

Three-Dimensional Edge-Preserving Image Enhancement for Computed Tomography

Nicolas Villain, Yves Goussard, *Member, IEEE*, Jérôme Idier, and Marc Allain

© 2003 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Abstract—Computed tomography (CT) images exhibit a variable amount of noise and blur, depending on the physical characteristics of the apparatus and the selected reconstruction method. Standard algorithms tend to favor reconstruction speed over resolution, thereby jeopardizing applications where accuracy is critical. In this paper, we propose to enhance CT images by applying half-quadratic edge-preserving image restoration (or deconvolution) to them. This approach may be used with virtually any CT scanner, provided the overall point-spread function can be roughly estimated. In image restoration, Markov random fields (MRFs) have proven to be very flexible *a priori* models and to yield impressive results with edge-preserving penalization, but their implementation in clinical routine is limited because they are often viewed as complex and time consuming. For these practical reasons, we focused on numerical efficiency and developed a fast implementation based on a simple three-dimensional MRF model with convex edge-preserving potentials. The resulting restoration method provides good recovery of sharp discontinuities while using convex duality principles yields fairly simple implementation of the optimization. Further reduction of the computational load can be achieved if the point-spread function is assumed to be separable. Synthetic and real data experiments indicate that the method provides significant improvements over standard reconstruction techniques and compares well with convex-potential Markov-based reconstruction, while being more flexible and numerically efficient.

Index Terms—Edge-preserving convex potentials, half-quadratic criteria, Markov random fields, three-dimensional image reconstruction, three-dimensional image restoration, x-ray tomography.

I. INTRODUCTION

MODERN medical imaging systems such as computed tomography (CT) or magnetic resonance imaging (MRI) provide the practitioner with high-resolution images that can be used as powerful diagnostic tools or—more recently—as sources of three-dimensional (3-D) anatomical data for appli-

Manuscript received November 25, 2002; revised June 11, 2003. This work was supported in part by the Natural Science and Engineering Research Council of Canada under Research Grant OGP013 8417 and under the France-Québec Cooperative Research Project PV P 40/7. The Associate Editor responsible for coordinating the review of this paper and recommending its publication was A. Manduca. *Asterisk indicates corresponding author.*

N. Villain was with the Biomedical Engineering Institute, École Polytechnique, Montreal, QC H3C 3A7, Canada. He is now with Philips Research France, 92156 Suresnes Cedex, France.

*Y. Goussard is with the Biomedical Engineering Institute, École Polytechnique, Montreal, QC H3C 3A7, Canada.

J. Idier was with the Laboratoire des Signaux et Systèmes, Supélec, 91192 Gif-sur-Yvette Cedex, France. He is now with Institut de Recherche en Communications et Cybernétique de Nantes, École Centrale de Nantes, 44321 Nantes Cedex 3, France.

M. Allain is with the Biomedical Engineering Institute, École Polytechnique, Montreal, QC H3C 3A7, Canada.

Digital Object Identifier 10.1109/TMI.2003.817767

cations like computer-aided surgery, model design for surgery planning, or custom-fitted articular implants. The precision one can expect from these data is limited by the resolution of the device used to collect them. Experimental investigation of the characteristics of a typical commercial CT scanner showed that the precision is of the order of 1.4 mm [1], which is satisfactory for most common clinical uses of scanners such as, e.g., diagnosis. However, in some situations, a greater precision is required. Such is the case in the application that motivated this study, where the geometry of the knee joint has to be acquired with a submillimeter accuracy in order to design and manufacture personalized prostheses. Therefore, the images need to somehow be enhanced.

To increase the precision of the images, one approach would consist of performing two-dimensional (2-D) or 3-D reconstruction directly from the projection data with methods that are more accurate than the backprojection-type techniques generally implemented in commercial scanners. For instance, recently proposed edge-preserving reconstruction methods based on Markov random field (MRF) models (see, e.g., [2]–[4]) demonstrate impressive results on simulated examples. Nevertheless, two main reasons could limit their clinical use.

The first reason is merely practical. Reconstruction methods require access to the projection data, which commercial CT scanners scarcely allow. The second reason is that reconstruction techniques are specific to a given device because they require knowledge of its exact characteristics, which may prove difficult to obtain from the vendor. For instance, although modern CT scanners usually share the same structure (rotating X-ray tube coupled with a circular detector array), reconstruction algorithms have to take into account the exact geometry of the imaging system (diameter of the detector array, distance between the source and the detectors, etc.), which is seldom available to the end-user. Both reasons are particularly important when one wants to work in several hospitals equipped with different imaging systems.

In order to achieve greater precision, another option consists of processing the already reconstructed images under adequate assumptions about the nature of the degradations that affect them. One such assumption is the linearity of the degradation process [5]. In this framework, distortions can be easily identified in an experimental manner with the use of an appropriate phantom, as explained in [1]. Therefore, as illustrated in Fig. 1, since the PSF can be estimated by the user, this approach becomes highly portable and may be applied on any system. This is the reason why it is adopted here.

Within the framework defined above, image enhancement becomes an image restoration problem. Our goal is to develop an

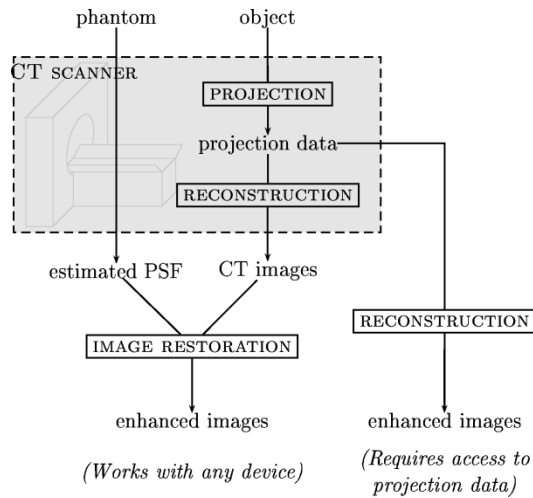


Fig. 1. Method comparison.

image restoration method that complies with the following constraints.

- *Accurate restoration of image discontinuities.* This is an obvious requirement in order to improve the precision of the images, and is even more important since edge location conveys important anatomical and geometrical information.
- *Ability to process 3-D data at a reasonable computational cost.* Most biological structures investigated with imaging systems such as CT scanners are 3-D in nature and are, therefore, better characterized in a 3-D framework. This becomes a strong requirement when one is interested in the geometry of such structures. In addition, as illustrated in the sequel, explicitly accounting for the 3-D nature of the objects under investigation tends to increase the quality of the results.

Image restoration is known to be an ill-posed problem [6], which needs to be regularized in order to yield good quality results. Regularization can be achieved by the addition of *a priori* constraints to the problem and its solution. Here, we adopt a Bayesian approach in which the *a priori* constraints are specified under the form of probability density functions (PDFs) of the original object to be restored and of the noise that unavoidably corrupts the data. In our problem, the PDF of the noise process does not have a strong impact on the quality of the results and a simple noise model is generally satisfactory. On the other hand, selecting an adequate *a priori* model of the object is more complicated. Non-Gaussian MRFs have proven to be adequate models for media composed of homogeneous regions separated by sharp discontinuities, such as many biological media, and to yield good edge preservation. In addition, these models are based upon local pixel interactions [7] that can be easily customized, thereby offering flexibility to tailor the model to the problem at hand. However, MRFs designed to restore sharp discontinuities often yield nonconvex criteria [8] and require complex optimization techniques, which is a severe drawback for efficient processing of 3-D data. In order to improve the numerical efficiency of MRF-based image restoration methods, Geman and Yang [2] and Geman and Reynolds [9] proposed simulated annealing techniques based on convex duality princi-

ples for maximizing the objective function. However, the computational cost still remains high. For similar purposes, Bouman and Sauer [10] proposed to use special *convex potential* MRFs that can still provide adequate edge preservation while yielding a convex criterion, thereby simplifying the optimization considerably. Nonetheless, in both cases, the numerical cost may depend directly on the spatial extent of the linear distortion.

Here, we propose a 3-D image restoration method that combines convex potential MRFs and duality-based maximization along the lines presented in [11]. First, we make assumptions on the image formation process so as to simplify its mathematical formulation without jeopardizing the quality of the results. Second, in order to fully account for the 3-D nature of the structures under investigation, we define a 3-D MRF prior model with convex potentials, which simplifies the optimization while yielding adequate edge preservation. Third, we derive a single site update (SSU) algorithm based on the above assumptions and on convex duality principles. As illustrated by results obtained with synthetic and real data, the resulting method is able to process 3-D data in a reasonable amount of time on low-end computers while providing adequate edge preservation.

The remainder of this paper is organized as follows. In Section II, we present some background material on Bayesian estimation and define the models of the image formation process and of the 3-D MRF priors. In Section III, we then focus on the choice of the minimization method. In Section IV, we present the algorithm we developed to cope with the increased numerical complexity of 3-D data. Results obtained on synthetic and real data are presented in Section V and conclusions are drawn in Section VI.

II. FRAMEWORK

A. Image Formation

Here, we consider a discrete real-valued 3-D object $\mathbf{x} = \{x_s\}_{s \in \mathcal{S}}$; $\mathcal{S} = \{1, \dots, M\} \times \{1, \dots, N\} \times \{1, \dots, Z\}$ observed through a CT scanner as a series of CT images $\mathbf{y} = \{y_s\}_{s \in \mathcal{S}}$. As often, we assume that the imaging system is linear with a point-spread function (PSF) \mathcal{H} . We then make the following important assumptions regarding the PSF of the imaging system.

- The PSF is assumed to be 2-D. Actually, the physical functioning of CT scanners implies that a given image depends on the object density inside the X-ray beam (and, likely, somewhat outside due to diffusion in neighboring regions). However, in practical cases, as long as the width of the beam is smaller than the distance between two consecutive images, we can consider that a given slice of the observed data (i.e., a given image) depends only on the corresponding slice of the original discrete object.
- The PSF is assumed to be spatially invariant. The reason is twofold. On the one hand, within the region of interest, which is small compared to the whole field of view of the CT scanner, the PSF does not vary much (less than 7% change in the full width at half maximum along both tangential and radial axes according to [1]). On the other hand, we have found the variations of the PSF induce negligible effects on restoration (see Section V for an illustra-

tion). This relatively weak assumption allows us to model the observation process by a simple space-invariant convolution operation.

Thus, the image formation process with additive noise can be written in matrix form as

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{n} \quad (1)$$

where \mathbf{H} is the convolution matrix constructed from the PSF \mathcal{H} .

B. Bayesian Estimation

In the Bayesian framework introduced in the previous section, the prior knowledge on the image to be restored is specified through PDF $p(\mathbf{x})$, while the information about the observed data and the image formation process is contained in conditional PDF $p(\mathbf{y}|\mathbf{x})$. Using Bayes rule, one can write

$$p(\mathbf{x}|\mathbf{y}) \propto p(\mathbf{y}|\mathbf{x})p(\mathbf{x}) \quad (2)$$

and PDF $p(\mathbf{x}|\mathbf{y})$ represents all knowledge available on the unknown quantity \mathbf{x} once all sources of information are accounted for. Since manipulation of a complete PDF may be intractable in practice, $p(\mathbf{x}|\mathbf{y})$ is often used to define a point estimator. Here, we choose the *maximum a posteriori* (MAP) estimator because it yields a very effective and adaptable criterion. The MAP estimate $\hat{\mathbf{x}}$ is defined as

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} [p(\mathbf{x}|\mathbf{y})] = \arg \max_{\mathbf{x}} [p(\mathbf{x})p(\mathbf{y}|\mathbf{x})] \quad (3)$$

and we now specify PDFs $p(\mathbf{y}|\mathbf{x})$ and $p(\mathbf{x})$.

$p(\mathbf{y}|\mathbf{x})$ only depends on the image formation process defined in (1) and on the noise PDF $p(\mathbf{n})$. Although the noise of reconstructed images turns out to be correlated, we noticed that the simplest model of additive white Gaussian noise provided satisfactory results for our restoration approach. In that case, the only free parameter is the noise variance σ_n^2 , which can be estimated from regions where biological tissues are uniform enough. In these regions, the intensity variations in CT images directly reflect the additive noise. With that model, $p(\mathbf{y}|\mathbf{x})$ can be expressed as

$$p(\mathbf{y}|\mathbf{x}) = K \exp \left\{ -\frac{\|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2}{2\sigma_n^2} \right\} \quad (4)$$

where K denotes the appropriate normalizing factor.

The choice of an *a priori* PDF of the object to be restored is more delicate. As explained in Section I, we selected MRF models in order to account for and restore sharp discontinuities present in the original image. Under a positivity condition, which is generally fulfilled in practice, a convenient way of specifying the PDF of an MRF is to use the Gibbs formulation that takes the form [12]

$$p(\mathbf{x}) = \frac{1}{G} \exp \left\{ -\lambda \sum_{c \in \mathcal{C}} V_c(\mathbf{x}) \right\} \quad (5)$$

where V_c denotes the Gibbs potential defined on each set c of voxels which are mutual neighbors (called a clique), and where G is the normalizing factor. Substituting (4) and (5) into (3)




Fig. 2. Chosen 3-D neighborhood of a given voxel (gray region) and corresponding two-element cliques.

yields the formal expression of the estimate $\hat{\mathbf{x}}$, which, after applying a neglog operation, is more conveniently expressed as

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} [J(\mathbf{x})] \quad (6)$$

with

$$J(\mathbf{x}) = \frac{\|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2}{2\sigma_n^2} + \lambda \sum_{c \in \mathcal{C}} V_c(\mathbf{x}). \quad (7)$$

Despite the formal simplicity of criterion J , its minimization may be either straightforward or quite unwieldy depending on the choices made for the set of cliques \mathcal{C} and the potentials V_c . Firstly, the choice of the clique system determines the range of interactions between pixels. For example, in a 2-D setting, the most studied neighborhoods are isotropic and ordered: first order for the nearest four pixels, second order for the nearest eight ones, etc. In most cases, third-order and higher order schemes yield complex criterion and that is the reason why they are scarcely used. Second and maybe more importantly, many kinds of potentials have been proposed; they correspond to various tradeoffs between quality of the restoration and ease of minimization of the resulting criterion. For continuous-valued MRFs, they fall into three main groups of increasing complexity and quality, i.e., quadratic, convex, and nonconvex. In practical cases, the choice of both the clique system and potentials determines the numerical cost of the method. We found that, in order to deal with 3-D data, we had to adopt the simplifying choices described below.

C. Simplified 3-D MRF Model

First, we define a neighborhood system with the corresponding set of cliques \mathcal{C} . Then, we describe the Gibbs potentials defined on the cliques.

As \mathbf{x} is modeled as a 3-D field, the neighborhood of a given voxel extends in 3-D into the neighboring slices. Since the cross-slice distance is greater than cross-voxel distances within the same slice, we choose the eight nearest voxels from the same slice and the single nearest voxels of the upper and lower slices, as shown in Fig. 2.

Even with a limited neighborhood like this one, the set of all possible cliques includes several three-element cliques and one four-element clique. Here, we choose to consider only two-element cliques because they make it easier to define Gibbs potentials conveying local smoothness. For larger cliques, more complicated interactions are also more difficult to interpret. For example, Geman and Geman [13] proposed an elaborated model with four-element cliques, but it involves line variables, thus increasing numerical complexity. Therefore, another reason why

using only two-element cliques is to keep the computational cost as low as possible.

We, thus, define Gibbs potentials only on \mathcal{C}_2 , the set of two-voxel cliques, as penalty functions φ of the gradient u_c inside the cliques

$$u_c = \frac{\Delta x_c}{d_c} \quad (8)$$

where Δx_c and d_c , respectively, denote the difference of the values and the distance of the two voxels in clique c . With our choice of cliques, which is shown in Fig. 2, d_c can only take three different values, i.e., d_1 for two nearest voxels of the same slice, $\sqrt{2}d_1$ for a diagonal clique in the same slice, and $d_2 > d_1$ for a clique between two different slices. The criterion becomes

$$J(\mathbf{x}) = \frac{\|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2}{2\sigma_n^2} + \lambda \sum_{c \in \mathcal{C}_2} \varphi(u_c). \quad (9)$$

The choice of penalty function φ is guided not only by the accuracy of the *a priori* model, but also by practical considerations, i.e., the simplicity of optimization of the resulting criterion J . Indeed, once J is defined, image restoration simply consists of solving $\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} [J(\mathbf{x})]$. With Z slices of $M \times N$ pixels, \mathbf{x} is composed of $Z \times M \times N$ voxels. Even with a convex criterion, optimization can be numerically burdensome.

In Section III, we show that optimization can be carried out with local iterative algorithms and that a low numerical cost can be achieved by replacing the original criterion by an equivalent augmented one.

III. HALF-QUADRATIC APPROACH

A. Local Iterative Optimization

Although we obtained the criterion in (9) in the Bayesian framework of MAP estimation, it could equivalently result from other stochastic (e.g., penalized maximum likelihood) or deterministic (e.g., constrained least squares) approaches. In our case, we carefully chose the model in order to ensure the convexity of J and the uniqueness of its minimum. However, due to the size of the objects involved, finding efficient minimization methods is a challenging numerical problem that is crucial for practical applications.

When global update algorithms such as standard, conjugate or pseudoconjugate gradient algorithms [14], [15] are used for minimizing (9), each iteration requires the computation of the gradient $\partial J(\mathbf{x})/\partial \mathbf{x}$ and at least one time the computation of the criterion $J(\mathbf{x})$ depending on the line search minimization method. Fast convolution techniques [using fast Fourier transform (FFT)] and separable PSFs can considerably reduce the computational cost, but memory requirements of global iterative algorithms remain high (roughly speaking, from three times to six times the image size) and domain constraints are difficult to introduce.

On the contrary, SSU algorithms like relaxation or Gauss–Siedel [16] are less memory demanding twice the image size in 2-D) and hard domain constraints can be easily introduced. Besides, local minimization is well suited to MRFs

TABLE I
OVERVIEW OF THE ALGORITHM

- | |
|---|
| <p>(a) Initialization.
 (b) For each slice $k = 1, \dots, Z$:
 (c) For each column $j = 1, \dots, N$:
 (e) For each pixel $s = (i, j, k)$, $i = 1, \dots, M$:
 (e1) Calculate m_s according to (20)
 (e2) Update x_s according to the over-relaxation scheme
 $x_s^{\text{new}} \leftarrow x_s + \alpha(m_s - x_s)$
 (e3) For every clique c including s:
 update l_c according to (21)
 Iterate step (e) until end of the current column.
 Iterate step (c) until end of the current slice.
 Iterate step (b) until convergence.</p> |
|---|

because there is no long-range interaction between voxels. The sphere of influence of a given voxel is thus determined by the size of the PSF. For these reasons, we use the same approach as Besag [7] and Bouman and Sauer [10] and choose an iterative relaxation method. In addition, we adopt an over-relaxed scheme since such a choice has been shown to yield faster convergence [11], [15].

The problem with local iterative methods is that, for any penalty function φ , there is no closed-form expression for the minimum of J with respect to a given voxel. The only well-known exception is when φ is quadratic, but it yields restored images that are too smooth for our application. However, in Section III-B, we show that simple results of convex analysis allow us to derive a local minimization strategy that is proven to converge, although it does not require any line search step.

B. Augmented Criterion

A recent approach has been to reformulate the restoration as the minimization of an equivalent criterion derived from the previous criterion (9) through the application of convex analysis results. Both criteria are equivalent in the sense that they share the same global minimum. However, the nice thing about the augmented criterion is that its local minimization becomes straightforward for a much larger class of penalty functions φ including edge-preserving functions quite suitable for our application.

Indeed, by introducing auxiliary variables \mathbf{l} , as in [2], [9], [11], and [17], the original criterion $J(\mathbf{x})$ can be replaced by an augmented criterion $K(\mathbf{x}, \mathbf{l})$ with the following properties:

$$\min_{\mathbf{l}} [K(\mathbf{x}, \mathbf{l})] = J(\mathbf{x}) \quad (10)$$

$$K \text{ is quadratic with respect to } \mathbf{x} \quad (11)$$

$$K \text{ can be easily minimized with respect to } \mathbf{l}. \quad (12)$$

Due to property (11), such a criterion is called half-quadratic and the corresponding approach is called half-quadratic regularization.

Here, we give only the assumptions and convex analysis results needed to carry out the minimization. For detailed mathematical justification and convergence proofs, see [18].




Fig. 3. Estimated widths of the PSF of a CT scanner in tangential and radial directions as a function of radius in polar coordinates.

The required assumptions on the penalty function are the following:

$$\varphi \text{ is even} \quad (13)$$

$$\varphi(\sqrt{\cdot}) \text{ is strictly concave on } \mathbb{R}_+ \quad (14)$$

$$\varphi \text{ is continuous near zero and is } C^1 \text{ on } \mathbb{R} \setminus \{0\}. \quad (15)$$

If these conditions are met, which is the case for several penalty functions found in the literature, we can define an auxiliary function ψ by the relationship

$$\psi(l) = -\min_x (lx^2 - \varphi(x)). \quad (16)$$

By construction, $-\psi$ is the concave conjugate of $\varphi(\sqrt{\cdot})$ [19]. Since the latter is concave (14), it is, in turn, the concave conjugate of $-\psi$, which yields

$$\varphi(x) = \min_l (lx^2 + \psi(l)). \quad (17)$$

From (17), it is obvious that, by introducing an auxiliary variable l_c for each two-element clique c , the augmented criterion defined by

$$K(\mathbf{x}, \mathbf{l}) = \frac{\|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2}{2\sigma_n^2} + \lambda \sum_{c \in \mathcal{C}_2} (l_c u_c^2 + \psi(l_c)) \quad (18)$$

satisfies (10). In addition, K is quadratic as a function of \mathbf{x} and we show in Section IV that minimization of K with respect to a given auxiliary variable l_c is a local and inexpensive operation defined by closed-form expression (21). Therefore, K fulfills conditions (10)–(12).

These properties ensure that local minimization of K with respect to \mathbf{x} and with respect to \mathbf{l} are simpler than the local min-

imization of the original criterion J with respect to \mathbf{x} alone. Concerning global minimization, if φ is convex, J is convex in \mathbf{x} , and its global minimum can be reached by local iterative minimization. The question is now whether this property is also shared by K . If φ is convex, it is obvious that K is both convex in \mathbf{x} when \mathbf{l} is fixed and in \mathbf{l} when \mathbf{x} is fixed, but this does not necessarily imply global convexity, which is often overlooked. Nevertheless, it is shown in [18] that when φ is convex and a few other technical requirements are satisfied, K is also convex up to a change of variables. Therefore, it is rigorously possible to reach its global minimum by alternate minimization with respect to the intensity variables \mathbf{x} and to the auxiliary variables \mathbf{l} .

Another very appealing aspect of this approach is that the various numerical recipes developed for quadratic optimization can be directly adapted to half-quadratic regularization. Depending on the simplifying assumptions available in a given practical case, tailored techniques can be used to achieve faster or more efficient convergence. This allowed us to develop an optimized recursive algorithm for separable PSFs—presented in Section IV—that requires significantly fewer operations per iteration, especially with large PSFs, than the corresponding algorithm without simplifications.

IV. ALGORITHM

A. Update Strategy

As a function of \mathbf{x} , K is a quadratic function; its minimum value can be expressed in closed form and is reached at m_s defined by

$$\left. \frac{\partial K(\mathbf{x}, \mathbf{l})}{\partial x_s} \right|_{x_s=m_s} = 0 \quad (19)$$




Fig. 4. (a) Test phantom made of enamel wires to estimate the PSF of a CT scanner. (b) CT image of the phantom with the two selected samples. (c) Sample PSF close to the center (radius ≈ 30 mm). (d) Sample PSF far from the center (radius ≈ 90 mm).

and solving for m_s yields an explicit expression [11]

$$m_s = x_s + \frac{[\mathbf{H}^t \mathbf{y}]_s - [\mathbf{H}^t \mathbf{H} \mathbf{x}]_s + 2\sigma_n^2 \sum_{\{r,s\} \in \mathcal{C}_2} \lambda_{\{r,s\}} (x_r - x_s)}{[\mathbf{H}^t \mathbf{H}]_{s,s} + 2\sigma_n^2 \sum_{\{r,s\} \in \mathcal{C}_2} \lambda_{\{r,s\}}} \quad (20)$$

where the sums extend to the neighborhood of the currently visited voxel s , and where $[\mathbf{y}]_i$ and $[\mathbf{M}]_{i,j}$, respectively, denote the i th element of vector \mathbf{x} and element (i, j) of matrix \mathbf{M} .

As a function of the auxiliary variables \mathbf{l} , $K(\mathbf{x}, \mathbf{l})$ reads $\sum_{c \in \mathcal{C}_2} K_c(l_c, u_c)$, which is independently minimized with respect to each auxiliary variable. The minimum is reached for

$$\hat{l}_c = \frac{\varphi'(u_c)}{2u_c} \quad (21)$$

according to a simple duality result [2], [9], [19, Ch. 7], [20].

The whole procedure alternates updates of \mathbf{x} and \mathbf{l} according to (20) and (21). A classical over-relaxation scheme may in-

crease the convergence speed. Instead of directly updating x_s with m_s , a larger correction is applied, i.e., $\alpha(m_s - x_s)$ with $\alpha > 1$ (in our experiments, α was set to 1.7). An overview of the whole algorithm in the unconstrained case is given in Table I. Note that like many SSU algorithms, this procedure lends itself well to the introduction of separable constraints through a simple modification of step (e2), e.g.,

$$(e2) \quad x_s^{\text{new}} \leftarrow \max\{0, x_s + \alpha(m_s - x_s)\}$$

corresponds to a positivity constraint. The proposed procedure performs SSU on the augmented criterion K . Recently, it has been put forward that the same procedure also has the structure of an SSU algorithm on the original criterion J [21, pp. 142–144]. More precisely, it is closely connected to iterative coordinate descent (ICD) procedures proposed by Bouman and Sauer in [22]. However, it has a significantly simpler structure since it neither requires line search, root extraction, nor any nested iteration, each update being performed using a constant step size.

TABLE II
DETAILED ALGORITHM

(a) Initialize: calculate and store $\mathbf{H}^t \mathbf{y}$, \mathbf{g} and \mathbf{w} .
(b) For each slice $k = 1, \dots, Z$:
(c) For each column $j = 1, \dots, N$:
(d) Initialize \mathbf{R}^{0j} (27)
(e) For each pixel $s = (i, j, k)$, $i = 1, \dots, M$:
(e1) Apply the mask $[\mathbf{C}^t \mathbf{C} \mathbf{x}]_{ij} = \mathbf{w}^t \mathbf{R}^{ij}$ (28)
(e2) Add Gibbs and back-filtered contributions to compute m_s (20)
(e3) Update x_s subject to domain constraints $x_s^{\text{new}} \leftarrow x_s + \alpha(m_s - x_s)$
(e2) For every clique c including s , update l_c : $l_c^{\text{new}} \leftarrow \varphi'(u_c)/2u_c$ (21)
(e4) Update $R_0^{ij} \leftarrow R_0^{ij} + w_0(x_s^{\text{new}} - x_s)$ to take x_s^{new} into account.
(e5) $R_k^{i+1,j} = R_{k+1}^{ij}$ for $k = -p + 1, \dots, p - 2$.
(e6) Compute $R_{p-1}^{i+1,j}$ (27)
Iterate step (e) until end of the current column.
Iterate step (c) until end of the current slice.
Iterate step (b) until convergence.

From an implementation standpoint, the auxiliary variables require only an insignificant amount of additional memory if the updates are carefully intertwined. For example, with a column-wise raster scan, it is never necessary to store more than one column of auxiliary variables for first-order cliques, and not more than three columns for second-order cliques. In addition, since $[\mathbf{H}^t \mathbf{y}]_s$ does not depend on the current value of \mathbf{x} , this quantity can be pre-computed and stored. Only the Gibbs part of the criterion (which includes only a few voxels) and quantity $[\mathbf{H}^t \mathbf{H} \mathbf{x}]_s$ have to be updated at each iteration. For a PSF of size $p \times q$, the latter requires $4pq$ multiplications for each visited voxel, which comes down to $2pq$, thanks to central symmetry of $\mathbf{H}^t \mathbf{H}$, which is the 2-D autocorrelation function of the PSF.

For large PSFs, evaluation of this term can be burdensome because it is repeated for each voxel and iteration. Even with PSFs of moderate sizes, this is the main limiting factor in the performance of the algorithm. To tackle this problem, we propose to use the assumption of separability of the PSF in order to reduce the amount of computation required for evaluation of the convolution term $[\mathbf{H}^t \mathbf{H} \mathbf{x}]_s$ and thereby allow to process large datasets in 3-D. In Section IV-B, we show how close to separable the PSF of our CT scanner is and we describe how to compute the convolution term in that case.

B. Efficient Convolution Using a “PSF Separability” Assumption

In many cases, the PSF can be assumed to be separable to make computation simpler. In [1], a phantom made of thin enamel wires (0.25 mm in diameter) was used to estimate the PSF of a CT scanner. As expected, the mechanical acquisition process and the reconstruction algorithm result in a rotational symmetry of image properties. Indeed, in this case, the PSF has two axes of symmetry, i.e., the radial and tangential axes. The full width at half maximum was estimated in those two




Fig. 5. (a) Computer-rendered image of model used for our experiments. (b) Profile of the model with the numbering of the 35 slices that were simulated and acquired with the CT scanner.




Fig. 6. Shape of the potential function (solid line) compared to its quadratic and linear asymptotes (dashed line).

TABLE III
NORMALIZED QUADRATIC ERRORS FOR ALL TESTED METHODS
(SIMULATED DATA SET, 35 127 × 127 SLICES)

Method	e_{Q3D}	$\min(e_{Q2D})$	$\text{mean}(e_{Q2D})$	$\max(e_{Q2D})$
FBP 2D recon.	2.79×10^{-1}	2.57×10^{-1}	2.88×10^{-1}	3.39×10^{-1}
Gauss 3D recon.	3.08×10^{-2}	2.95×10^{-2}	3.19×10^{-2}	4.20×10^{-2}
Half-quad. 2D resto.	6.76×10^{-3}	5.33×10^{-3}	7.18×10^{-3}	9.86×10^{-3}
Half-quad. 3D resto.	5.94×10^{-3}	3.80×10^{-3}	6.23×10^{-3}	9.09×10^{-3}
Half-quad. 3D recon.	4.27×10^{-3}	2.68×10^{-3}	4.45×10^{-3}	7.10×10^{-3}

directions by fitting a 2-D Gaussian curve to the acquired PSF. The results were that the PSF was slightly more elongated in the tangential direction and that this asymmetry grew larger as the distance from the center increased. However, for a 20-cm field of view (large enough for the knee), the maximum difference between PSF widths estimated in the two directions was less than 3%, as reported in Fig. 3. As an example, Fig. 4 shows the phantom used for this experiment and two sample PSFs, one close to the center, the other one at the periphery of the field of view. Both PSFs are almost circular and can, therefore, be assumed to be separable.

In the following, we assume that the PSF is separable, i.e., that \mathcal{H} can be expressed as

$$\mathcal{H} = \mathbf{v} \mathbf{h}^t$$

where \mathbf{v} (respectively, \mathbf{h}) denotes the vertical (respectively, horizontal) generating vector of \mathcal{H} . We then show that separa-




Fig. 7. (a) Graph of 2-D normalized quadratic errors e_{Q2D} for all tested methods applied to the simulated data set. (b) Zoom on the lower part of previous graph.

bility property $\mathcal{H} = \mathbf{v}^t \mathbf{h}$ yields an efficient way of computing $[\mathbf{H}^t \mathbf{H} \mathbf{x}]_s$ from the autocorrelation vectors (\mathbf{g}, \mathbf{w}) of components (\mathbf{h}, \mathbf{v}) defined by

$$g(i) = \sum_{j=i}^q h(j)h(j-i), \quad \text{for } i = -(q-1), \dots, q-1 \quad (22)$$

$$w(i) = \sum_{j=i}^p v(j)v(j-i), \quad \text{for } i = -(p-1), \dots, p-1. \quad (23)$$

This scheme could also be used for a 3-D separable PSF, but, according to our model, we only consider a 2-D PSF and present the optimized algorithm in each slice.

In model definition, we assumed that the image formation process was completely 2-D, which means that a given tomographic image depends only on the corresponding slice of the object. Moreover, since we model the tomograph as a linear system with spatially invariant PSF \mathcal{H} , the relationship between an observed image $\mathbf{Y}^{(z)}$ and the corresponding slice of the (discretized) object $\mathbf{X}^{(z)}$ is given by a simple 2-D convolution

$$\mathbf{Y}^{(z)} = \mathcal{H} * \mathbf{X}^{(z)} + \mathbf{N}^{(z)}. \quad (24)$$

Lining up the columns (or rows) of the matrices into vectors yields the equivalent expression

$$\mathbf{y}^{(z)} = \mathbf{C} \mathbf{x}^{(z)} + \mathbf{n}^{(z)}. \quad (25)$$

where \mathbf{C} is the convolution matrix built from the PSF \mathcal{H} . Finally, concatenating the vectors corresponding to all the slices results in the global model we used thus far (1).

Although it appears to be a 3-D relationship, it is merely a collection of 2-D relationships. Therefore, \mathbf{H} is built from \mathbf{C} and it is easy to show that

$$\mathbf{H}^t \mathbf{H} \mathbf{x} = \begin{bmatrix} \mathbf{C}^t \mathbf{C} \mathbf{x}^{(1)} \\ \mathbf{C}^t \mathbf{C} \mathbf{x}^{(2)} \\ \vdots \\ \mathbf{C}^t \mathbf{C} \mathbf{x}^{(Z)} \end{bmatrix}. \quad (26)$$

Thus, the computation of $[\mathbf{H}^t \mathbf{H} \mathbf{x}]_s$ for a given voxel corresponding to pixel (i, j) in slice z reduces to the computation of $[\mathbf{C}^t \mathbf{C} \mathbf{x}^{(z)}]_{ij}$. In the following equations, we omit reference to the current slice z for the sake of simplicity and develop the convolution product using the matrix \mathbf{X} representing the current slice.

Let us introduce an auxiliary vector \mathbf{R}^{ij} of length $2p-1$ defined by

$$R_k^{ij} = \sum_{l=-(q-1)}^{q-1} g_l X_{i+k, j+l} \quad (27)$$

for $k = -p+1, \dots, p-1$. It is easy to check that

$$[\mathbf{C}^t \mathbf{C} \mathbf{x}]_{ij} = \sum_{k=-(p-1)}^{p-1} w_k R_k^{ij} = \mathbf{w}^t \mathbf{R}^{ij}. \quad (28)$$

As autocorrelation vectors, \mathbf{g} and \mathbf{w} are symmetric. Thus, (27) and (28) only involve q and p multiplications, respectively. The point in introducing this auxiliary vector is the shift invariance $R_k^{i+1, j} = R_{k+1}^{ij}$, which is a straightforward consequence of the definition of R (27). Therefore, during the scan of a whole column, for all $k < p-1$, $R_k^{i+1, j} = R_{k+1}^{ij}$ was already computed for the previous voxel. Indeed, after each update of a voxel x_s , one only has to update R_0^{ij} accordingly and to calculate $R_{p-1}^{i+1, j}$.




Fig. 8. Slice #22 from the test data set (127×127 images). (a) Slice of the simulated 3-D object. (b) Filtered back-projection 2-D reconstruction. (c) Gaussian MRF 3-D reconstruction. (d) Half-quadratic MRF 3-D reconstruction. (e) Half-quadratic MRF 2-D restoration. (f) Half-quadratic MRF 3-D restoration.

Finally, this provides a very efficient way of updating the voxels by recursive scanning of the columns, as described in Table II (a transposed form can also be derived).

V. RESULTS

A. Simulated Images

Our first concern was to validate the restoration approach against the reconstruction approach. In Section I, we advocated the advantages of restoration in terms of portability and ease of use for clinical applications. However, quantitative comparison with existing reconstruction methods is necessary to assess the interest of our approach.

In order to investigate the performance of the selected methods in reasonably realistic 3-D configurations, we simulated a 3-D object intended to roughly approximate the distal part of a femur with simple geometric features while introducing significant estimation difficulties (smooth areas, edges, grooves, convex and concave parts). The resulting design is presented in Fig. 5(a). From this model, slices were computed in 35 planes evenly spaced every 2 mm along its main axis and digitized into 127×127 images. For each slice, 90 projections using 141 parallel rays were then evaluated along 90 evenly spaced angles. Finally, Gaussian noise with a standard deviation of 3.0 was added to the projections; this corresponds to a signal-to-noise ratio of approximately 25 dB,

where the signal-to-noise ratio is defined as the ratio between the power of the noise and the power of the centered noiseless projection data. This data set actually represented a whole scan of the object and was the common starting point of the three reconstruction methods that we selected as follows.

- *Filtered back-projection (FBP 2-D recon.)*: This technique is very fast and can be considered the standard since it is implemented in most commercial CT scanners. In our experiments, the filter cutoff frequency was set to half the projection sampling frequency.
- *Inversion of the projection operator with 3-D quadratic regularization or equivalently 3-D Gaussian priors (Gauss 3-D recon.)*: Such a technique can be implemented very efficiently, but tends to smooth out the images, which can blur the edges and reduce the actual resolution.
- *Inversion of the projection operator with 3-D priors identical to the one presented in Section II-C (half-quad. 3-D recon.)*.

Our restoration method was then applied to the images most resembling those produced by a standard CT scanner, i.e., images obtained by the filtered back-projection method. In all simulations, the PSF was a 2-D Gaussian function whose characteristics were derived from the simulated reconstructions of point objects. For comparison purposes, both 2-D (half-quad. 2-D




Fig. 9. Slice #32 from the test data set (127×127 images). (a) Slice of the simulated 3-D object. (b) Filtered back-projection 2-D reconstruction. (c) Gaussian MRF 3-D reconstruction. (d) Half-quadratic MRF 3-D reconstruction. (e) Half-quadratic MRF 2-D restoration. (f) Half-quadratic MRF 3-D restoration.

resto.) and 3-D (half-quad. 3-D resto.) approaches were tested. For all half-quadratic methods, the *a priori* model was either a 2-D or 3-D version of the MRF described in Section II-C. The potential φ selected for our experiments is quadratic toward zero and linear for large values (Fig. 6). Such potentials have been found to present a good tradeoff between convexity and edge preservation. Several functions with the same general behavior have been proposed in the literature (see, e.g., [10]) and they all produce quite similar results. The following form, which obviously fulfills conditions (13)–(15), was implemented:

$$\varphi(u) = \sqrt{1 + \left(\frac{u}{\delta}\right)^2} \quad (29)$$

where δ is a scaling factor determining the transition between the quadratic and linear regions. In our experiments, the values of *tuning parameters* δ and $\sigma_n^2 \lambda$ [see (9)] were determined empirically so as to produce the best visual results ($\delta = 1.0 \times 10^{-3}$ and $\sigma_n^2 \lambda = 8.75 \times 10^{-5}$ in 2-D and 6.65×10^{-5} in 3-D).

With this experimental setup, each slice of the original 3-D object was known. For comparison purposes, the normalized quadratic error between theoretical and reconstructed or restored slices was evaluated. This normalized quadratic error between a given experimental result x and the original object x^{ref} can be evaluated either in 2-D for each slice (e_{Q2D}) or

in 3-D for the whole set of slices (e_{Q3D}) according to the following expression:

$$e_{Q3D}(x) = \sum_{s \in \mathcal{S}} (X_s - X_s^{\text{ref}})^2, \quad \text{where } X_s = \frac{x_s}{\sqrt{\sum_{m \in \mathcal{S}} x_m^2}}. \quad (30)$$

The computed error values for all tested methods are summarized in Table III and Fig. 7. As expected, the lowest error values are obtained with the half-quadratic 3-D reconstruction method. However, the results from half-quadratic reconstruction and restoration are very close to each other both numerically (same order of magnitude) and visually. The behavior of the two techniques is illustrated in Fig. 8, where both of them produce similar results, and in Fig. 9, where half-quadratic 3-D reconstruction performs noticeably better than half-quadratic 3-D restoration. Nonetheless, this comparison validates the restoration approach when access to projection data is hardly possible.

The higher numerical complexity of half-quadratic methods with respect to quadratic approaches is fully justified when looking at the images in Figs. 8 and 9. With Gaussian reconstruction, substantial noise reduction could be achieved only at the expense of smoothing out the edges in an unacceptable way. Arguably, this poor result is partly due to the fact that the noise level used in our simulations may be higher than in typical CT images. This choice was made in order to discriminate

TABLE IV
 NORMALIZED QUADRATIC ERRORS WHEN PSF SIZE IS CHANGED (SIMULATED DATA SET, 35 127×127 SLICES)

Method	PSF size	e_{Q3D}	$\min(e_{Q2D})$	$\text{mean}(e_{Q2D})$	$\max(e_{Q2D})$
2D resto.	0.25 mm	7.27×10^{-3}	5.99×10^{-3}	7.66×10^{-3}	1.01×10^{-2}
2D resto.	0.5 mm	6.76×10^{-3}	5.33×10^{-3}	7.18×10^{-3}	9.86×10^{-3}
2D resto.	1 mm	7.08×10^{-3}	5.51×10^{-3}	7.50×10^{-3}	1.04×10^{-2}
3D resto.	0.25 mm	6.47×10^{-3}	4.26×10^{-3}	6.72×10^{-3}	9.89×10^{-3}
3D resto.	0.5 mm	5.94×10^{-3}	3.80×10^{-3}	6.23×10^{-3}	9.09×10^{-3}
3D resto.	1 mm	6.24×10^{-3}	4.09×10^{-3}	6.53×10^{-3}	9.52×10^{-3}

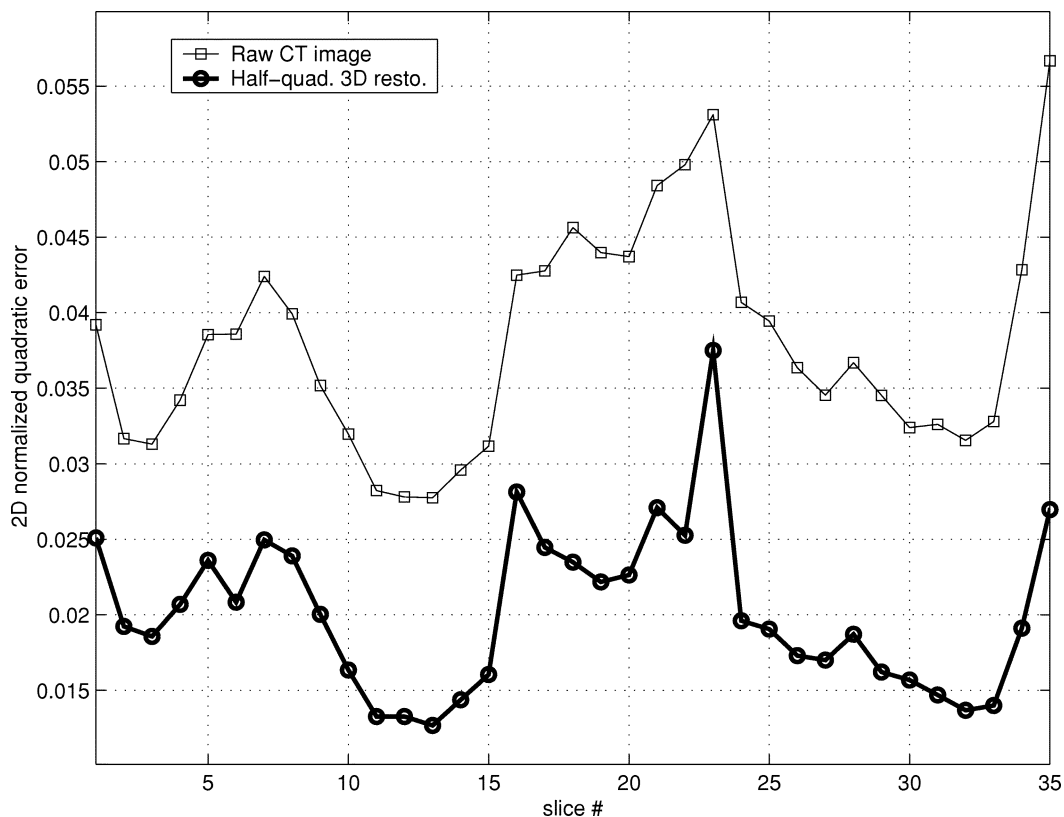


Fig. 10. Graph of 2-D normalized quadratic errors e_{Q2D} for the proposed restoration method applied to CT images of the PVC phantom in water.

better between the various methods, and our results illustrate the ability of half-quadratic potentials to produce a high level of regularization without smoothing out the edges. Therefore, such techniques are very well suited to the reconstruction or restoration of CT images.

Another issue is the comparison between 2-D and 3-D restorations. On the one hand, the overhead for 3-D processing amounts to approximately 20% of the total computing time. On the other hand, although the restored images in 2-D and 3-D look quite similar (see Fig. 8 and 9), the error is noticeably lower in 3-D. This improvement can be measured by the overall 3-D error (-12%) and by the mean 2-D error (-28%). However, the results are very variable from one slice to another. By comparing the profile of the phantom [see Fig. 5(b)] with the error curve (Fig. 7), one can notice that the 3-D approach is particularly efficient when two successive slices are similar (slices 1–5 and 20–25). This can be interpreted as the effect of regularization along the z -axis, which is included in the 3-D MRF model, but not in the 2-D MRF model.

The dependency of the results with respect to the size of the PSF was also tested by running the same examples with two other PSF sizes (double or half the size estimated from point objects). The results summarized in Table IV show that, even with such a large range of variation, the performances of restoration are only slightly disturbed (around 10% increase of 2-D error, even less for 3-D error). Visual estimation of the size of the PSF (from sharp edges or bright spots) would be more precise than this and, thereby, would enable accurate restoration in practical cases.

B. Real CT Images

Our restoration method was also applied to real CT data. The images were obtained with a Picker PQ 5000 tomograph used in planar mode on a phantom made from the 3-D design already described in Section V-A [see Fig. 5(a)]. Among easy-to-process materials, polyvinyl chloride (PVC) was selected because its opacity to X-rays is very close to that of bone tissue.




Fig. 11. Samples from the CT scan of the PVC phantom in water (256×256 images): slice #23 with an air bubble and slice #26 with a zoom to show the sharpness of the edges. (a) Raw CT image. (b) Restored image with half-quadratic 3-D MRF. (c) Ground truth from the phantom design.

The phantom was placed in water in order to approach X-ray opacity of soft tissues. The CT scanner was used in planar acquisition mode and the reconstruction settings were set to *bone* because this mode is designed for orthopedic applications and because it gives the sharpest images (at the expense of a slightly higher level of noise than the *smooth* or *standard* settings). Precise alignment of the phantom, elimination of positioning ambiguities, and monitoring of the accuracy of the table movements were made possible by the presence of two grooves perpendicular to the main axis of the phantom [see Fig. 5(a)]. This setup allowed us to compare theoretical phantom slices with experimentally reconstructed ones.

The values for the normalized 2-D error e_{Q2D} are plotted in Fig. 10 and sample images are shown in Fig. 11 ($\delta = 1$ and $\sigma_n^2 \lambda = 80$). The noise is almost completely suppressed and, as expected, the edges are preserved. On the average, the error is almost divided by a factor of two between the raw CT images and restored images, except when air bubbles are stuck

in one of the phantom corners, as illustrated in Fig. 11. It can be observed that the improvement is not as large as with simulated images, which may be explained by at least two reasons. First, there is always a residual error due to imperfect registration between the real CT images and the 3-D model, as well as imperfect dimensions of the actual PVC phantom. Second, in real CT images, nonlinear distortions appear, such as partial volume effects that blur the edges when the size of slices change rapidly, which could explain the larger errors in slices 16–23. 3-D half-quadratic restoration partly corrects for these deficiencies and, therefore, may appear as a flexible and inexpensive way of improving the quality of standard CT images for dedicated purposes.

VI. CONCLUSION

In this paper, we have proposed a method to enhance CT images in order to increase the accuracy of edge localization. We

have chosen to use image restoration (or deconvolution) of CT images rather than reconstruction from X-ray projections essentially for practical reasons. Indeed, in most cases, the practitioner only has access to already reconstructed (and filtered) CT images. In order to estimate the PSF of the imaging system, one can use a simple phantom made with enamel wires, as presented in Fig. 4. However, since restoration results are not very sensitive to the exact parameters of the PSF, as shown on simulated examples in Table IV, rough estimation from edges and bright spots of CT images may be sufficient.

The restoration method that we have proposed implements half-quadratic regularization because it allows to use convex edge-preserving potentials that exhibit quite an interesting tradeoff between image quality and computational efficiency. On the one hand, experimental results show that the edges are indeed very much sharper in the restored images. On the other hand, since the bulk of the computations occurs in the convolution term depending on the PSF size, simplifying assumptions, such as the PSF separability, have to be made to decrease the computational burden significantly.

Reducing the numerical complexity was the key factor that allowed us to use a 3-D MRF as an *a priori* model. The clique structure was simplified as much as possible so as to make the algorithm easier to interpret and to tune without jeopardizing image quality. Significantly better results were obtained with a 3-D prior model than with a 2-D MRF because of regularization along the three directions. Using 3-D processing, CT images from the whole scan of a patient can be manipulated as a unique 3-D volume and enhanced as a whole. The tuning parameters can be set once for the whole 3-D object, thereby simplifying the use of the method. Further investigations aiming at making the technique fully unsupervised are currently under way. Other avenues of research encompass extension of these techniques to spiral and multislice CT scanners, which are becoming prevalent in clinical settings.

REFERENCES

- [1] S. Doré and Y. Goussard, "Experimental determination of CT point spread function anisotropy and shift-variance," in *Proc. 19th Annu. Int. IEEE Engineering in Medicine and Biology Society Conf.*, Chicago, IL, Oct. 1997, pp. 788–791.
- [2] D. Geman and C. Yang, "Nonlinear image recovery with half-quadratic regularization," *IEEE Trans. Image Processing*, vol. IP-4, pp. 932–946, July 1995.
- [3] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud, "Deterministic edge-preserving regularization in computed imaging," *IEEE Trans. Image Processing*, vol. 6, pp. 298–311, Feb. 1997.
- [4] A. H. Delaney and Y. Bresler, "Globally convergent edge-preserving regularized reconstruction: An application to limited-angle tomography," *IEEE Trans. Image Processing*, vol. 7, pp. 204–221, Feb. 1998.
- [5] A. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. New York: IEEE Press, 1987.
- [6] B. Hunt, "The application of constrained least squares estimation to image restoration by digital computer," *IEEE Trans. Commun.*, vol. COM-22, pp. 805–812, Sept. 1973.
- [7] J. E. Besag, "On the statistical analysis of dirty pictures (with discussion)," *J. Roy. Statist. Soc. B*, vol. 48, no. 3, pp. 259–302, 1986.
- [8] A. Blake and A. Zisserman, *Visual Reconstruction*. Cambridge, MA: MIT Press, 1987.
- [9] S. Geman and G. Reynolds, "Constrained restoration and recovery of discontinuities," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 14, pp. 367–383, Mar. 1992.
- [10] C. Bouman and K. Sauer, "A generalized Gaussian image model for edge-preserving MAP estimation," *IEEE Trans. Image Processing*, vol. 2, pp. 296–310, July 1993.
- [11] S. Brette and J. Idier, "Optimized single site update algorithms for image deblurring," in *Proc. Int. Image Processing Conf.*, Lausanne, Switzerland, 1996, pp. 65–68.
- [12] J. E. Besag, "Spatial interaction and the statistical analysis of lattice systems (with discussion)," *J. Roy. Statist. Soc. B*, vol. 36, no. 2, pp. 192–236, 1974.
- [13] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-6, pp. 721–741, Nov. 1984.
- [14] F. Beckman, "The solution of linear equations by the conjugate gradient method," in *Mathematical Methods for Digital Computers*, A. Ralston, H. Wilf, and K. Enslein, Eds. New York: Wiley, 1960.
- [15] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical Recipes in C, the Art of Scientific Computing*, 2nd ed. New York: Cambridge Univ. Press, 1992.
- [16] D. M. Young, "Iterative solution of large linear system," in *Computer Science and Applied Mathematics*. New York: Academic, 1971.
- [17] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud, "Two deterministic half-quadratic regularization algorithms for computed imaging," in *Proc. IEEE Int. Conf. Image Processing*, vol. 2, 1994, pp. 168–172.
- [18] J. Idier, "Convex half-quadratic criteria and interacting auxiliary variables for image restoration," *IEEE Trans. Image Processing*, vol. 10, pp. 1001–1009, July 2001.
- [19] D. Luenberger, *Optimization by Vector Space Methods*, 1st ed. New York: Wiley, 1969.
- [20] R. T. Rockafellar, *Convex Analysis*. Princeton, NJ: Princeton Univ. Press, 1970.
- [21] M. Allain, "Approche pénalisée en tomographie hélicoïdale. Application à la conception d'une prothèse personnalisée du genou," Ph.D. dissertation, Univ. Paris-Sud, Orsay, France, 2002.
- [22] C. A. Bouman and K. D. Sauer, "A unified approach to statistical tomography using coordinate descent optimization," *IEEE Trans. Image Processing*, vol. 5, pp. 480–492, Mar. 1996.