelwise evaluation of the signal time-course in a large number of repeated acquisitions (10,15,17).

The statistical distributions mentioned above describe the probability density of real-valued signal intensities, i.e. signal intensities represented by floating-point numbers. Image data, however, are commonly quantized to integers before visualization or archiving in the DICOM format (18). Depending on the signal-to-noise ratio and the scaling factors applied before the quantization, very low noise levels with substantial artifacts due to the quantization process can occur. In this case, the evaluation of the signal mean value or the standard deviation in a background region, in subtraction data, or in a pixel time-course will lead to inaccurate estimations of the true noise level if the quanti-

zation of the signal intensity is ignored. Consequently, the original noise standard deviation as well as the SNR determined from these parameters with the methods mentioned above will be over- or underestimated. The purpose of this study was to analyze the consequences of the signal quantization and to evaluate an improved method for the exact determination of very low noise levels in discrete image data.

Theory

Continuous signal distributions

If we assume normally distributed noise superimposed on the signal of the real and imaginary part of raw data in each receiver channel, then real and imaginary part of the complex image data of each channel reconstructed by Fourier transform are also superimposed by white Gaussian noise (3,4). The statistical distribution of the signal, x , is described by the probability density

$$P_{\text{Gauss}}(x; \mu, \sigma) = \frac{1}{\sqrt{2\pi} \sigma} \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right) \quad [1]$$

with original standard deviation σ and mean value μ (the mean values, μ , of the real and imaginary part will in general be different; in background areas, this mean value is zero). We call σ the "original" standard deviation, since the actual standard deviation measured in the final (e.g., magnitude) image will be different from σ , in general.

MR images are usually presented as magnitude data. The most important probability densities for magnitude image signals in MRI are the Rician distribution (2,3,7), which describes signal statistics after magnitude calculation of complex data,

$$P_{\text{Rice}}(x; \mu, \sigma) = \frac{1}{\sigma^2} x \exp\left(-\frac{x^2 + \mu^2}{2\sigma^2}\right) I_0\left(\frac{x\mu}{\sigma^2}\right), \quad [2]$$

and the non-central χ -distribution of signal intensities after RSS reconstruction of n channels (9),

$$P_{\text{nc}\chi}(x; \mu, \sigma, n) = \frac{\mu}{\sigma^2} \left(\frac{x}{\mu}\right)^n \exp\left(-\frac{x^2 + \mu^2}{2\sigma^2}\right) I_{n-1}\left(\frac{x\mu}{\sigma^2}\right), \quad [3]$$

where $I_n(z)$ is the modified Bessel function of the first kind of n th order. In the image background with zero original signal, $\mu = 0$, Eq. [2] simplifies to

the Rayleigh distribution (1):

$$P_{\text{Rayleigh}}(x; \sigma) = P_{\text{Rice}}(x; 0, \sigma) = \frac{1}{\sigma^2} x \exp\left(-\frac{x^2}{2\sigma^2}\right) \quad [4]$$

and Eq. [3] can be simplified to (9)

$$P_{\text{nc}\chi}(x; 0, \sigma, n) = \frac{1}{\Gamma(n)\sigma^2} \left(\frac{x}{2\sigma^2}\right)^{n-1} x^n \exp\left(-\frac{x^2}{2\sigma^2}\right), \quad [5]$$

where $\Gamma(n)$ is the gamma function. If, on the other hand, the signal-to-noise ratio is sufficiently high, then both the Rician and the non-central χ -distribution can be well approximated by Gaussian distributions with the original standard deviation, σ (Eq. [1]). This is important for noise measurements based on the pixel time-course in repeated acquisitions or on data of a difference image. In both cases, the SNR in the analyzed region should be sufficiently high to provide approximately normally distributed signal intensities. The standard deviations of the considered data are either the original standard deviation, σ , if a single pixel time-course is considered or the original standard deviation increased by a factor of $\sqrt{2}$ if a region of subtraction data is evaluated.

Discrete signal distributions

The last step of the image reconstruction commonly includes the transition from floating-point to (non-negative) integer signal intensities, i.e. the quantization (also called digitalization or discretization) of the intensity data, in order to map signal intensities to the, e.g., 12-bit integer range from 0 to 4095, which is frequently used in DICOM data. Prior to the quantization, the floating-point intensities will generally be scaled to an appropriate range; however, we will subsequently assume without loss of generality that all floating-point data are already in an appropriate range for quantization, i.e., the original noise level, σ , and image intensity, μ , refer to the data directly before the quantization process.

For real numbers, x , the corresponding quantized integer value, k , is obtained by applying the quantization function, Q , i.e., $k = Q(x)$. Different functions, $Q(x)$, are available for quantization such as simple truncation of the fractional part, $k = Q(x) = \text{trunc}(x)$, or symmetric arithmetic rounding (i.e., rounding to the nearest integer), $k = Q(x) = \text{round}(x)$. The quantization error, e , is the

difference between the original value, x , and the quantized value, k , i.e., $e = x - Q(x)$; for truncation, the quantization error is between 0 and 1; and for symmetric rounding, it is between -0.5 and 0.5 . It is also useful to define quantization boundaries, b_k , which separate intervals that are quantized to different values, i.e., two real numbers x_1, x_2 are quantized to the same value $k = Q(x_1) = Q(x_2)$, if and only if they lie between two successive quantization boundaries, $b_k \leq x_1, x_2 < b_{k+1}$. For instance, the quantization boundaries for truncation of non-negative numbers are 0, 1, 2, 3, ..., while the quantization boundaries for arithmetic rounding are 0, 0.5, 1.5, 2.5, ...

If the original image noise level, σ , is sufficiently large in comparison to the quantization error, i.e. $\sigma \gg 1$, then quantized MR image data are well described by the continuous probability density, $P(x; \mu, \sigma, n)$; here, n shall denote any other parameters such as number of channels. However, the continuous description by $P(x; \mu, \sigma, n)$ is not longer appropriate if σ is in the order of 1 or below, i.e. if image noise and quantization noise become comparable. In this case, the signal intensity cannot longer be considered as (almost) continuous quantity. Instead of characterization by a probability density, the distribution of the discrete signal intensities, k , is much better described by the relative frequency, $F(k; \mu, \sigma, n)$, i.e. we use a histogram (with natural bin size 1) for data characterization instead of a continuous density function.

If the mathematical expression of $F(k; \mu, \sigma, n)$ for the relative frequency of the signal k is known, then the original noise level, σ , can be estimated by comparing $F(k; \mu, \sigma, n)$ regarded as a function of σ with the observed counts, n_k , or the relative frequencies, f_k , of pixels with the discrete intensity, k . These counts, n_k , or relative frequencies, f_k , can be determined either in an image region of interest (ROI) or for a single pixel over the signal time-course in a series of repeated acquisitions, i.e. $n_k = \text{"# pixels with intensity } k\text{"}$ found in either a ROI or over the measured time-course. The relative frequencies, f_k , are obtained by dividing n_k by the total number, N , of pixels in the ROI or of acquired repetitions, i.e. $f_k = n_k / N$. This is done for $k = 0, 1, 2, 3, \dots, k_{\text{max}}$, where k_{max} is the maximum pixel intensity in the considered sample. We can then estimate an approximation of the original

standard deviation, σ , using a least-squares fit,

$$\sigma = \operatorname{argmin}_{\sigma} \sum_{k=0}^{k_{\max}} (F(k; \mu, \sigma, r) - f_k)^2. \quad [6]$$

Mathematically more rigid than this least-squares approach is the maximization of the corresponding likelihood function. This approach was recently proposed by Sijbers et al. for the estimation of the noise variance of Rayleigh-distributed MR data (8). The general likelihood function for multinomially distributed counts, $\{n_k\}$, is given by

$$L(\sigma, \mu; \{n_k\}, r) = \frac{M!}{\prod_{k=0}^{k_{\max}} n_k!} \prod_{k=0}^{k_{\max}} F(k; \mu, \sigma, r)^{n_k}. \quad [7]$$

and is typically maximized by minimizing $-\ln L$. Thus, the original standard deviation, σ , can be determined as

$$\begin{aligned} \sigma &= \operatorname{argmin}_{\sigma} \left[-\ln \left(\frac{M!}{\prod_{k=0}^{k_{\max}} n_k!} \prod_{k=0}^{k_{\max}} F(k; \mu, \sigma, r)^{n_k} \right) \right]. \quad [8] \\ &= \operatorname{argmin}_{\sigma} \left[-\sum_{k=0}^{k_{\max}} n_k \ln F(k; \mu, \sigma, r) \right] \end{aligned}$$

In order to apply this approach to image data, we must derive expressions for the relative frequency, $F(k; \mu, \sigma, r)$, for different situations, e.g. in the background noise of single-channel or RSS multi-channel data, in the foreground signal of MR image data, or in the signal of subtraction images. A general expression of the relative frequency, $F(k; \mu, \sigma, r)$, of quantized image data is the integral of the probability density with lower and upper integral limits, b_k and b_{k+1} :

$$F(k; \mu, \sigma, r) = \int_{b_k}^{b_{k+1}} P(x; \mu, \sigma, r) dx. \quad [9]$$

Hence, the relative frequencies of signal intensities, k , can be calculated from the cumulative distribution function of the considered probability density. The relative frequencies for the evaluation of background noise are derived from Eqs. [4] and [5], i.e.

$$F_{\text{Rayleigh}}(k; \sigma) = \int_{b_k}^{b_{k+1}} P_{\text{Rayleigh}}(x; \sigma) dx = -\exp\left(-\frac{x^2}{2\sigma^2}\right) \Big|_{b_k}^{b_{k+1}} \quad [10]$$

$$\begin{aligned} F_{\text{nc}\chi}(k; 0, \sigma, n) &= \int_{b_k}^{b_{k+1}} P_{\text{nc}\chi}(x; 0, \sigma, n) dx \\ &= -\exp\left(\frac{-x^2}{2\sigma^2}\right) \sum_{l=0}^{n-1} \frac{1}{l!} \left(\frac{x^2}{2\sigma^2}\right)^l \Big|_{b_k}^{b_{k+1}} \quad [11] \end{aligned}$$

If the signal intensities are evaluated in the image foreground, then the situation is slightly more complex since the discrete frequencies also depend on the original signal, μ , (and not only on the standard deviation). If, for instance, the pixel time-course is analyzed in repeated acquisitions, then the relative frequencies are:

$$F_{\text{Gauss}}(k; \mu, \sigma) = \int_{b_k}^{b_{k+1}} P_{\text{Gauss}}(x; \mu, \sigma) dx = \frac{1}{2} \operatorname{erf} \frac{x - \mu}{\sqrt{2}\sigma} \Big|_{b_k}^{b_{k+1}}, \quad [12]$$

where $\operatorname{erf}(z)$ is the error function. The dependence on μ becomes obvious, if we consider a very small standard deviation, $\sigma \ll 1$. If μ is somewhere in the middle between two quantization boundaries, e.g., $\mu = (b_k + b_{k+1})/2$, then almost all signal intensities will be quantized to k . If on the other hand, the mean value is very close to a quantization boundary, e.g., $\mu = b_k$, then about half of all signal intensities will be quantized to $k-1$ and the other half to k . In order to estimate the noise level, σ , based on Eq. [12], one can either determine the values of σ and μ that minimize the bracketed expression in Eq. [8] or one can first estimate μ from the original data such that only a one-dimensional minimization is required to find σ . A reasonable estimate of μ in the case of arithmetic rounding is the mean value of all pixel intensities in a series of repeated acquisitions. If data is quantized by truncation then the mean value must be increased by an offset of 0.5 to obtain an estimate for μ .

If foreground data of difference images are analyzed, it is useful to restrict the analysis to sufficiently large regions, in which the original (noise-free) signal intensity, μ , is approximately uniformly (or normally) distributed and covers an intensity range that is greater than the quantization distance (these are reasonable assumptions in typical MR

images, in which the signal always varies to a certain extent due to e.g. spatially variable coil sensitivities). In this case, the distribution of the subtracted discrete intensities does not longer depend on the original signal intensity, μ , due to averaging over all values of μ . Thus, μ does not appear in the final expression for the relative frequencies of the signal intensities, which is (independently of the quantization method) given by

$$\begin{aligned}
 F_{\text{Subt}}(k; \sigma) &= \int_{k-1}^k P_{\text{Gauss}}(x; 0, \sqrt{2}\sigma)(1-k+x)dx \\
 &+ \int_k^{k+1} P_{\text{Gauss}}(x; 0, \sqrt{2}\sigma)(1-x+k)dx \\
 &= \frac{\sigma}{\sqrt{\pi}} \exp\left(-\frac{x^2}{4\sigma^2}\right) \Big|_{k-1}^k - \frac{\sigma}{\sqrt{\pi}} \exp\left(-\frac{x^2}{4\sigma^2}\right) \Big|_k^{k+1} \\
 &+ \frac{(1-k)}{2} \operatorname{erf}\left(\frac{x}{2\sigma}\right) \Big|_{k-1}^k + \frac{(1+k)}{2} \operatorname{erf}\left(\frac{x}{2\sigma}\right) \Big|_k^{k+1}
 \end{aligned} \quad [13]$$

for $k = 0, \pm 1, \pm 2, \pm 3, \dots$ as derived in the Appendix.

Minimal detectable noise level

Based on Eqs. [10] to [13], it is possible to determine the minimum noise level, σ_{\min} , that can be estimated from a sample of size, N , of discrete image data. Obviously, if all signal intensities in a sample are quantized to the same value, k_0 , (e.g., $k_0 = 0$ in the background), then the standard deviation cannot be estimated. In the case of asymmetric distributions, the sample must include at least one pixel with an intensity, k_1 , different from the most frequently found intensity, k_0 , i.e.

$$N \sum_{k \neq k_0} F(k; \mu, \sigma_{\min}, r) = 1. \quad [14]$$

If background noise is evaluated and the most frequently found intensity, k_0 , is 0, then Eq. [14] can be transformed to

$$N \sum_{k=1}^{\infty} F(k; \sigma_{\min}, r) = N \int_{b_1}^{\infty} P(x; \sigma_{\min}, r) dx = 1. \quad [15]$$

For Rayleigh-distributed noise, this is equivalent to

$$N \left(-\exp\left(-\frac{x^2}{2\sigma_{\min}^2}\right) \right) \Big|_{b_1}^{\infty} = N \exp\left(-\frac{b_1^2}{2\sigma_{\min}^2}\right) = 1, \quad [16]$$

i.e., $\sigma_{\min} = b_1 / \sqrt{2 \ln N}$. A similar calculation is possible for multi-channel RSS data; however, the polynomial expression in Eq. [11] will generally require a numerical approach to determine σ_{\min} . In addition, it must be taken into account that the

most frequently found background intensity, k_0 , may be greater than 0 even for very low noise levels in the case of a large number of channels.

In the case of symmetric intensity distributions such as the Gaussian distribution of foreground data, a sample with at least three different signal intensities may be required to accurately estimate the noise level. If, for instance, the mean value, μ , of the normally distributed signal intensity is very close to a quantization boundary, then the intensities are approximately equally distributed to two discrete values k_0 and k_1 , independently of the standard deviation, σ . Thus the condition in Eq. [14] must be modified to

$$N \sum_{k \neq k_0, k_1} F_{\text{Gauss}}(k; \mu, \sigma_{\min}) = 1. \quad [17]$$

Using Eq. [12] and setting $k = k_0 + 1 = k_1$, this is equivalent to

$$\begin{aligned}
 1 &= N \left(1 - \int_{b_{k-1}}^{b_{k+1}} P_{\text{Gauss}}(x; \mu, \sigma_{\min}) dx \right) \\
 &= N \left(1 - \frac{1}{2} \operatorname{erf}\left(\frac{x-\mu}{\sqrt{2}\sigma_{\min}}\right) \Big|_{b_{k-1}}^{b_{k+1}} \right).
 \end{aligned} \quad [18]$$

The solution, σ_{\min} , of this equation depends strongly on the original signal, μ . Since the original signal is typically continuously distributed and not exactly known, a practical approach for the determination of σ_{\min} is to calculate the maximum over all possible values of μ , i.e. $\sigma_{\min} = \max_{\mu}(\sigma_{\min}(\mu))$.

In foreground data of difference images, it can be assumed that the analyzed region of interest contains a sufficiently broad uniform (or normal) distribution of (real-valued) signals prior to the superposition of noise as discussed above. Thus, there are always original signals very close to quantization boundaries resulting in a non-trivial distribution (i.e. a distribution containing more than a single signal intensity) even for extremely low noise levels. This distribution can be exploited to estimate the standard deviation of noise for practically all values of σ , i.e., $\sigma_{\min} = 0$.

Methods

Four typical experiments corresponding to the most frequently used techniques for SNR measurements were simulated to analyze the conse-

quences of signal quantization for the determination of the noise level (and thus eventually for SNR measurements): The noise level was estimated based on the statistical properties of (a) the background signal in single-channel data, (b) the background signal in RSS multi-channel data, (c) the time-course of a single pixel in the image foreground evaluated over a large number of repeated identical acquisitions, and (d) a foreground region in a difference image of two repeated (identical) acquisitions. The details of the simulations used for these four different experiments are described below. We compared the conventional approaches, which exploit the standard deviation or mean value of the simulated data to estimate the original noise level, σ , with the proposed maximum-likelihood method (Eqs. [8] and [10] – [13]), which estimates the original noise level, σ , from the counted frequencies, n_k , of signal intensities. The relative deviation of the estimated noise level, σ_{est} , and the actual original noise level, σ , was assessed as the ratio $\sigma_{\text{est}}/\sigma$. This ratio was evaluated and compared in all simulations for the conventional and the proposed maximum-likelihood technique.

In a first series of simulations, we analyzed the relative deviation, $\sigma_{\text{est}}/\sigma$, for different original noise levels, σ , ranging from 0.02 to 2.00 based on a relatively large sample size of $N = 25\,000$ in order to minimize statistical errors. The ratios, $\sigma_{\text{est}}/\sigma$, were plotted over the originally applied noise level, σ , to demonstrate the deviation from the ideal ratio of 1. We then determined for all noise-estimation methods the minimal noise level, $\sigma_{5\%}$, for which the absolute value of the relative deviation $|(\sigma_{\text{est}}/\sigma) - 1|$ is lower than 5 % for all $\sigma > \sigma_{5\%}$. I.e., the considered noise-estimation method is accurate within a 5-% range for all $\sigma > \sigma_{5\%}$. We compared these results with the theoretical minimum, σ_{min} , that can be estimated from a sample of size $N = 25\,000$ of discrete image data according to Eqs. [14] to [18]. Finally, we averaged the deviation $|(\sigma_{\text{est}}/\sigma) - 1|$ over the interval between σ_{min} and 2.0 in order to compare the mean deviations obtained by the different noise-estimation methods.

In a second series of simulations, we analyzed the dependence of the noise estimation on the sample size, N , (i.e. the number of pixels in the considered region or in the pixel time-course) by

varying N from 16 to 25 600. These simulations were performed with a fixed original noise level of $\sigma = 0.6$ and were repeated 100 times in order to calculate the mean value and the standard deviation of the ratio $\sigma_{\text{est}}/\sigma$. In addition to the maximum-likelihood approach, we used also the least-squares fit (Eq. [6]) in these simulations in order to compare both techniques.

These two series of simulations were performed for the two most frequently applied quantization techniques: simple truncation of the fractional part and symmetric arithmetic rounding.

In all performed simulations, the noise was estimated using the conventional approaches and the proposed maximum-likelihood technique. To analyze the noise-estimation approaches based on background data of single-channel (experiment a) and RSS 16-channel measurements (experiment b), we generated synthetic complex background data with Gaussian noise and mean value $\mu = 0$. We then simulated a single-channel magnitude and a 16-channel RSS reconstruction followed by the quantization to integer values. We calculated the mean value, m , and standard deviation, s , of the quantized data and estimated the originally applied noise level, σ_{est} , using

$$\begin{aligned}\sigma_{\text{est}}(m) &= \sqrt{2/\pi} m \quad \text{and} \\ \sigma_{\text{est}}(s) &= \sqrt{2/(4-\pi)} s\end{aligned}\quad [19]$$

for single-channel data (2) and

$$\begin{aligned}\sigma_{\text{est}}(m) &= m/\beta(n) \quad \text{and} \\ \sigma_{\text{est}}(s) &= s/\sqrt{2n - \beta(n)^2}\end{aligned}\quad [20]$$

with $\beta(n) = \sqrt{\pi/2} (1 \cdot 3 \cdot 5 \cdots (2n-1)) / (2^{n-1} (n-1)!)$ and $n = 16$ for 16-channel RSS data (9). These results were compared with the noise estimations based on the maximum-likelihood approach proposed in Eq. [8] together with Eqs. [10] and [11].

To analyze the pixelwise determination of the noise level (experiment c), the signal time-course of a single pixel in N repeated single-channel acquisitions was simulated. In this case, the results of the noise estimations depend on the signal mean value, μ . Thus, we generated 20 data sets with signal mean values (over the time-course) varying linearly between $\mu = 100$ in data set #1 and $\mu = 100.95$ in data set #20. The SNR of the pixel time-courses varied between about 50 (for $\sigma = 2.0$) and 5050 (for $\sigma = 0.02$). After quantization, the originally applied

noise level was estimated from the standard deviation, s , of the N repetitions as $\sigma_{\text{est}}(s) = s$ and from the counts of signal intensities applying the maximum-likelihood approach based on Eqs. [8] and [12] with a two-parameter fit for σ and μ . Finally, the results of all 20 data sets (with varying μ) were averaged.

To analyze the noise-level measurements in subtraction data (experiment d), we simulated pairs of single-channel magnitude data sets with a foreground signal (before adding the noise) of $100+R$, where R was randomly and uniformly distributed between 0 and 1, but identical in both data sets to be subtracted, resulting again in SNRs between about 50 and 5050. The data sets were first quantized to integer values and then subtracted. We calculated the standard deviation, s , of the subtraction data and estimated the originally applied noise level as $\sigma_{\text{est}}(s) = \sqrt{1/2} s$. This result was compared with the noise level estimated from the counts of pixel intensities in the subtraction data based on Eqs. [8] and [13].

Finally, we evaluated an image example taken from our MR image archive: we determined the noise level in the background of a T_1 -weighted MR mammography image (quantized by truncation) acquired with a 2-channel mammography coil and a routine protocol (Magnetom Symphony, Siemens Medical Sol., Erlangen, Germany) using the conventional and the proposed maximum-likelihood approach.

Results

The deviations of the noise level, $\sigma_{\text{est}}/\sigma$, plotted over the original noise level in the range between 0.02 and 2 are shown in Figs. 1 and 2. Generally, all noise estimations approach the true value with increasing noise level, while larger deviations between σ_{est} and σ occur for smaller noise levels. The deviations of the conventionally estimated noise levels are substantially larger than those of the maximum-likelihood estimations, for which the ratio, $\sigma_{\text{est}}/\sigma$, remains very close to 1 for all but very small values of σ . This behavior is illustrated by the values of $\sigma_{5\%}$ in Tables 1 and 2, i.e., by the lower limits for acceptable noise estimations with the different evaluated methods: after truncation (Ta-


ble 1), $\sigma_{5\%}$ varies between 0.89 and 2.00 for conventional noise measurements, and between 0.02 and 0.23 for noise estimations with the maximum-likelihood approach. After symmetric rounding (Table 2), $\sigma_{5\%}$ varies between 0.26 and 1.41 for conventional noise measurements, and between 0.02 and 0.24 for noise estimations with the maximum-likelihood approach. We found a very good agreement between $\sigma_{5\%}$ determined with the maximum-likelihood approach and the theoretical lower limit of noise estimations, σ_{min} . The corresponding deviations, $|\left(\sigma_{\text{est}}/\sigma\right) - 1|$, averaged over the interval between σ_{min} and 2.0, vary between 3.39 % and 45.45 % for conventionally estimated noise levels and between 0.08 % and 0.58 % for noise estimations using the maximum-likelihood approach (Tables 1 and 2).

Particularly large deviations even for relatively large noise levels greater than 1.0 are observed for the noise levels estimated from the mean value of background noise in the single-channel data after truncation (Fig. 1a, dash-dotted line). In 16-channel RSS data (Figs. 1b and 2b), the deviations show characteristic oscillations for small original noise levels, $\sigma < 0.6$. The most obvious difference between noise estimations after truncation and after symmetric rounding is the considerably better noise estimation based on the mean value of background noise if symmetric rounding was used (Figs. 2a and 2b, dash-dotted line).


The dependence of the noise estimations on the sample size (i.e. the size of the used ROI or the number of repeated acquisitions) is shown in Fig. 3 for data quantization by truncation. The mean values of the noise estimations remain approximately constant even for relatively small samples, while the associated standard deviations increase with decreasing sample size. The standard deviation of noise estimations based on single-channel background data is similar for all four noise-estimation methods (Fig. 3a). For very small sample sizes below about 100, a certain systematic deviation was found for the results of the maximum-likelihood approach. Based on 16-channel RSS data, a substantially larger standard deviation is observed for noise estimations using the standard deviation of background data than for those using the alternative approaches (Fig. 3b). For noise measurements based on difference images (Fig. 3d), the standard

deviation of the least-squares approach is slightly larger (by a factor of about 1.5) than the standard deviation of the conventional and the maximum-likelihood approaches. Similar results as just described were obtained for the dependence of the noise estimations on the sample size after quantization by arithmetic rounding.


The application in original MRI data is demonstrated in Fig. 4. With the maximum-likelihood approach, the determined noise level is 0.67 in contrast to 0.41 and 0.86 derived conventionally from the mean value and standard deviation, respectively; this corresponds to deviations of the conventionally determined noise levels of about 39 % and 28 %.



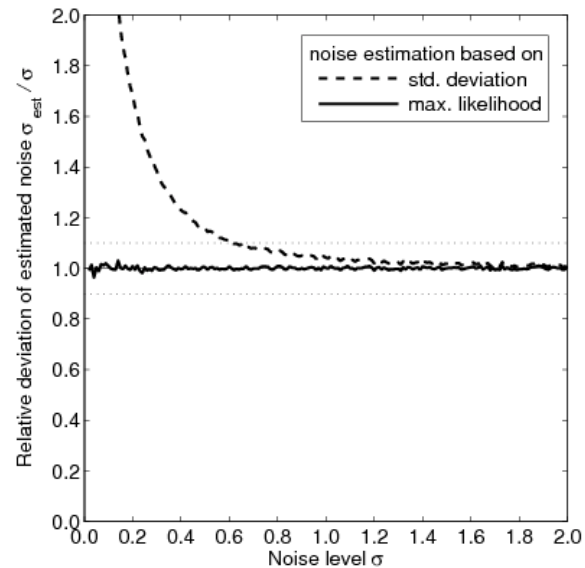
(a) single-channel background data



(b) 16-channel background data



(c) pixelwise evaluated foreground data



(d) subtracted foreground data

Figure 1: Relative deviation of the estimated noise, σ_{est}/σ , in noise measurements after data quantization by truncation; sample size $N = 25\,000$. The ratio, σ_{est}/σ , is plotted over the true noise level, σ , for four different noise-measurement experiments (a–d) and different noise-estimation techniques. The solid line represents the proposed maximum-likelihood approach, while the dashed and dash-dotted lines represent the conventional approaches based on the standard deviation or the mean value of the analyzed data sample, respectively.

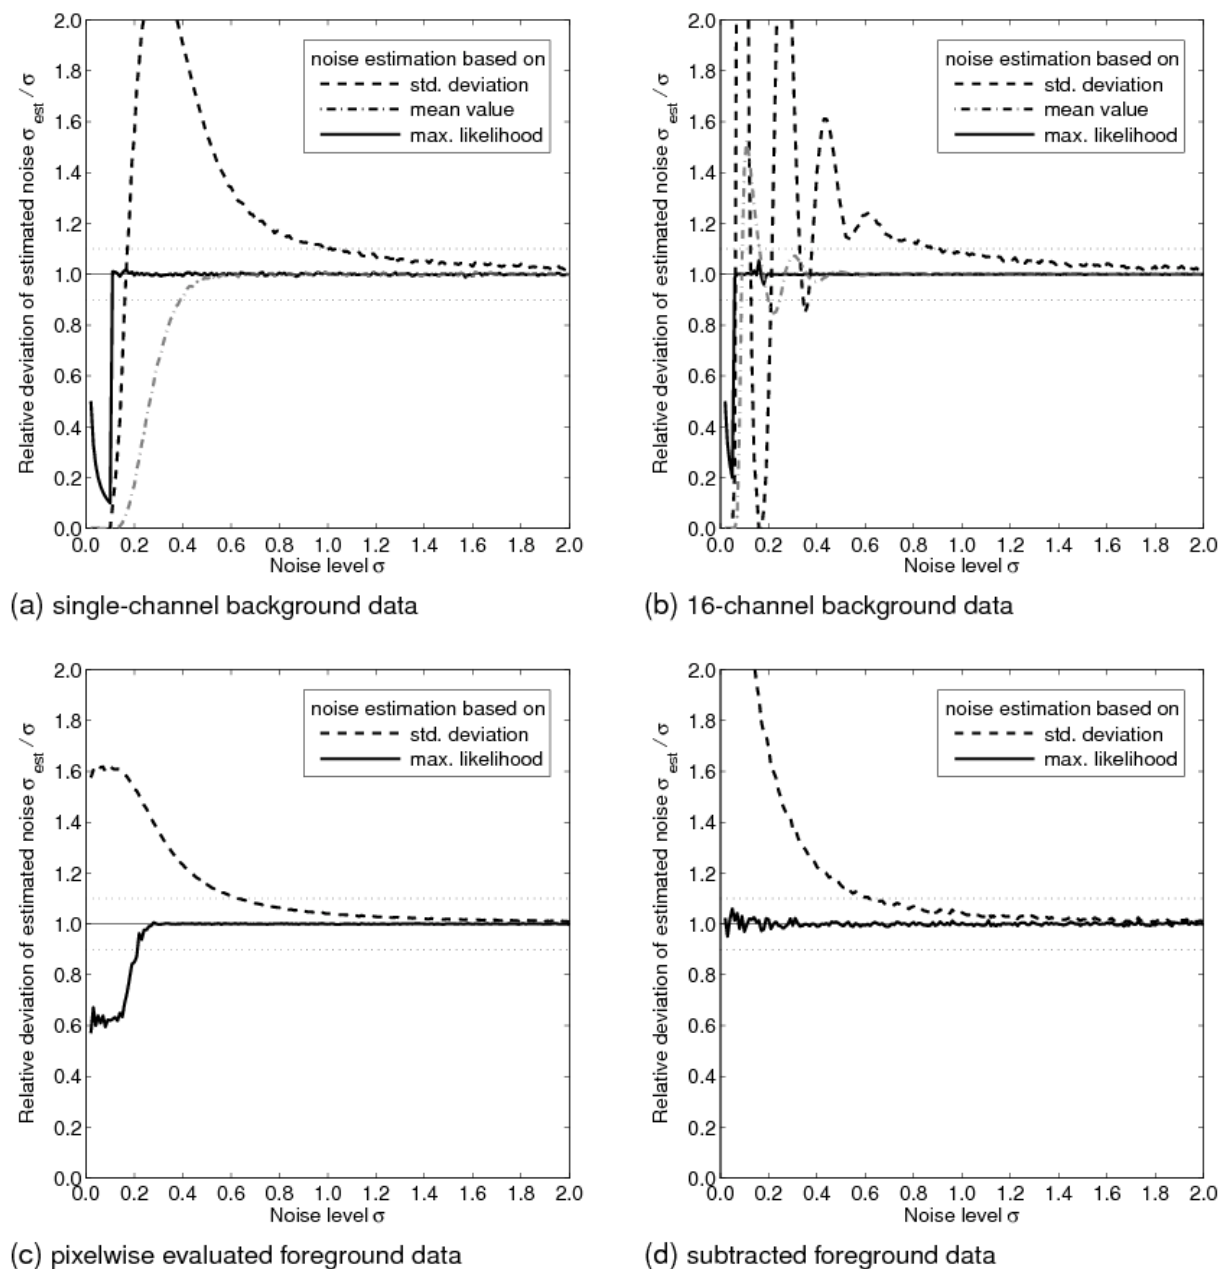


Figure 2: Relative deviation of the estimated noise, $\sigma_{\text{est}}/\sigma$, in noise measurements after data quantization by symmetric arithmetic rounding; sample size $N = 25\,000$. The ratio, $\sigma_{\text{est}}/\sigma$, is plotted over the true noise level, σ , for four different noise-measurement experiments (a–d) and different noise-estimation techniques. The solid line represents the proposed maximum-likelihood approach, while the dashed and dash-dotted lines represent the conventional approaches based on the standard deviation or the mean value of the analyzed data sample, respectively.

Table 1: Noise measurements after quantization by truncation; sample size $N = 25\,000$


	σ_{\min}	Noise estimation based on					
		standard deviation		mean value		max. likelihood	
		rel. dev. (%) ^a	$\sigma_{5\%}$	rel. dev. (%) ^a	$\sigma_{5\%}$	rel. dev. (%) ^a	$\sigma_{5\%}$
Single-channel data, background	0.222	12.05 (18.15)	1.26	45.45 (24.81)	2.00	0.313 (0.317)	0.21
RSS 16-channel data, background	0.117	21.90 (44.92)	1.24	13.95 (16.02)	1.77	0.083 (0.086)	0.11
Multiple-acquisition data, foreground	0.243	7.44 (9.56)	0.89	–	–	0.124 (0.245)	0.23
Subtraction data, foreground	0	23.26 (54.92)	0.89	–	–	0.486 (0.439)	0.02

^a Given are the mean value and the standard deviation in parentheses of the relative deviation, $|(\sigma_{est}/\sigma) - 1|$, averaged over the interval $\sigma_{\min} \leq \sigma \leq 2.0$ and expressed in percent.

Table 2: Noise measurements after quantization by symmetric arithmetic rounding; sample size $N = 25\,000$

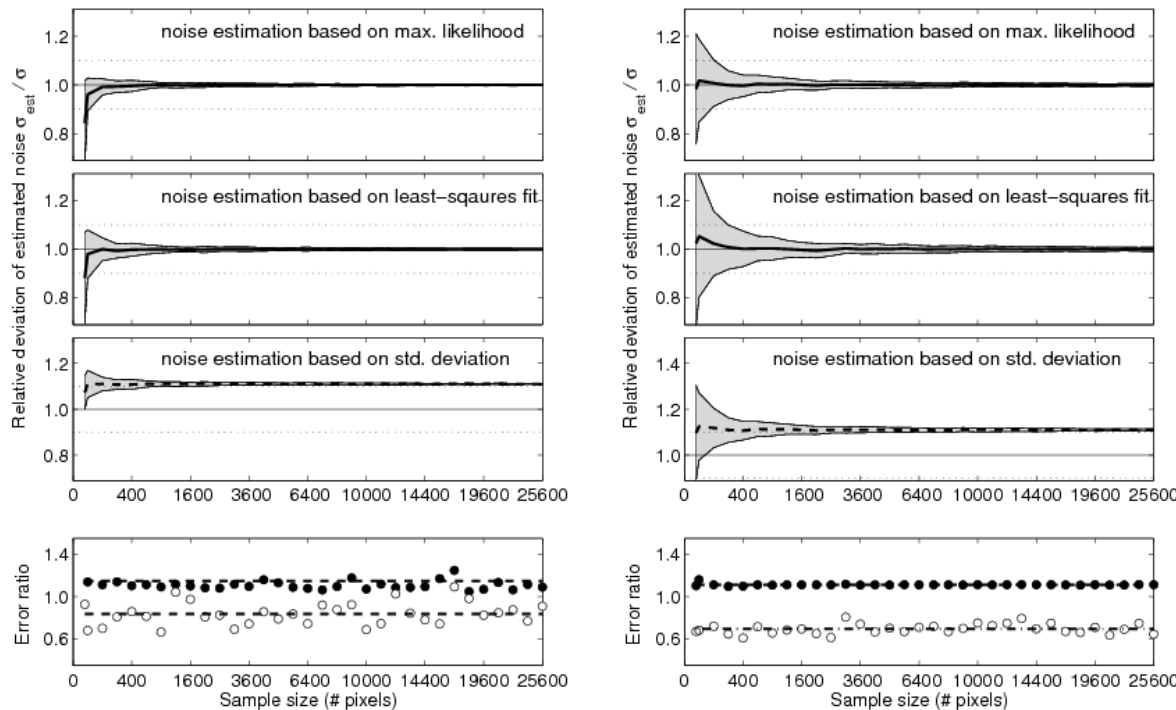
	σ_{\min}	Noise estimation based on					
		standard deviation		mean value		max. likelihood	
		rel. dev. (%) ^a	$\sigma_{5\%}$	rel. dev. (%) ^a	$\sigma_{5\%}$	rel. dev. (%) ^a	$\sigma_{5\%}$
Single-channel data, background	0.111	25.53 (33.59)	1.41	8.80 (23.22)	0.42	0.338 (0.288)	0.11
RSS 16-channel data, background	0.0583	29.35 (77.92)	1.30	3.39 (12.65)	0.26	0.100 (0.138)	0.05
Multiple-acquisition data, foreground	0.243	7.45 (9.57)	0.90	–	–	0.100 (0.097)	0.24
Subtraction data, foreground	0	23.13 (54.28)	0.88	–	–	0.581 (0.727)	0.02

^a Given are the mean value and the standard deviation in parentheses of the relative deviation, $|(\sigma_{est}/\sigma) - 1|$, averaged over the interval $\sigma_{\min} \leq \sigma \leq 2.0$ and expressed in percent.



(a) single-channel background data

(b) 16-channel background data



(c) pixelwise evaluated foreground data

(d) subtracted foreground data

Figure 3: Accuracy and precision of noise measurements after data quantization by truncation; original noise level $\sigma = 0.6$. The mean relative deviation, σ_{est}/σ , and the corresponding standard deviation (gray area) are plotted over the sample size, N , for four different noise-measurement experiments (a–d) and different noise-estimation techniques. In the bottom plot of each experiment, the ratio of the standard deviations of the maximum-likelihood approach and the alternative approaches is shown (open circles: least-squares fit, black filled circles: standard deviation, gray filled circles: mean value).




Figure 4: Image example, demonstrating the occurrence of a very low noise level in clinical routine MRI. (a) T_1 -weighted MR mammography, (b) magnified view of the 64×64-pixel background region shown in (a), a very narrow intensity window is used to visualize the noise distribution (black pixels: intensity 0, gray pixels: intensity 1, white pixels: intensity 2 and 3). (c) The relative frequencies of the signal intensities in (b); the filled circles indicate the relative frequencies of the fitted distribution.

Discussion

The results demonstrate that the conventional approaches for noise measurements should not be applied in quantized image data with noise levels that are comparable to or lower than the quantization errors, i.e., if the noise level, σ , is approximately equal to or lower than 1. Noise estimations based on the standard deviation of the noisy foreground or background data tend to overestimate the true noise level, σ , while noise estimations based on the mean value of background data rather underestimate σ . The accuracy of the noise estimation can be substantially improved with a maximum-likelihood approach, which exploits the counted frequencies of the quantized signal intensities and the analytically determined distributions derived in Eqs. [10] – [13]. This is important for many typical applications of SNR measurements such as the comparison of different pulse sequences or protocols for a certain diagnostic question. While, e.g. in technical comparisons of different RF coils, sequence parameters could easily be changed in order to avoid very low signal or noise levels, this option is generally not available for a clinical protocol that has already been optimized in terms of contrast or spatial resolution. It is also noteworthy that particularly high-SNR applications are influenced by the described noise quantization effects. Scaling factors are generally chosen depending on the maximum signal intensity in an image; thus, scaling the originally high image in-

tensity to intermediate values will, at the same time, decrease the noise level to rather low values.

Even the improved approach for noise measurements, however, requires a certain minimum noise level. If the original noise level is too small ($\sigma < \sigma_{\min}$) then the noise level cannot longer be estimated accurately from the existing data. This is illustrated by the remaining deviations of the noise estimation using the maximum-likelihood approach for noise levels between 0 and about 0.2 in Figs. 1 and 2.

We used truncation and symmetric arithmetic rounding in this study to quantize the floating-point data to integer values. These approaches are by far the most frequently applied quantization techniques. However, the proposed technique for noise measurements can also be applied for other quantization techniques such as, e.g., the ceil-function; the only difference to our presented work is that the integration boundaries, b_k , in Eqs. [10] – [12] must be chosen differently. Comparing the results of truncation and symmetric rounding, a relatively similar behavior is generally observed with respect to the deviations of the estimated noise levels. The most obvious differences are found in the evaluation of background data, for which the estimations based on the signal mean value are much more accurate after rounding than after truncation. This is explained by the systematic bias (an offset of -0.5) of the mean value of truncated floating-point data distributed around a quantization boundary.

The evaluation of the standard deviations of the estimated noise levels in dependence on the sample size demonstrates that the precision of the proposed noise-estimation technique is generally comparable to the conventional techniques (which, however, provide a substantially lower accuracy as indicated by the deviations of the mean value from 1.0 in Fig. 3). The standard deviations of all evaluated approaches decrease with the expected dependence on the square root of the sample size, $1/\sqrt{N}$. The evaluation of samples with 400 to 1000 data points (corresponding to relatively small regions with about 20 to 35 pixels diameter) should be sufficient to estimate the noise level with an error below 5 %. The maximum-likelihood and least-squares approaches gave very similar results in most cases. In some of our simulated experiments, the precision of the maximum-likelihood approach was better (by about 50 %) than the one of the least-squares fit. On the other hand, the maximum-likelihood technique resulted in slightly biased (increased) noise estimations at very small sample sizes, e.g. for the single-channel background experiment (Fig. 3a).

Although somewhat surprising, we found a certain number of datasets in our clinical image archive with very low quantized signal intensities, i.e. data sets in which the standard deviation of noise becomes comparable to the magnitude of quantization errors. The image shown in Fig. 4 is representative for several data sets acquired in different patients with the identical protocol and RF coil configuration, all of which exhibit similarly low noise levels. If possible, such low signal intensities should be avoided, but typically the scaling factors used for the image reconstruction cannot easily be modified by the user. Of course, it would be generally desirable that the image reconstruction system checks the absolute signal intensities before quantization and applies appropriate scaling factors if necessary.

Generally, the effects of signal quantization should be considered not only for noise measurements, but also for other image post-processing steps or data evaluation procedures. As long as filters are applied by the scanner itself during image reconstruction, they should not be affected by quantization effects. Signal quantization is typically the last step of image reconstruction, and all filters

such as rawdata-based noise-reduction filters or intensity-normalization filters, which compensate for coil-sensitivity or B_1 effects, are applied *before* signal quantization. If, however, image data are post-processed after archiving in an integer format, then quantization effects may influence the obtained results. This could be the case e.g. in fMRI analyses of the time-course of a single pixel, for which statistical properties such as the autocorrelation of the signal-time curve may be significantly influenced by data quantization.

Conclusions

At very low noise levels, the quantization of floating-point image data can bias the measurement of image noise and, consequently, the accurate determination of SNR or CNR in MR images. This effect is observed when using conventional techniques for noise estimation based on the assumption of a continuous distribution of signal intensities. We extended and evaluated a recently suggested maximum-likelihood approach for noise measurements that takes the discrete distribution of signal intensities into account. Our results demonstrate that noise levels that are comparable to or lower than the quantization errors can be accurately estimated with the proposed approach.

Appendix: Signal frequencies in subtracted integer image data

Consider two real-valued image data sets of identical acquisitions and a foreground region with sufficiently high SNR. Then the signal statistics of each acquisition can be approximated by a Gaussian distribution with standard deviation, σ , and the statistics of the difference, x , of the signal intensities (for each pixel) is also described by a Gaussian distribution with a standard deviation increased by a factor of $\sqrt{2}$ (due to the subtraction operation). If we calculated the quantized signal intensities *after* subtraction, then it would be described by $F_{\text{Gauss}}(k; 0, \sqrt{2} \sigma)$, i.e. by Eq. [12].

However, quite often image data is first quantized and subtracted only afterwards (e.g., during retrospective data analysis based on archived data), resulting in very different signal statistics. In this

case, the condition for an integer pixel intensity, k , in the final subtraction is that there are exactly k quantization boundaries between both original floating-point values. This is possible for signal differences, x , that lie between $(k - 1)$ and $(k + 1)$ as visualized in Fig. 5. The probability density for the difference x itself is $P_{\text{Gauss}}(x; 0, \sqrt{2}\sigma)$. If we now employ the assumption that the original signal intensities, μ , in the evaluated sample (region) are approximately uniformly distributed over a range greater than the quantization interval, then we can obtain a second expression that describes the probability to find exactly k boundaries (with unit distance) for a given difference interval, x . This expression is directly connected to the geometrical relation between the interval size and the difference between x and k as illustrated in Fig. 5, and is given by: $1 - |x - k|$, i.e. $1 - (k - x)$ for $(k-1) < x < k$ and $1 - (x - k)$ for $k < x < (k+1)$. Thus, the overall probability (i.e. the relative frequency) for the integer pixel intensity, k , in the final subtraction data set is

$$F_{\text{Subt}}(k; \sigma) = \int_{k-1}^k P_{\text{Gauss}}(x; 0, \sqrt{2}\sigma)(1 - k + x)dx + \int_k^{k+1} P_{\text{Gauss}}(x; 0, \sqrt{2}\sigma)(1 - x + k)dx \quad [21]$$

for $k = 0, \pm 1, \pm 2, \pm 3, \dots$. By introducing the term that describes the probability to find k quantization boundaries, the result is already implicitly averaged over all possible original mean values, μ , and, thus, does not longer depend on μ . It is also interesting to note that the expression for $F_{\text{Subt}}(k; \sigma)$ is inde-

pendent of the chosen quantization procedure, i.e. it is the same for symmetric rounding or truncation; only the distance of quantization boundaries must be 1. By rearrangement of the terms and subsequent integration, Eq. [21] can be expressed as:

$$F_{\text{Subt}}(k; \sigma) = \int_{k-1}^k x P_{\text{Gauss}}(x; 0, \sqrt{2}\sigma) dx - \int_k^{k+1} x P_{\text{Gauss}}(x; 0, \sqrt{2}\sigma) dx + (1 - k) \int_{k-1}^k P_{\text{Gauss}}(x; 0, \sqrt{2}\sigma) dx + (1 + k) \int_k^{k+1} P_{\text{Gauss}}(x; 0, \sqrt{2}\sigma) dx \quad [22]$$

$$= \frac{-\sigma}{\sqrt{\pi}} \exp\left(-\frac{x^2}{4\sigma^2}\right) \Big|_{k-1}^k + \frac{\sigma}{\sqrt{\pi}} \exp\left(-\frac{x^2}{4\sigma^2}\right) \Big|_k^{k+1} + \frac{(1-k)}{2} \text{erf}\left(\frac{x}{2\sigma}\right) \Big|_{k-1}^k + \frac{(1+k)}{2} \text{erf}\left(\frac{x}{2\sigma}\right) \Big|_k^{k+1}$$

The same expression can be derived alternatively based on the relative frequencies of the quantized Gaussian distributions, $F_{\text{Gauss}}(k; \mu, \sigma)$, prior to the subtraction and on the probability distribution of the mean value, $P(\mu)$, by a direct (but lengthy) transformation of:

$$F_{\text{Subt}}(k; \sigma) = \int P(\mu) \left(\sum_{k'=-\infty}^{\infty} F_{\text{Gauss}}(k'; \mu, \sigma) F_{\text{Gauss}}(k'-k; \mu, \sigma) \right) d\mu \quad [23]$$




Figure 5: Calculation of the relative frequencies of signal intensities in difference images that were subtracted after quantization. To obtain a final integer pixel intensity of, e.g., $k = 5$, exactly 5 quantization boundaries (numbered from 1 to 5 in this example) are required between the two original floating-point values, which are assumed to lie in the light-gray-shaded intervals. The difference, x , (before quantization) between these two values can vary between 4 and 6 as indicated by the black horizontal brackets. Obviously, the probability to find exactly 5 boundaries between the two values is close to zero if x is approximately 4 or 6, since most configurations with $x \approx 4$ or $x \approx 6$ will spread over either only 4 or over 6 quantization boundaries as shown by the gray brackets. The probability increases to one if the distance, x , approaches 5. The probability to find exactly 5 boundaries for a given distance, x , is visualized by the ratio of the number of black and gray brackets in the example, and is expressed by $1 - |x - 5|$.

References

1. Edelstein WA, Bottomley PA, Pfeifer LM. A signal-to-noise calibration procedure for NMR imaging systems. *Med Phys* 1984;11:180–185.
2. Henkelman RM. Measurement of signal intensities in the presence of noise in MR images. *Med Phys* 1985;12:232–233.
3. Holden JE, Halama JR, Hasegawa BH. The propagation of stochastic pixel noise into magnitude and phase values in the Fourier analysis of digital images. *Phys Med Biol* 1986;31:383–396.
4. Bernstein MA, Thomasson DM, Perman WH. Improved detectability in low signal-to-noise ratio magnetic resonance images by means of a phase-corrected real reconstruction. *Med Phys* 1989;16:813–817.
5. Miller AJ, Joseph PM. The use of power images to perform quantitative analysis on low SNR MR images. *Magn Reson Imaging* 1993;11:1051–1056.
6. McGibney G, Smith MR. An unbiased signal-to-noise ratio measure for magnetic resonance images. *Med Phys* 1993;20:1077–1078.
7. Gudbjartsson H, Patz S. The Rician distribution of noisy MRI data. *Magn Reson Med* 1995;34:910–914. (Erratum in: *Magn Reson Med* 1996;36:332.)
8. Sijbers J, Poot D, den Dekker AJ, Pintjens W. Automatic estimation of the noise variance from the histogram of a magnetic resonance image. *Phys Med Biol* 2007;52:1335–1348.
9. Constantinides CD, Atalar E, McVeigh ER. Signal-to-noise measurements in magnitude images from NMR phased arrays. *Magn Reson Med* 1997;38:852–857. (Erratum in: *Magn Reson Med* 2004;52:219.)
10. Kellman P, McVeigh ER. Image reconstruction in SNR units: a general method for SNR measurement. *Magn Reson Med* 2005;54:1439–1447.
11. Dietrich O, Raya JG, Reeder SB, Ingrisch M, Reiser MF, Schoenberg SO. Influence of multi-channel combination, parallel imaging, and other reconstruction techniques on MRI noise characteristics. *Magn Reson Imaging* 2008;26:754–762.
12. Roemer PB, Edelstein WA, Hayes CE, Souza SP, Mueller OM. The NMR phased array. *Magn Reson Med* 1990;16:192–225.
13. Larsson EG, Erdogmus D, Yan R, Principe JC, Fitzsimmons JR. SNR-optimality of sum-of-squares reconstruction for phased-array magnetic resonance imaging. *J Magn Reson* 2003;163:121–123.
14. Kaufman L, Kramer DM, Crooks LE, Ortendahl DA. Measuring signal-to-noise ratios in MR imaging. *Radiology* 1989;173:265–267.
15. Dietrich O, Raya JG, Reeder SB, Reiser MF, Schoenberg SO. Measurement of signal-to-noise ratios in MR images: Influence of multichannel coils, parallel imaging, and reconstruction filters. *J Magn Reson Imaging* 2007;26:375–385.
16. Murphy BW, Carson PL, Ellis JH, Zhang YT, Hyde RJ, Chenevert TL. Signal-to-noise measures for magnetic resonance imagers. *Magn Reson Imaging* 1993;11:425–428.
17. Reeder SB, Wintersperger BJ, Dietrich O, Lanz T, Greiser A, Reiser MF, Glazer GM, Schoenberg SO. Practical approaches to the evaluation of signal-to-noise ratio performance with parallel imaging: application with cardiac imaging and a 32-channel cardiac coil. *Magn Reson Med* 2005;54:748–754.
18. National Electrical Manufacturers Association (NEMA). Digital Imaging and Communication in Medicine (DICOM). Rosslyn: National Electrical Manufacturers Association (NEMA); 2007.