Lehrstuhl für Bildverarbeitung
Institute of Imaging & Computer Vision

Noise-Resistant Weak-Structure Enhancement for Digital Radiography

Martin Stahl and Til Aach and Sabine Dippel and Thorsten Buzug and Raphael Wiemker and Ulrich Neitzel

in: SPIE Vol. 3661: Medical Imaging 99: Image Processing. See also BIBTEX entry below.

BIBTEX:

```
@inproceedings{STA99a,
author = {Martin Stahl and Til Aach and Sabine Dippel
   and Thorsten Buzug and Raphael Wiemker and
   Ulrich Neitzel},
title = {Noise-Resistant Weak-Structure Enhancement for
   Digital Radiography},
booktitle = {{SPIE Vol. 3661: Medical Imaging 99: Image Processing}},
editor = {K. M. Hanson},
publisher = {SPIE},
address = {San Diego, USA},
month = {February 20--26},
year = {1999},
pages = {1406--1417}}
```


document created on: December 19, 2006
created from file: spiemi99coverpage.tex
cover page automatically created with CoverPage.sty
(available at your favourite CTAN mirror)

RWTH AACHEN
RHEINISCH-WESTFÄLISCHE TECHNISCHE HOCHSCHULE AACHEN

PROCEEDINGS OF SPIE
SPIE—The International Society for Optical Engineering

Medical Imaging 1999

Image Processing

Kenneth M. Hanson
Chair/Editor

22–25 February 1999
San Diego, California

Sponsored by
SPIE—The International Society for Optical Engineering

Cooperating Organizations
AAPM—American Association of Physicists in Medicine
APS—American Physiological Society
FDA Center for Devices and Radiological Health
IS&T—The Society for Imaging Science and Technology
NEMA—National Electrical Manufacturers Association/Diagnostic Imaging and Therapy Systems Division
RSNA—Radiological Society of North America
SCAR—Society for Computer Applications in Radiology

Published by
SPIE—The International Society for Optical Engineering

SPIE
P
PROCEEDINGS SERIES

Volume 3661
Part Two of Two Parts

MI POSTER AWARD
Honorable Mention

Noise-resistant weak-structure enhancement for digital radiography

M. Stahl[a], T. Aach[b], T. M. Buzug[a], S. Dippel[a] and U. Neitzel[c]

[a] Philips Research, Roentgenstrasse 24–26, D-22335 Hamburg, Germany

[b] Medical University of Lübeck, D-23569 Lübeck, Germany

[c] Philips Medical Systems, Roentgenstrasse 24–26, D-22335 Hamburg, Germany

ABSTRACT

Today's digital radiography systems mostly use unsharp masking-like image enhancement techniques based on splitting input images into two or three frequency channels. This method allows to enhance very small structures (edge enhancement) as well as enhancement of global contrast (harmonization). However, structures of medium size are not accessible by such enhancement. We develop and test a nonlinear enhancement algorithm based on hierarchically repeated unsharp masking, resulting in a multiscale architecture allowing consistent access to structures of all sizes. The algorithm is noise-resistant in the sense that it prevents unacceptable noise amplification. Clinical tests performed in the radiology departments of two major German hospitals so far strongly indicate the superior performance and high acceptance of the new processing.

Keywords: digital radiography, image enhancement, multi-resolution techniques, multiscale enhancement, noise resistance

1. INTRODUCTION

In X-radiographs often large differences in X-ray density dominate over the diagnostically relevant detail information. Hard- or softcopy display then faces the dilemma of having to reproduce subtle details with sufficient contrast and at the same time large variations without clipping. In digital radiography, these conflicting requirements can be reconciled by suitable processing.[1–6] The large density variations occur predominantly over low frequencies and diagnostically relevant information is mainly high frequent. Therefore, most present systems separate the image into two (or three) frequency bands allowing selective enhancement of the high frequent information. Receiver operating characteristic (ROC) studies show that detail perception can indeed be improved by appropriate processing.[2]

While these image enhancement methods have been tuned towards considerable performance over the years, improved image decomposition techniques, like wavelets[7] and pyramids[8] have emerged. These *multiscale* techniques decompose images into several frequency channels of varying spatial resolution. The advantage of these decomposition techniques is that structures of different size appear separately in different scales, and can be processed independently, as done for television images in.[9] However, when applied to (digital) radiographs (see e.g.[10–12]), it turns out that these approaches are as sensitive to noise as unsharp masking-based algorithms. The reason for this is that the finest scale is more or less equivalent to the high-frequency channel of the described standard processing, where noise is most critical.

In this paper we therefore focus on a noise-resistant multiscale enhancement algorithm. We show that multiscale processing can be regarded as hierarchically repeated unsharp masking, and describe how existing unsharp masking-based algorithms can be mapped on a multiscale structure. We then exploit the additional opportunities of this architecture for contrast enhancement, and describe how to prevent uncontrolled noise amplification. Finally, results of a first clinical study are presented.

Further author information: (Send correspondence to M. Stahl)
T. M. Buzug: Present address: RheinAhrCampus Remagen, University of Applied Sciences, Suedalle 2, 53424 Remagen, Germany

2. FROM UNSHARP MASKING TO MULTISCALE PROCESSING

As a reference standard algorithm to be mapped later on a multiscale pyramid, we use a two-step unsharp masking-based algorithm (Figure 1), which is applicable e.g. to storage phosphor radiographs[1]: First, a slightly blurred version I_{blurred}, generated by filtering with a small kernel, is subtracted from the original I_{org}. The resulting high-frequency image I_{edge} contains mostly edge information, and is enhanced by a constant *sharpness factor* SF_{us}. Adding the enhanced highpass data to I_{blurred} provides an image which exhibits increased perceived sharpness (*Edge Enhancement*). Then, a large box kernel — e.g. 201 × 201 pixel for a $2k \times 2k$-radiograph — is applied to separate middle and high spatial frequencies (contrast information, I_{contrast}) from the very low ones. Subjecting I_{contrast} to another constant factor — termed contrast factor CF_{us} — amplifies contrast information relative to the large and low-frequent density variations I_{density} (*Harmonization*). Note that nonlinear amplification like[4] is easily introduced into this structure by choosing density- or contrast-dependent gain factors (see dotted lines in Figure 1). This holds e.g. for the reference standard algorithm that is used in 3.3 for chest images.[13]

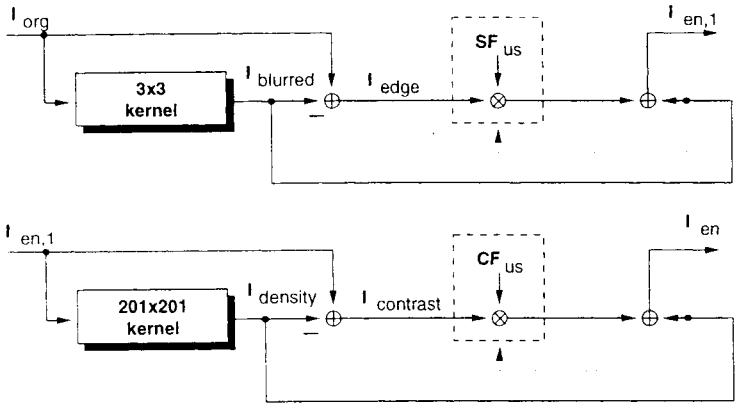

Figure 1. Two step unsharp-masking-based processing

Obviously, this algorithm allows separate access only to fine detail and low-frequent information. Medium-sized structures remain grouped together in the broad frequency band I_{contrast}. The algorithm's performance can hence be considerably improved by accessing structures of arbitrary size in the same way as fine structures. For convenience, we seek to hierarchically continue the Edge Enhancement operation, where, to avoid an unpractical increase in the amount of data, each lowpass image I_{blurred} is subsampled by a factor of two in each direction (↓). Consequently for subtraction from the original, as well as addition in the reconstruction path, an expansion step must be introduced, i.e. zeros are inserted between the available samples (↑) followed by another lowpass.

Ideally, the lowpass preceding the subsampling should cut off at half the Nyquist frequency. For practical reasons, only small binomial filter kernels are used, thus introducing some aliasing into the subsampled images. However, this is not critical, because the alias cancels out during reconstruction. This balance between produced and cancelled alias is disturbed by subband processing but generally not to a severe degree.

The subsampling provides an image $I_{\text{org},S1}$ of half height and half width compared to the original, representing the spatial frequencies up to half the Nyquist frequency of the original image. The image itself, however, contains spatial frequencies up to the original Nyquist frequency, because subsampling has doubled all actually present frequencies. This means that the subsampled image is a sharp but smaller image, where edges correspond to larger structures in the original. Therefore, the same filter as already used in the upper part of Figure 1 can be applied again to process the edges of $I_{\text{org},S1}$. Repeating this nesting of unsharp-masking operations n times leads to the algorithm in Figure 2, which is known as the *Laplacian Pyramid*.[8] In the following, the difference images will be referred to as subband images or subbands. To achieve backward compatibility of this algorithm to our reference processing, we now map the two-step linear filtering operation on the pyramid. If 3 × 3 binomial filter kernels are used at each level of the pyramid, the finest, uppermost level of the pyramid S_0 is almost equivalent to the edge image I_{edge} of Figure 1. The number of pyramid levels needed is determined by the requirement that the (expanded) lowpass level generated

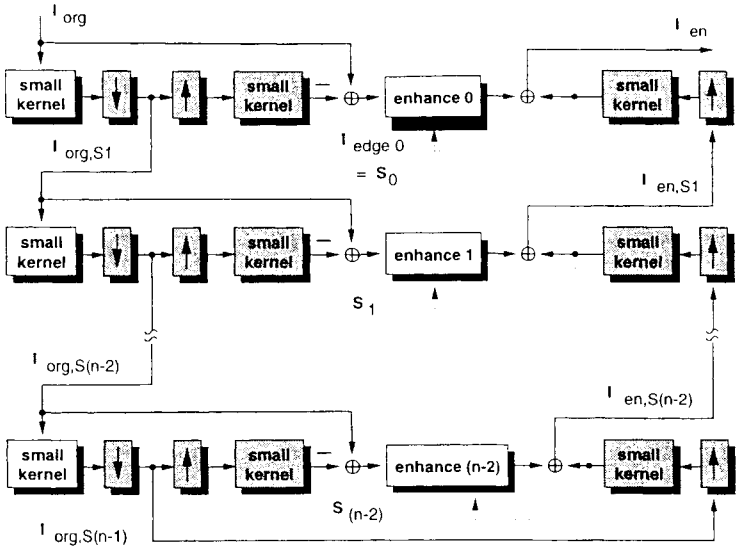

Figure 2. The Laplacian Pyramid with subband processing modules

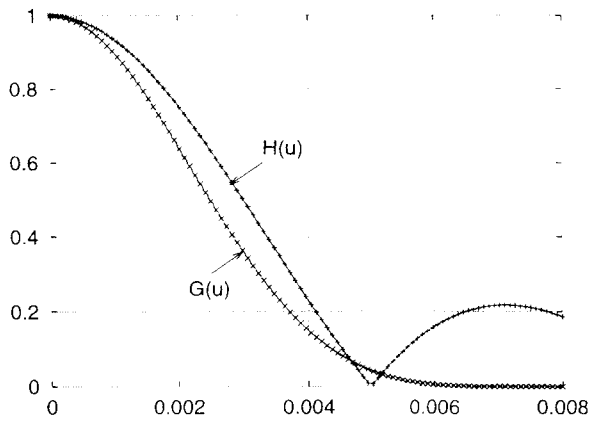

Figure 3. Comparison of the MTF of a one dimensional box kernel $H(u)$ to a pyramid transfer function $G(u)$ considering only the lowpass level. The length of the box kernel is $N = 201$ pixels. The pyramid consists of 8 levels, each using a binomial kernel with a length of 3 pixels.

within the pyramid is approximately equivalent to I_{density} in Figure 1. This is met if the magnitude transfer function (MTF) of the pyramid, considering only the lowpass level, is a good approximation of the MTF of the large box kernel. The MTF $H(u)$ of a box kernel filter of size $N \times N$ is given by

$$H(u) = |\frac{1}{N} \cdot \frac{\sin(N \cdot \pi \cdot u)}{\sin(\pi \cdot u)}| \quad . \tag{1}$$

The lowpass filter chosen for the pyramid is a separable binomial filter of size 3×3 pixels, based on the one dimensional kernel [0.25 0.5 0.25]. Its MTF $G_L(u)$ is given by

$$G_L(u) = (\cos(\pi \cdot u))^2 \quad . \tag{2}$$

Applying this kernel recursively in an n-level pyramid and including the effects of spectral spreading caused by

subsampling yields the transfer function for the lowpass level of the pyramid. For a comparison with the transfer function of the box kernel, it must be taken into account that the reconstruction path of the pyramid applies the same sequence of filters for interpolation. Therefore, the transfer function for the lowpass level of the pyramid must be squared to obtain the transfer function for the expanded lowpass level:

$$G(u) = \left(\prod_{i=0}^{n-2} G_L(2^i \cdot u) \right)^2 \quad , \tag{3}$$

where alias was neglected to simplify the calculation. For $N = 201$, the spectrum of the box kernel can be approximated by a pyramid consisting of 8 levels, as shown in Figure 3. Mapping of the linear enhancement shown in Figure 1 is now easily done by applying the factor $\mathrm{CF}_0 = \mathrm{SF}_{\mathrm{us}} \cdot \mathrm{CF}_{\mathrm{us}}$ to the highpass image and $\mathrm{CF}_i = \mathrm{CF}_{\mathrm{us}}$, $0 < i < 7$, to the other subband images. The lowpass subband image remains unchanged. Unsharp masking algorithms with density or contrast dependent gain factors are mapped similarly.

3. SUBBAND PROCESSING

Clinical acceptance depends decisively on an appropriate design of the subband enhancement and careful tuning of its parameters. In this respect backward compatibility as achieved by the above mapping, which allows re-use of introduced parameters, is an important issue. Moreover we implemented and evaluated a set of new features, described in the following sections.

3.1. Structure Boost

This feature is provided by amplitude-dependent contrast amplification applied to high- and bandpass subband images. The contrast amplification function has been designed according to the following requirements:

- For large contrasts the function should approach the constant value CF_i according to (2).
- For contrast values close to zero the function should approach $CF_i + G_i$, $G_i \geq 0$. G_i defines the additional gain for weak contrasts within subband image i. For $G_i = 0$, $0 \leq i \leq n-2$ Structure Boost is disabled and our algorithm behaves like linear unsharp-masking in Figure 1.

A suitable amplification function is given by

$$CF_i(c) = \begin{cases} G_i \cdot \left(1 - \frac{|c|}{c_0}\right)^{p_i} + CF_i & \text{for } |c| \leq c_0 \\ CF_i & \text{for } |c| > c_0 \end{cases} \quad , \tag{4}$$

The additional enhancement is restricted to weak contrasts between zero and the transition amplitude c_0. The gain G_i allows a smooth transition towards the linear unsharp-masking processing. The exponent p defines how fast the amplification factor decreases towards CF_i. Modification of the exponent does not influence the amplification of contrast values close to zero where also most of the noise can be found. Therefore, the exponent parameter can be used to control the amount of additional weak structure enhancement while being to some extent robust with respect to noise amplification. The dependence of the contrast amplification on gain G is shown in Figure 4. The clinical evaluation showed that modification of G and p is of practical use, whereas for the transition amplitude c_0 a constant value can be chosen, which depends only on the number of quantization levels (typically between 10 and 16 bit) of the radiographs.

The effect of Structure Boost is demonstrated in figure 5, showing a skull image processed with standard processing and multiscale processing. The image was obtained in one of our clinical studies and was taken to examine the state of the patient after a fall. The radiograph exhibits an occipital fracture of the calotte. According to the Radiologists' comments the multiscale processed image provides a superior representation of small structures rendering an almost 3-d like impression of the skull. The visibility of the fracture line has been clearly improved. This holds also for the structure of vessels on the calotte. Enhancement of the highpass subband image was reduced compared to the other subband images in order to avoid noisy or spider net like appearance of the bone structure of the calotte. The latter might feign a pathology.

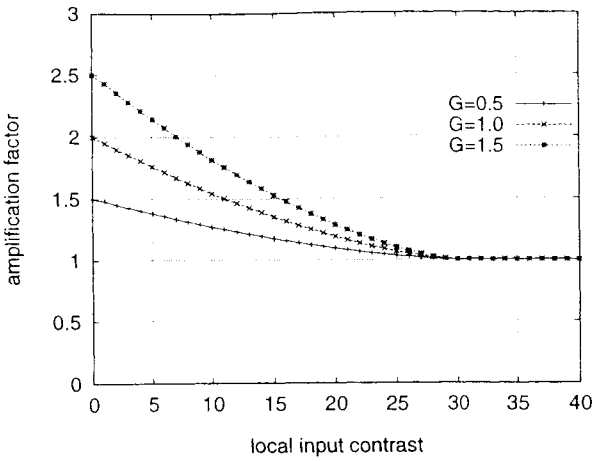

Figure 4. Contrast amplification function according to (4) for $c_0 = 30$, $p = 1.5$, $CF_i = 1.0$, and various values of G.

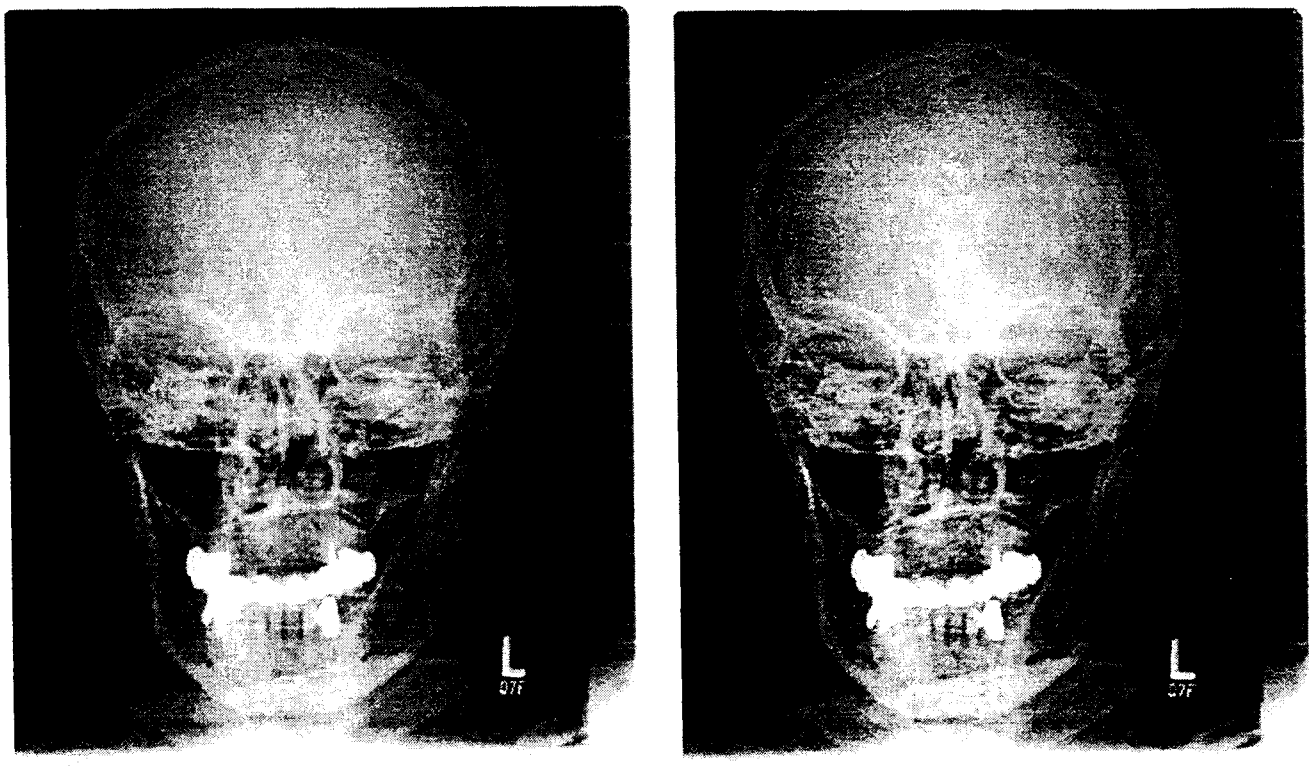

Figure 5. Scull processed with standard unsharp-masking algorithm (left) and multiscale algorithm (right).

3.2. Noise Containment

The Noise Problem. As X-ray images are usually subjected to a logarithmic gain curve during detection, noise power is strongest in areas with low incident X-ray dose, which are displayed as bright areas. Additionally, noise is mainly high frequent. Hence, noise affects the radiograph predominantly in high-frequency bands of low optical density. Subjectively, noise is most strongly visible in areas of weak or no texture-caused intensity fluctuations ("activity"). The 3-d representation in Figure 6 combines this perceptual effect with the physical dependence of

noise, indicating the region where noise is most disturbing. Basically, enhancement of (weak) details also amplifies

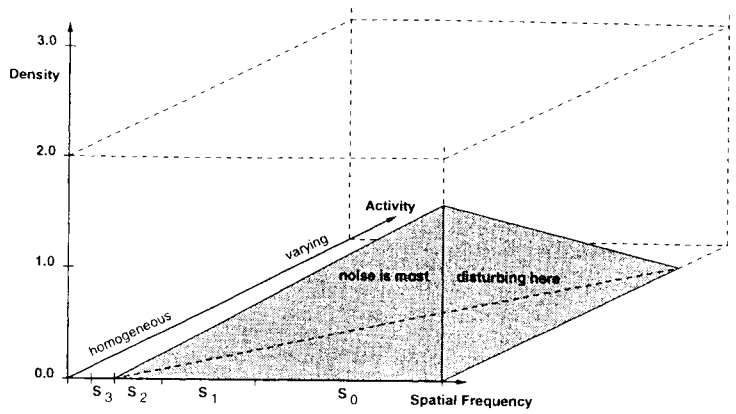

Figure 6. Qualitative plot of subjective noise visibility dependent on signal activity taking into account physical presence of noise. The spatial frequency is subdivided into n octaves according to the subband images S_0, S_1, ..., S_{n-2}, $I_{S(n-1)}$.

noise. Based on the above considerations about noise sensitive regions, our noise containment strategy seeks to keep noise within acceptable limits without losing the benefits of weak structure enhancement.

Realization of Noise Containment. Measures for local density and local activity allow to locate each subband image pixel within the diagram of figure 6. Pixel values mainly affected by noise are located close to the edge given by lowest density, lowest activity and highest spatial frequencies (grey tretrahedreon). The effectiveness of our noise control algorithm is focussed on this most noise sensitive region by integrating a density-, activity-, and scale-dependent control into the enhancement. Only the amount of enhancement is controlled and no information is taken away from the original image. We hence refer to it as *Noise Containment.*

A convenient way of implementation is to first apply the enhancement function to the subband pixel and then weight the difference value between enhanced and original pixel by an attenuation factor $b(x,y)$ that depends on the subband index i (= spatial frequency), local image density and local activity. The attenuation factor $b(x,y)$ is given by the product of density and activity dependent factors $b_A(x,y)$ and $b_D(x,y)$ such that it is smaller than 1 in the noise sensitive region and equal to 1 elsewhere. A block diagram of the algorithm is given in Figure 7. The enhanced subband $S_{i,fully\,enh.}$ and the original subband $S_{i,orig}$ are weighted and added according to

$$S_{i,nr\,enh.}(x,y) = b(x,y) \cdot S_{i,fully\,enh.}(x,y) + (1 - b(x,y)) \cdot S_{i,orig}(x,y) \tag{5}$$

providing a noise robust pixel value $S_{i,nr\,enh.}(x,y)$.

The described method is easy to implement and the concept works regardless of the type (linear, nonlinear, ...) of enhancement. Guidance of Noise Containment by imaging physics as well as perceptual aspects allows a well focussed restriction to the most noise sensitive regions which normally constitute a relatively small part of the image (figure 6). On the rest of the image the full potential of contrast enhancement can be exploited.

The local density measure M_D can easily be retrieved from the expanded and enhanced image of the next lower level of the pyramid $I_{en,S(i+1),exp}$. An example of b_D versus M_D is given in Figure 7. The density-dependent weighting factor b_D is one in dark areas (high optical density), corresponding to full contrast enhancement. Starting at a critical density value $M_{D,c}$, it decreases over the interval ΔM_D towards a value $b_{D,min}$ corresponding to a residual enhancement. This transition, however, depends on the subband: whereas in high-frequency subbands for low densities a reduction of enhancement might be necessary, it will not be needed in lower-frequency subbands, which provide better signal to noise ratios. For b_A versus M_A similar considerations hold. Activity measures can be calculated using, for instance, the standard deviation or local one-dimensional histogram-based measures like

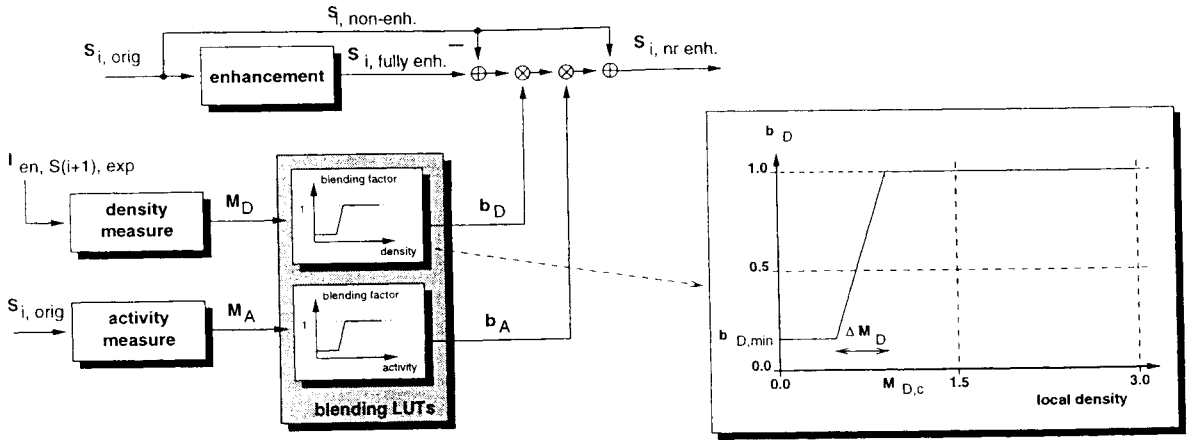

Figure 7. Diagram Φ of noise containment algorithm.

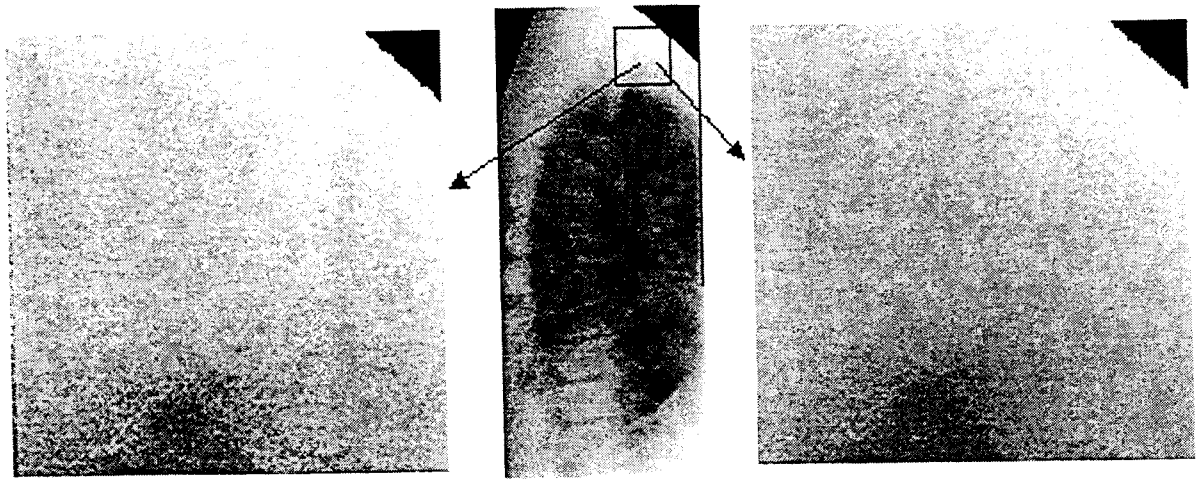

Figure 8. Dorsal spine processed with multiscale algorithm without (left) and with Noise Containment (right).

information (negative entropy) or energy. The lateral radiograph of a dorsal spine demonstrates the effect of Noise Containment (Figure 8). A good representation of all regions of the lateral spine must cope with the large density variations. In the upper boundary area the shoulder-blades cause disturbing superpositions, while in the lower boundary area the large density transition between lung and abdominal cavity must be handled. In both boundary areas, high noise levels degrade the image impression. Unacceptable noise boosting on the finest scales, particularly in the bright regions of the upper and lower parts of this radiograph, could successfully be prevented through Noise Containment. Expert comments included:

- Structure Boost enables better exploitation of the available dynamic range of the film together with excellent representation of weakly contrasting details such as the posterior edges of the bodies of the vertebra.
- Noise Containment provides distinct quality improvement in the lower and upper boundary area of the dorsal spine.
- Noise Containment does not lead to any loss of detail information.

3.3. Density-Controlled Contrast Enhancement

Requirements on contrast enhancement may depend on the region of the image because of e.g. the presence of noise or the risk of feigned pathologies due to over-enhancement. The latter is the case e.g. for chest images. Here an amount of contrast enhancement that could well be applied to the region of the mediastine could be unacceptable for the region of the lung.

Concerning the lung area, chest image processing should provide good contrast of the pulmonary structures, render visible weakly contrasting pathological structures, e.g. lung nodules, and visualize continuity of bone structure. The processing should avoid to over-emphasize the pulmonary vascular system which could feign pneumonia. The retro-cardiac space should become more transparent e.g. to improve visibility of superposed lung structure. Another goal is a good transparency of the mediastine to allow distinction of bone structure and vessels, trachea and bronchi, positioning of catheter, wire, artificial cardiac valves, spinal column and the shadow of the heart. The processing should avoid noise amplification in abdomen, mediastine and retro-cardiac space.

A clinical study at the Medical University of Hannover (see also section 3.5) confirmed that density-, and scale-dependent control of enhancement similar to the density dependent part of Noise Containment is well suited to cope with these processing goals. This is because different anatomical details of chest images are separable by the criteria of detail size and local image density. It has to be noted that a prerequisite for the separation of anatomical regions of the chest image by local density is a ranging algorithm that automatically adapts the density curve to the actual exposure of the detector such that characteristic regions are always mapped to the same local density. This is state of the art in current products.[13]

region, detail	local density	scales	required contrast enhancement
abdomen, noise	≤ 0.5	0, 1	no enhancement
abdomen, relevant details	≤ 0.5	≥ 2	moderate to strong enhancement
mediastine & retro-cardiac, noise	≤ 1.0	0, 1	no enhancement
mediastine & retro-cardiac, relevant details	≤ 1.0	≥ 2	moderate to strong enhancement
lung, rib contours	$\simeq 1.7$	0, 1	moderate enhancement
lung, pulmonary vascular system	$\simeq 1.7$	2, 3	low enhancement
lung, large weakly contrasting nodules	$\simeq 1.7$	≥ 4	moderate to strong enhancement

Table 1. Identification of anatomical regions within chest images by scale and local density, and the requirements on contrast enhancement, respectively. The statements hold for a spatial resolution of 2.5 lp/mm.

The identification of anatomical details of chest images is listed in Table 1. The requirements on contrast enhancement result from the processing goals mentioned above. Also, noise sensitive regions can be identified, which leads to a quite satisfactory implementation of Noise Containment as a spin-off feature of Density-Controlled Contrast Enhancement. Density-controlled contrast enhancement is realized by integrating a density- and scale-dependent control into the enhancement which influences the amount of enhancement. This algorithm corresponds to Noise Containment (figure 7) without the activity-dependent control path. The subband pixel is first enhanced according to the contrast amplification function (4). Then the difference value between enhanced and original pixel is weighted by an attenuation factor b that depends on the subband index i (= spatial frequency) and local image density. This weighted difference value is finally added to the original subband pixel. The attenuation factor b is selected such that it is smaller than 1 in regions of reduced enhancement and equal to 1 in regions of maximum enhancement. This concept works regardless of the type (linear, nonlinear, ...) of enhancement. Results of the density-controlled multiscale processing were compared with a standard chest processing.[13] The weighting factors for multiscale contrast enhancement were determined at the Medical University of Hannover (figure 9). Contrast was generally judged to be better in the multiscale processed image, especially in the brighter regions. Due to the selection of the density curves, the processing did not disturb the balanced image impression, and did not feign non-existing pathologies. A small drawback was a slight increase in noise that could not be avoided.

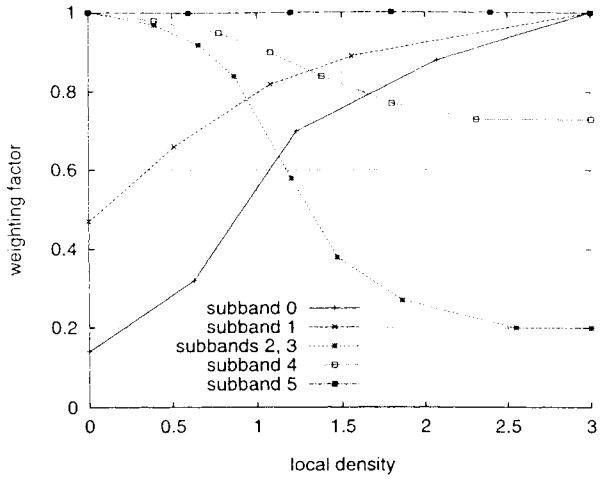

Figure 9. Weighting factor b for density-controlled multiscale chest processing.

3.4. Equalization

We define *Equalization* as scale-dependent variation of the contrast amplification function. Contrary to density-controlled contrast enhancement, which can also vary from scale to scale but only affects certain regions of the image, *Equalization* influences the overall impression of the image. Apart from the scale-dependent factors CF_i, we have for practical reasons restricted scale dependence to the gain factor G_i in (4). While easy to handle, this produces clinically valuable results without causing visible artifacts. Preference of higher spatial frequencies, preference of lower spatial frequencies and equal treatment of all scales are three typical Equalizer settings for the gain that we determined during our clinical studies.

The preference of the subband image corresponding to the smallest details provides a sharper image impression. This was the preferred parameterization for images where the diagnostic interest was exclusively in small details like bone texture. Stronger enhancement of lower scales would only produce disturbing density variations resulting from superposed soft tissue. In such images, the additional boosting of noise could be accepted because noise is masked by high frequent texture information.

The equal treatment of all scales was preferred for images where the diagnostic interest was equally spread from very small to large details. Superposition effects should not be reduced but structures of different sizes should become easier to distinguish. Boosting of the noise level could irritate in some regions of the image. In this case the noise compensation mechanism proved to be a suitable countermeasure. Typical examples are lateral spine images (see figure 10). Equal treatment of all scales is appropriate because diagnostic interest targets small details, such as the posterior edges of a vertebra, as well as larger details like a body of a vertebra. Unacceptable noise boosting on the finest scales, particularly in the bright regions of the upper and lower parts of this radiograph, could successfully be prevented through the Noise Containment options.

If the image contains few small sized details of diagnostic interest, preferred enhancement of large sized details is a good means to prevent too much noise amplification (see Figure 6). Clinicians found this to give a three dimensional impression to the images. Pelvis images, for instance, were found to look similar to conventional (analogue) radiographs without loosing the advantage of better exploitation of the available dynamic range. Enhancement of weakly contrasting medium sized details is desired, while the same enhancement applied to the finest scales was said to generate an "unnatural" image impression. This led to the contrast amplification functions in Figure 10. The smooth transition of contrast gain values over scale was judged to give a balanced image impression.

3.5. Clinical Studies

The multiscale algorithm was evaluated on a storage phosphor radiography system in the radiology departments of Fulda Municipal Hospital, and on a selenium-based chest-radiography system at the Medical University of Hannover. A clinical prototype incorporating the complete functionality of a diagnostic workstation was connected to the

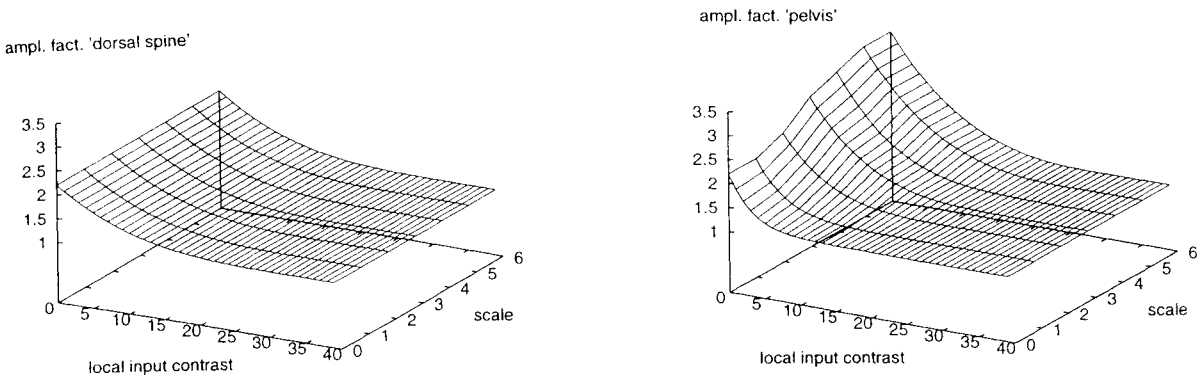

Figure 10. Contrast amplification function for dorsal spine (left) and pelvis (right) as dependent on contrast amplitude and scale index.

radiography systems, allowing continuous operation in the diagnostic routine. A comprehensive archiving solution allowed consistent storage of images, processing parameters and radiologists' comments. For printing, images were fed back to the hospital workstation. Evaluation of the images was exclusively made on film printouts. On the prototype

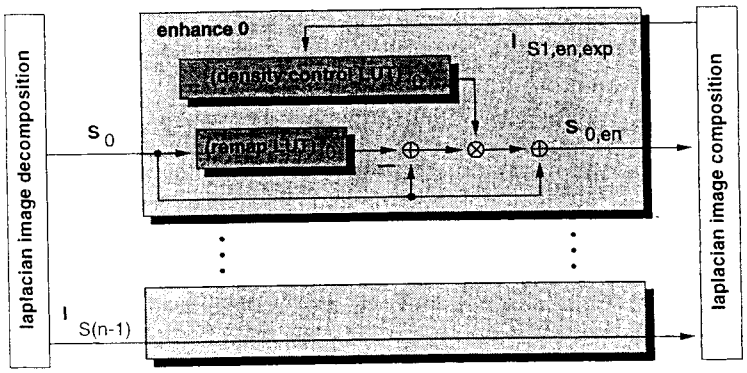

Figure 11. Design of subband enhancement.

the described algorithm was realized as follows (see Figure 11). Prior to multiscale contrast enhancement, a film-like look-up table was applied to the radiographs. Then the images were subjected to a pyramidal decomposition according to the compatibility criteria described in 2. For linear contrast enhancement and additional Structure Boost all subband images S_i, $0 \leq i < (n-1)$ were subjected to a remapping operation. The corresponding look-up table transforms the contrast values according to:

$$c_{i,remapped}(x,y) = CF_i(c_i(x,y)) \cdot c_i(x,y) \tag{6}$$

$CF_i(c_i(x,y))$ obeys (4). Noise Containment is controlled by the local density measure only, so that this feature could be realized via density-controlled contrast enhancement. (see section 3.2). The density measure can easily be retrieved from the expanded reconstructed image of the next lower level. The clinical studies were divided into two phases: one for parameter optimization under a radiologist's guidance, and one to collect feedback on the processing results from a larger number of radiologists.

During the first phase the image impression was optimized by exploiting the modification options provided by our algorithm. Stability of the obtained parameters was verified by throughput of a larger number of radiographs.

We covered the major part of bone diagnostics as well as chest and abdominal regions. This way we collected a considerable number of routine radiographs, optimized parameter sets, and expert comments on case histories, image quality, and processing performance. The comments were arranged in an HTML-based documentation, with the images serving as active links. Clicking into the images automatically loads them into the prototype and activates the corresponding parameter set, thus permitting recreation of the clinical trials.

In the case of the study on the storage phosphor system, twenty images processed by both standard processing and multiscale processing, using the optimized parameters, were presented side-by-side to 14 radiologists, without identifying to them which ones were multiscale-processed, and which ones were standard-processed. The comparative evaluations were based on the criteria in Figure 12, defined by the radiologists well before the session started. The

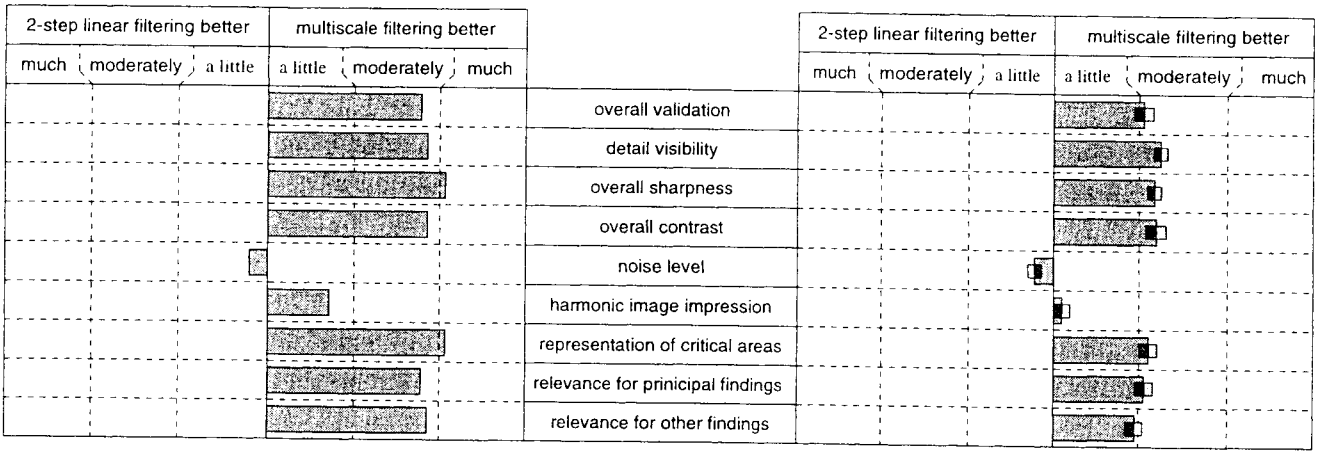

Figure 12. Opinion scores for dorsal spine, averaged over all radiologists (left). Scores averaged over all radiologists and all evaluated examination types (right).

right part of diagram 12 summarizes the results, showing grand averages over all images and all readers (broad bar). The averages are based on numerical values ranging from -3 to $+3$. Additionally the standard deviation of each grand average value is given (narrow bar). Multiscale processing outperforms the standard algorithm almost everywhere, without affecting the subjective criterion of "harmonic image impression". The increased enhancement potential did evidently not result in a related boost of noise. Scores for the last three criteria (critical areas, principal and other findings) show that there is indeed good reason to believe that the new processing eases the diagnostic process. The study also revealed that the achievable quality enhancement depends on the imaged anatomical region. For some regions like the hand, results from standard processing are already close to the optimum, since almost no hidden details could be made better visible. For other regions such as dorsal spine, multiscale processing significantly improved image quality (see left part of diagram 12), much better in fact than the averaged scores show. Hence, as long as the focus is on harmonization and fine detail enhancement in not too noisy conditions, the described reference algorithm is already a powerful processing tool. Due to backward compatibility, multiscale processing should always be at least as good as this reference algorithm.

In the case of the study on the selenium-based chest system, images of an anthropomorphic chest phantom that were randomly superimposed with simulated lesions were processed by chest standard processing[13] and multiscale processing. Detection of lesions was assessed by five radiologists and evaluated by Jackknife ROC-methodology. The study showed that multiscale enhancement was significantly superior for the detection of micro-nodules in the lung area of adipose patients (A_z 0.82 vs. 0.76). For other lesion types investigated, differences were not significant, though there was a general trend towards superior detection with multiscale enhancement.[14]

4. DISCUSSION

We developed a new nonlinear multiscale algorithm that can be derived from unsharp masking-based standard processing. Backward compatibility enables re-use of processing know-how acquired on current standard processing and eases acceptance of the new algorithm. Particular attention was paid to noise robustness.

A clinical study confirms that a suitable design of the subband processing yields a significantly improved detail visibility for a broad spectrum of routine radiographs without unacceptable boosting of noise and distortion of a natural image impression.

Detectors of the next generation as e.g. solid state detectors will bring further improvements with respect to versatility, signal to noise ratio, dynamic range, and spatial resolution.[15,16,1] We expect that the advantages of the new processing, especially the specific enhancement of weakly contrasting structures of any size, will be indispensable to exploit the full potential of these detectors.

ACKNOWLEDGMENTS

We thank the director, Prof. M. Galanski, and the staff of the Institute of Diagnostic Radiology of the Medical University of Hannover for their support. We thank the director, Prof. J.-P. Haas, and the staff of the radiology department of Fulda Municipal Hospital for their support. We are especially grateful to Dr. E. Dencker and Dr. E. Müller for their helpful contributions and their patience during many extended sessions. We furthermore thank our colleagues at Philips Medical Systems, Hamburg, for fruitful discussions, technical help and for providing applicational background.

REFERENCES

1. T. Aach, U. Schiebel, and G. Spekowius, "Digital image acquisition and processing in medical x-ray imaging," *Journal of Electronic Imaging (Special Section on Biomedical Image Representation)* **8**(1), 1999 (in press).
2. M. Ishida, *Computed Radiography (Tateno, Iinuma, Takano, eds.)*, ch. Image Processing, pp. 25–30. Springer Verlag, Tokyo, 1987.
3. I. Maack and U. Neitzel, "Optimised image processing for routine digital radiography," in *Proceedings Computer Assisted Radiology*, U. Lemke, ed., pp. 108–114, Berlin, Springer Verlag, 1991.
4. A. R. Cowen, A. Giles, A. G. Davies, and A. Workman, "An image processing algorithm for PPCR imaging," in *SPIE Vol. 1898, Image Processing*, pp. 833–843, 1993.
5. T. L. Ji, M. K. Sundareshan, and H. Roehrig, "Adaptive image contrast enhancement based on human visual properties," *IEEE Transactions on Medical Imaging* **13**(4), pp. 573–586, 1994.
6. A. G. Davies, A. R. Cowen, G. J. S. Parkin, and R. F. Bur, "Optimising the processing and presentation of PPCR images," in *SPIE Vol. 2712, Image Processing*, pp. 189–195, 1996.
7. A. N. Akansu and R. A. Haddad, *Multiresolution Signal Decomposition*, Academic Press, Boston, 1992.
8. P. J. Burt and E. H. Adelson, "The Laplacian pyramid as a compact image code," *IEEE Transactions on Communications* **31**(4), pp. 532–540, 1983.
9. R. F. Bessler and J. H. Arbeiter, "Contrast enhancement using Burt pyramid processing," in *SID International Symposium, Digest of Technical Papers*, pp. 352–353, (San Diego), May 6–8 1986.
10. P. Vuylsteke and E. Schoeters, "Multiscale image contrast amplification (MUSICA^{TM})," in *SPIE Vol. 2167, Image Processing*, pp. 551–560, 1994.
11. S. Ranganath and H. Blume, "Hierarchical image decomposition and filtering using the S-transform," in *SPIE Vol. 914*, pp. 799–814, 1988.
12. S. Ranganath, "Image filtering using multiresolution representations," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **13**(5), pp. 426–440, 1991.
13. H. Blume, "Digital processing of THORAVISION images: The Algorithms," product description, Philips Medical Systems, 1996.
14. E. Dencker, M. Stahl, S. Dippel, C. M. Schaefer-Prokop, S. Baus, S. Goehde, D. Hoegemann, A. Leppert, and M. Galanski, "Nonlinear multiscale processing in selenium-based chest radiography," in *ECR, Book of Abstracts*, 1999.
15. N. Jung, P. L. Alving, F. Busse, N. Conrads, H. M. Meulenbrugge, W. Rütten, U. Schiebel, M. Weibrecht, and H. Wieczorek, "Dynamic x-ray imaging system based on an amorphous silicon thin-film array," in *SPIE Vol. 3336*, 1998.
16. U. Schiebel, N. Conrads, N. Jung, M. Weibrecht, H. Wieczorek, T. Zaengel, M. J. Powell, I. D. French, and C. Glasse, "Fluoroscopic imaging with amorphous silicon thin-film arrays," in *SPIE Vol. 2163*, pp. 129–140, 1994.