

Efficient binary reconstruction for non-destructive evaluation using gammagraphy

Marc Allain^{1,3} and Jérôme Idier²

¹ Institut Fresnel, Domaine Universitaire de St-Jérôme, 13397 Marseille Cedex 20, France

² Institut de Recherche en Communications et en Cybernétique de Nantes (IRCCyN), 1 rue de la Noë, BP 92101, 44321 Nantes Cedex 03, France

E-mail: marc.allain@fresnel.fr and Jerome.Idier@irccyn.ec-nantes.fr

Received 15 January 2007, in final form 3 May 2007

Published 13 June 2007

Online at stacks.iop.org/IP/23/1371

Abstract

The localization and the sizing of 3D flaws within a homogeneous metallic medium is a major task for non-destructive evaluation (NDE). This paper addresses the problem of the reconstruction of such flaws using an efficient binary algorithm. Basically, the method rests on the fact that a simple binary constraint suffices for an accurate and robust reconstruction in the context of NDE. A heuristic minimization, computationally attractive, is designed in order to provide fast reconstructions. The proposed algorithm is compared with standard binary (the *iterated conditional mode* algorithm) and non-binary (penalized approach with convex potential Gibbs random fields) reconstruction techniques.

1. Introduction

Reconstruction of binary volumes from noisy projections is an attractive tool for non-destructive evaluation (NDE) of metal structures [2, 3]. It amounts to the simultaneous detection and sizing of small defects in a homogeneous material. However, such a tomographic reconstruction problem is challenging when only *limited-angle* projections are available, since it is then very hard to recover the geometry of the defects along the average direction of the rays. As far as *in situ* NDE is considered, this unfavourable observation context is often a direct consequence of the operational constraints imposed on the NDE setup. In [1], for instance, the radiographic inspection of in-service pipes in the nuclear industry is considered with a total angular excursion of 30°, see figure 1.

In limited-angle tomography, faithful reconstructions of the 3D attenuation map cannot be expected without adequate prior constraints on the solution [4–6]. One possible approach is to adopt a parametric model for the shape of the defect (for instance, that it is spherical),

³ Earlier with the Institut de Recherche en Communications et en Cybernétique de Nantes (IRCCyN), 1 rue de la Noë, BP 92101, 44321 Nantes Cedex 03, France.




Figure 1. Left: 3D representation of the instrumental setup. Right: schematic representation of the measurement context for the synthetic problem considered in this paper. Note that all the dimensions are drawn from a real NDE setup in the nuclear industry, cf [1, figure 1].

and to estimate the few corresponding parameters, say, in the least-square sense [7]. However, this approach relies on strong assumptions, which are not natural in the NDE context, where the number and/or the shape of the hypothetical defects are not known. In [8], a genetic-algorithm strategy is proposed to handle an unknown number of defects, but the constraint of a parametrized shape remains. A potentially more flexible extension is obtained using a description of objects in terms of *contours*. Again, parametric contour-based methods [9, 10] hardly handle the case of an unknown number of contours, at completely unknown positions. Fully non-parametric modelling of the contour via a *level-set* formulation [11] yields a more suited approach. This leads to gradient-based iterations that aim to compute a local minimizer of the 3D partitioning problem [12].

A large part of the literature rather assumes a voxel decomposition for the attenuation map. Then, the tomographic reconstruction \hat{x} is usually defined as the minimizer of a penalized least-square criterion, such as

$$\begin{cases} \min_x \|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2 + \psi(\mathbf{x}; \boldsymbol{\theta}) & (1a) \\ \text{subject to } \mathbf{x} \in \mathbb{X} \subseteq \mathbb{R}^N, & (1b) \end{cases}$$

where $\boldsymbol{\theta}$ is a vector of hyperparameters. Within this framework, the solution \hat{x} is clearly a trade-off: whereas the first term favours the solutions that fit the data \mathbf{y} , the second one discards the solutions that are not compatible with prior information.

A usual choice is to consider ψ as the energy of a Markov random field with positive interactions, in order to favour clusters of voxels in the reconstruction [13]. More specifically, the well-known *Ising model* is of special interest here, since the spatial information to encode is binary: each voxel either belongs to an inclusion (of negligible attenuation), or to a sound region of the inspected object (of known attenuation). Only a few works have been based on the Ising model to tackle NDE issues. Among the notable exceptions let us mention [3], where it clearly appears that a major numerical obstacle must then be faced, of combinatorial nature, and that no simple, satisfying solution is available. Several authors have rather chosen to relax the binary constraint and to cope with NDE in the framework of continuous-valued Markovian reconstruction [2, 1]. The resulting problem is easier to deal with from the numerical viewpoint, since gradient-based optimization algorithms are then available. Unfortunately, the regularizing character of the binary constraint is lost, so that the reconstructed defects appear

stretched along some directions [1]. Let us remark here that the resort to Markov models strengthens such a tendency to produce oversized defects. This motivates our choice of an independent, identically distributed (iid) binary model for the voxels: the binary constraint is then the only regularizing ingredient, while the iid assumption corresponds to the simplest possible spatial model. Although the resulting optimization problem is still intractable under an exact form, it lends itself to a simple and satisfactory solution. Its main features are as follows.

- (i) Firstly, a *region of interest* (ROI) is determined through straightforward computations, such that the maximum posterior mode of voxels outside the ROI is zero, i.e., $\Pr(X_n = 0 \mid \mathbf{y}) > \Pr(X_n = 1 \mid \mathbf{y})$. This results in a dramatic reduction in the dimensionality of the problem.
- (ii) Secondly, we rely on a modified version of the *single most likely replacement* algorithm (SMLR) introduced in [14] to explore efficiently the binary configurations in the ROI. Whereas the original algorithm considers a ‘single site’ update in the volume, the *block most likely replacement* (BMLR) proposed here considers ‘blockwise’ updates within cubes of $2 \times 2 \times 2$ voxels. For limited-angle reconstructions, it is empirically observed that such a strategy allows us to escape from local minima where the SMLR would be trapped.

Based on a realistic synthetic NDE problem, our numerical procedure is shown to be very efficient, both in terms of computational load and accuracy of a reconstructed attenuation map. Moreover, the proposed method does not depend on any hyperparameter in contrast with many high-resolution tomographic reconstruction methods.

The paper is organized as follows. In section 2, the modelling of the tomographic measurement in NDE is introduced, and several Bayesian estimators are envisaged to solve the reconstruction problem. This section ends with some remarks concerning the computation of these estimators that motivate the design of a computationally efficient heuristic. Section 3 deals with the construction of the ROI that allows a reduction of the voxels involved in the reconstruction process. In section 4, the problem of binary object reconstruction within a ROI is stated. Standard binary reconstruction algorithms are introduced and the BMLR is derived. In section 5, the performances of the BMLR are compared to the standard *iterated conditional mode* algorithm (ICM) and to a penalized approach with convex potential Gibbs random fields. Finally, a discussion of this work and some conclusions are presented in section 6.

2. Binary constraint for NDE

Let $\mathcal{S} = \{1, \dots, N\}$ be an arbitrary enumeration of the N voxels that constitute the unknown three-dimensional attenuation map \mathbf{x} . Under usual assumptions, the tomographic measurement along the m th ray follows the Beer–Lambert relation [15, section 4.1.1]: for all $m = 1, \dots, M$,

$$y_m \stackrel{\text{def}}{=} -\log \frac{\tau_m}{\tau_s} = \mathbf{h}_m^t \mathbf{x} + e_m, \quad (2)$$

where $\mathbf{h}_m \in \mathbb{R}_+^N$ stands for the projection process along the m th ray (see for instance [15, section 7.1]), τ_s and τ_m are the photon emission rate of the source and the counting rate of detector m , respectively. e_m is a noise term that accounts for the statistical fluctuations in the measurement. Stacking all the measurements into a vector $\mathbf{y} \stackrel{\text{def}}{=} (y_m)$ yields the following observation model:

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{e}, \quad (3)$$

where \mathbf{H} is a matrix of size $M \times N$, so that \mathbf{h}_m^t is the m th line of \mathbf{H} .

In this paper, it will be assumed that both the geometry and composition of the inspected object are known, up to the possible presence of defects, of course. In more precise terms, we restrict ourselves to the case where the object would be perfectly known if it were free of any defect. Let $\mathbf{r} \in \mathbb{R}_+^N$ denote the corresponding defect-free attenuation map, and let us introduce a binary map $\tilde{\mathbf{x}} \in \{0, 1\}^N$ to encode the geometry of the defects: voxels that belong to a defect are assigned the value ‘1’, whereas the others are assigned the value ‘0’. Given that the defects are assumed to be composed of air, of negligible attenuation, the actual attenuation map \mathbf{x} is made of voxels such that

$$x_m = \begin{cases} 0 & \text{if } \tilde{x}_m = 1, \\ r_m & \text{otherwise,} \end{cases}$$

for all $m = 1, \dots, M$. In compact notation, we have $\mathbf{x} = \mathbf{r} - \text{diag}(\mathbf{r})\tilde{\mathbf{x}}$, so that a new linear observation model can be deduced from (3):

$$\tilde{\mathbf{y}} = \tilde{\mathbf{H}}\tilde{\mathbf{x}} + \tilde{\mathbf{e}},$$

where $\tilde{\mathbf{y}} = \mathbf{H}\mathbf{r} - \mathbf{y}$, $\tilde{\mathbf{H}} = \mathbf{H}\text{diag}(\mathbf{r})$ and $\tilde{\mathbf{e}} = -\mathbf{e}$. Thus, the reconstruction problem can be reformulated as a linear inverse problem with binary unknowns. In the rest of the paper, the tilde signs will be omitted for the sake of notational simplicity, so that the attenuation map itself will be considered as binary.

In the following, \mathbf{e} is assumed to be an outcome of a normal random vector \mathbf{E} with zero mean and iid components

$$\mathbf{E} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}), \quad (4)$$

where \mathbf{I} is the identity matrix and $\sigma > 0$ is the standard deviation of the noise. Provided that the counting rates at each detector are high and of the same magnitude, the Gaussian additive observation model is of sufficient accuracy. These assumptions are realistic in NDE and the Gaussian model will be adopted in this paper. For a low signal-to-noise ratio (SNR), however, a Poisson model is more appropriate to describe the random fluctuations in the measurements, see [16] for details.

Under assumptions (3) and (4), the likelihood of the observation vector reads

$$f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y} | \mathbf{x}) \propto \exp \left\{ -\frac{1}{2\sigma^2} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2 \right\}, \quad (5)$$

where ‘ \propto ’ stands for ‘proportional to’. Considering the Bayesian framework, a simple iid model for the prior distribution is introduced:

$$\Pr(\mathbf{X} = \mathbf{x}) = \prod_{n=1}^N \Pr(X_n = x_n) = p_1^k p_0^{N-k},$$

where $\forall n \in \mathcal{S}$, $\Pr(X_n = 1) = p_1$, $\Pr(X_n = 0) = p_0 = 1 - p_1$ and k is the number of voxels that are equal to ‘1’. Then the posterior probability reads

$$\Pr(\mathbf{X} = \mathbf{x} | \mathbf{y}) = \frac{f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y} | \mathbf{x})\Pr(\mathbf{X} = \mathbf{x})}{\sum_{\mathbf{x}' \in \mathbb{S}} f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y} | \mathbf{x}')\Pr(\mathbf{X} = \mathbf{x}')}, \quad (6)$$

where $\mathbb{S} = \{0, 1\}^N$ is the set of binary configurations. The *maximum a posteriori* (MAP) solution

$$\hat{\mathbf{X}}^{\text{MAP}} \stackrel{\text{def}}{=} \arg \max_{\mathbf{x} \in \mathbb{S}} \Pr(\mathbf{X} = \mathbf{x} | \mathbf{y}) \quad (7)$$

amounts to minimizing the following quadratic criterion:

$$\mathcal{J}(\mathbf{x}) \stackrel{\text{def}}{=} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2 + 2\sigma^2 \mu \mathbf{1}^t \mathbf{x} \quad (8)$$

under the binary constraint $\mathbf{x} \in \mathbb{S}$, where $\mathbf{1} \stackrel{\text{def}}{=} (1, \dots, 1)^t \in \mathbb{R}^N$ and $\mu \stackrel{\text{def}}{=} \log(p_0/p_1)$.

Beside the MAP solution, the Bayesian setting allows us to define alternate estimators. In particular, the *posterior mean* of voxel n reads

$$E[X_n | \mathbf{Y}] = \sum_{x_n \in \{0,1\}} x_n \Pr(X_n = x_n | \mathbf{y}) = P_n,$$

where

$$P_n \stackrel{\text{def}}{=} \Pr(X_n = 1 | \mathbf{y}).$$

Following [17], let us also define the *marginal posterior mode* (MPM) estimator according to

$$\widehat{X}_n^{\text{MPM}} \stackrel{\text{def}}{=} \arg \max_{x_n \in \{0,1\}} \Pr(X_n = x_n | \mathbf{y}), \quad (9)$$

which is easy to deduce from P_n :

$$\widehat{X}_n^{\text{MPM}} = \begin{cases} 1 & \text{if } P_n > 1/2 \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

Unfortunately, none of the latter solutions are easily computable. In particular, (7) is a *binary quadratic programming problem*. In the terms and notation of statistical physics, it amounts to minimizing the following energy:

$$H(\mathbf{s}) = - \sum_{(m,n) \in \mathcal{E}} J_{mn} s_m s_n + \sum_{n \in \mathcal{S}} F_n s_n, \quad (11)$$

where $s_n = 2x_n - 1 \in \{-1, +1\}$ are Ising spins,

$$J_{mn} = -\mathbf{h}_{\cdot m}^t \mathbf{h}_{\cdot n}, \quad \mathcal{E} = \{(m, n) \in \mathcal{S}^2, m < n, \mathbf{h}_{\cdot m}^t \mathbf{h}_{\cdot n} \neq 0\}$$

is the set of interacting spins, and

$$F_n = -\mathbf{h}_{\cdot n}^t \left(2\mathbf{y} - \sum_{m \in \mathcal{S}} \mathbf{h}_{\cdot m} \right) + 2\sigma^2 \mu$$

plays the role of the n th component of an exterior magnetic field. In some particular cases, like the one encountered in binary image denoising [18], solving such a problem can be done in an acceptable number of operations, that is a polynomial function of N . In contrast, our case falls within the class of NP-hard problems, for which it is very unlikely that a polynomial algorithm exists [19]. The first reason for NP-hardness is the fact that J_{mn} takes negative values (i.e., the interactions are antiferromagnetic) in a nonzero magnetic field. Moreover, the non-planarity of the neighbourhood graph \mathcal{E} is another source of NP-hardness [20].

On the other hand, computing $\widehat{X}_n^{\text{MPM}}$ requires the computation of partition functions involving summations over \mathbb{S} or large subsets of \mathbb{S} . More specifically, the evaluation of the key quantity P_n can be envisaged as follows. Let $\mathbb{S}_n^i \stackrel{\text{def}}{=} \{\mathbf{x} \in \mathbb{S} : x_n = i\}$, so that

$$P_n = \sum_{\mathbf{x} \in \mathbb{S}_n^1} \Pr(\mathbf{X} = \mathbf{x} | \mathbf{y}) = \frac{\sum_{\mathbf{x} \in \mathbb{S}_n^1} \Pr(\mathbf{X} = \mathbf{x} | \mathbf{y})}{\sum_{\mathbf{x} \in \mathbb{S}} \Pr(\mathbf{X} = \mathbf{x} | \mathbf{y})} = \frac{1}{1 + R_n(\mathbf{y})},$$

where

$$R_n(\mathbf{y}) \stackrel{\text{def}}{=} \frac{\sum_{\mathbf{x} \in \mathbb{S}_n^0} \Pr(\mathbf{X} = \mathbf{x} | \mathbf{y})}{\sum_{\mathbf{x} \in \mathbb{S}_n^1} \Pr(\mathbf{X} = \mathbf{x} | \mathbf{y})}. \quad (12)$$

Again, the computation of the sums entering the expression of R forms NP-hard problems [19, 20].

Despite the NP-hardness of the problem, some contributors resort to Markov chain–Monte Carlo (MCMC) techniques to solve such combinatory problems. In [17], a Gibbs sampler is used to compute the posterior mean estimator for binary denoising problems. In [21], the authors resort to simulated annealing in order to minimize a least-square criterion under binary constraint. By nature, such MCMC techniques converge in the asymptotic limit. When strong spatial correlations exist in the posterior probability, efficient sampling is difficult and the resulting algorithms may ask for unrealistic computation time before a useful estimate is available. In practice, we found that this problem is a major obstacle in designing MCMC-based techniques for binary reconstruction in limited-angle tomography. Hence, our main concern in this paper is to develop an efficient heuristic that computes a local maximizer of $\Pr(\mathbf{X} = \mathbf{x} \mid \mathbf{y})$ in a finite number of iterations. For this purpose, our strategy is twofold. Firstly, we aim at reducing the dimensionality of the problem by an appropriate selection of the voxels that should be discarded in the reconstruction process, see section 3. Then, a local maximizer of $\Pr(\mathbf{X} = \mathbf{x} \mid \mathbf{y})$ is computed via a deterministic iterative procedure, see section 4.

3. Region-of-interest (ROI) selection

As far as large 3D problems are considered, a first step towards a computationally attractive algorithm is to reduce the number of voxels involved in the reconstruction. As shown in the following section, it can be established that $\widehat{\mathbf{X}}_n^{\text{MPM}} = \mathbf{0}$ for a large amount of voxels, from costless preprocessing operations involving the data.

3.1. MPM based ROI in the Gaussian data case

Although P_n is not straightforward to calculate, it admits an easily computable upper bound \overline{P}_n that will be fruitful to manipulate.

Let $\mathbf{0}$ denote the null vector of length N and \mathbf{e}_n the n th canonical vector of \mathbb{R}^N . Indeed, the following proposition holds in the case of centred white Gaussian observation noise.

Proposition 1. *The posterior probability P_n admits the following upper bound:*

$$P_n \leq \overline{P}_n \stackrel{\text{def}}{=} \frac{1}{1 + \underline{R}_n(\mathbf{y})} \quad (13)$$

where

$$\underline{R}_n(\mathbf{y}) \stackrel{\text{def}}{=} \frac{\Pr(\mathbf{X} = \mathbf{0} \mid \mathbf{y})}{\Pr(\mathbf{X} = \mathbf{e}_n \mid \mathbf{y})} \quad (14)$$

$$= \exp(-\overline{\Delta}_n), \quad (15)$$

with

$$\overline{\Delta}_n \stackrel{\text{def}}{=} \frac{1}{2\sigma^2} (2\mathbf{h}_{\cdot n}^t \mathbf{y} - \|\mathbf{h}_{\cdot n}\|^2) - \mu. \quad (16)$$

Proof. See appendix A. □

Interestingly, it is clear from (10) that

$$\overline{P}_n \leq 1/2 \implies \widehat{\mathbf{X}}_n^{\text{MPM}} = \mathbf{0}. \quad (17)$$

The computation of \bar{P}_n plays a pivotal role in our method since it allows a significant reduction of the effective dimension of the reconstruction problem. Let us introduce the set \mathcal{I} of voxels for which implication (17) is inactive, i.e.,

$$\mathcal{I} = \{n \in \mathcal{S} : \bar{P}_n > 1/2\}.$$

The latter set defines a ROI in the sense that the reconstruction algorithm will be restricted to \mathcal{I} , while the other voxels will be assumed to match their MPM estimation, i.e., $\hat{X}_n^{\text{MPM}} = 0, \forall n \in \mathcal{S} \setminus \mathcal{I}$.

Let us remark that

$$\begin{aligned} \bar{P}_n \leq 1/2 &\iff \bar{\Delta}_n \leq 0 \\ &\iff \mathbf{h}_{,n}^t \mathbf{y} \leq \frac{1}{2} \|\mathbf{h}_{,n}\|^2 + \sigma^2 \mu, \end{aligned} \quad (18)$$

so that a threshold test on the backprojection $\mathbf{H}^t \mathbf{y}$ of the data allows us to determine \mathcal{I} , where the threshold value is generally not spatially invariant. Moreover, the test does not depend on the noise variance provided that the binary states are considered equally probable, i.e., $\mu = 0$.

A direct extension of proposition 1 can be established in the general Gaussian case of a noise with mean \mathbf{m} and covariance \mathbf{R} . A less trivial but richer extension is provided in subsection 3.3 to encompass some cases of log-concave likelihoods.

Finally, let us note that proposition 1 admits a formal counterpart, which yields a sufficient condition to ensure $\hat{X}_n^{\text{MPM}} = 1$. The latter condition is hardly met in our situation, because too few voxels are equal to 1, but it would be interesting to implement condition (18) in more balanced situations. The proof is omitted, since it is almost a paraphrase of that of proposition 1.

Proposition 2. *The posterior probability P_n admits the following lower bound:*

$$P_n \geq \underline{P}_n \stackrel{\text{def}}{=} \frac{1}{1 + \bar{R}_n(\mathbf{y})},$$

where

$$\bar{R}_n(\mathbf{y}) \stackrel{\text{def}}{=} \frac{\Pr(\mathbf{X} = \mathbf{1} - \mathbf{e}_n \mid \mathbf{y})}{\Pr(\mathbf{X} = \mathbf{1} \mid \mathbf{y})} = \exp(-\underline{\Delta}_n),$$

with

$$\underline{\Delta}_n \stackrel{\text{def}}{=} \frac{1}{2\sigma^2} (2\mathbf{h}_{,n}^t (\mathbf{y} - \mathbf{H}\mathbf{1}) + \|\mathbf{h}_{,n}\|^2) + \mu.$$

3.2. Synthetic example

For illustrative purposes, let us compute the threshold value in a simple, yet realistic context. The inspected object is made of a homogeneous material, with two homogeneous defects of spherical shape. Its discrete representation contains $64 \times 64 \times 64$ voxels, while each of the two defects spreads over 32 voxels. Some schematic illustrations of the measurement context are provided in figure 1. The data \mathbf{y} consist in seven 2D projections (each one being a 128×128 image) that were corrupted by zero mean iid Gaussian perturbations with standard deviation $\sigma = 0.01$, which is close to the maximum amplitude of the projection of a single voxel. Both the ROI \mathcal{I} and the backprojection were deduced from this data set. The former was computed with the additional assumption that $p_1 = p_0$ (i.e., $\mu = 0$).

Figures 2(left) and (right) depict the ‘central’ 2D cross-section of the backprojection $\mathbf{H}^t \mathbf{y}$ and of the upper bound \bar{P}_n , respectively.




Figure 2. Central 2D section of the backprojection (left) and of the map of the upper bound \bar{P}_n (right), respectively, obtained from noisy projections ($\sigma = 0.01$), for a pair of superimposed defects ($L = 26$ mm).

In this case, 17.3% of the 64×64 voxels of the central cross-section happen to belong to the ROI, and only 1.5% of the $64 \times 64 \times 64$ voxels of the object belong to it, while the ROI incorporates all 64 voxels of the defects. This example shows that ROI selection can lead to an impressive reduction of the problem dimension. However, it is worth stressing that such a preliminary operation does not raise the main ambiguity concerning the position of the defects, since the voxels that belong to the ROI are organized along the vertical axis, which is the average direction of the rays. It is only aimed at lowering the dimensionality of the reconstruction problem, not at solving it. However, as far as combinatory issues are involved, the size of the vector of unknowns is known as a crucial parameter.

From (18), it should be clear that the achievable reduction depends on the number of flaws in the volume, the noise in the data, and the number and the angular excursion of the projections. Therefore, it is not easy to assess quantitative results concerning the size of the ROI. However, our experience indicates that increasing the number of projections and/or their angular excursion decreases the size of the ROI. Likewise, a smaller SNR yields a larger ROI.

3.3. Generalization to log-concave likelihoods

Proposition 1 can actually be stated in more general terms, provided that the observations \mathbf{y} are statistically independent given the object \mathbf{x} , and that the likelihood of y_m is a log-concave function⁴ of $\mathbf{h}_m^t \mathbf{x}$.

Proposition 3. *Let the likelihood function take the following form:*

$$f_{Y|X}(\mathbf{y} | \mathbf{x}) = \prod_{m=1}^M \phi_m(\mathbf{h}_m^t \mathbf{x}), \quad (19)$$

where ϕ_m are log-concave functions. Then (13)–(14) still hold.

Proof. See appendix B. □

⁴ By definition, ϕ is a log-concave function if $\log \phi$ is a concave function.

The main non-Gaussian model we have in mind is the Poisson data model, which is the acknowledged model to describe that the measurements are related to a counting process. In the present case of transmission tomography, the corresponding data likelihood reads

$$\Pr(\mathbf{Y} = \mathbf{y} \mid \mathbf{X} = \mathbf{x}) = \prod_{m=1}^M \frac{(z_m)^{y_m}}{y_m!} \exp(-z_m). \quad (20)$$

In the latter expression, $z_m = \tau_s \exp(-\xi_m)$, where τ_s is the emission rate of the source and ξ_m is related to the attenuation along the m th ray according to

$$\xi_m = \mathbf{h}_m^t \mathbf{x} + g_m,$$

where $\mathbf{g} = (g_m)$ stands for a background component that accounts for the Compton scattering effect. Taking the scattering into account ($g_m \neq 0$) leads to non-logconcave functions. However, when the scattering can be neglected the Poisson likelihood (20) is a product of log-concave functions of \mathbf{x} and proposition 3 applies with the corresponding expression for $\underline{R}_n(\mathbf{y})$:

$$\underline{R}_n(\mathbf{y}) = \prod_{m=1}^M \exp(\tau_s (\exp(-h_{mn}) - 1) + y_m h_{mn}). \quad (21)$$

Here we shall not investigate the Poisson data model further since our main concern is to consider gamma or x-ray tomographic measurements with high counting rates, which are well approximated by the simple Gaussian described in section 2.

4. Reconstruction using a BMLR approach

In the case of limited-angle tomography, designing a computationally attractive algorithm that leads to high quality binary reconstructions is a challenge. Here, the proposed approach is twofold. Firstly, the reconstructed volume is reduced to the ROI as defined in the previous section. Our aim is then to minimize the penalized least-square criterion (8) with respect to the binary voxels that pertain to the ROI, while the others are forced to zero. For this purpose, we propose a deterministic procedure called block most likely replacement (BMLR). It is a heuristic procedure comparable to ICM. However, it provides significantly better results than ICM, while being still of a moderate computational cost for the considered NDE problems.

4.1. ROI constrained formulation

Let us introduce the set $\mathbb{I} = \{\mathbf{x} \in \mathbb{S} : x_n = 0 \text{ if } n \notin \mathcal{I}\}$. Hereafter, the assumption $\mathbf{x} \in \mathbb{I}$ is taken for granted. It is then natural to consider the problem of maximizing (8) under the constraint $\mathbf{x} \in \mathbb{I}$. Let us remark here that such a constrained formulation is not necessarily equivalent to the unconstrained one. Actually, it is easy to check that the two problems admit the same solution if $\widehat{\mathbf{X}}^{\text{MAP}} \in \mathbb{I}$ only. Unfortunately, it is far from easy to check the latter condition since $\widehat{\mathbf{X}}^{\text{MAP}}$ is unknown. Nonetheless, maximizing (8) under the constraint $\mathbf{x} \in \mathbb{I}$ remains coherent and likely to produce a meaningful solution. Moreover, it can be checked by practice that the proposed maximization procedure with or without the ROI provides close reconstruction results—see section 5.

4.2. ICM and SMLR heuristics

From the sake of low computational complexity, the ICM procedure introduced in [13] is of particular interest since it amounts to simple successive scalar relaxation of voxels. Let

$\mathcal{N} \stackrel{\text{def}}{=} \{n_1, \dots, n_{|\mathcal{I}|}\}$ be an arbitrary enumeration of the set \mathcal{I} . From any initial solution $\mathbf{x}^{(1)} \in \mathbb{I}$, the k th iteration of ICM can be formally defined by

$$\begin{aligned} \mathbf{x}^{(k,1)} &= \mathbf{x}^{(k)} \\ \ell = 1 \dots |\mathcal{I}|, \quad \mathbf{x}^{(k,\ell+1)} &= \begin{cases} \boldsymbol{\zeta}^{(k,\ell)} & \text{if } \mathcal{J}(\boldsymbol{\zeta}^{(k,\ell)}) < \mathcal{J}(\mathbf{x}^{(k,\ell)}), \\ \mathbf{x}^{(k,\ell)} & \text{otherwise,} \end{cases} \\ \mathbf{x}^{(k+1)} &= \mathbf{x}^{(k,|\mathcal{I}|+1)}, \end{aligned}$$

where $\boldsymbol{\zeta}^{(k,\ell)}$ is a *proposition vector* defined by

$$\boldsymbol{\zeta}^{(k,\ell)} \stackrel{\text{def}}{=} \mathbf{x}^{(k,\ell)} \oplus \mathbf{e}_{n_\ell}, \quad (22)$$

and \oplus stands for the componentwise *logical OR exclusive*.

The ICM iterations can be implemented efficiently in a recursive fashion, as follows. Let \mathbf{x} and \mathbf{x}' be two binary N -length vectors that only differ at index n , i.e., $x'_n = x_n \pm 1$. Then, it is easy to establish that

$$\mathcal{J}(\mathbf{x}') = \mathcal{J}(\mathbf{x}) + \mathbf{h}_{,n}^t (\mathbf{h}_{,n} \mp 2\varepsilon) \pm 2\sigma^2\mu, \quad (23)$$

where $\varepsilon = \mathbf{y} - \mathbf{H}\mathbf{x}$. Moreover, we have also

$$\varepsilon' = \mathbf{y} - \mathbf{H}\mathbf{x}' = \varepsilon \mp \mathbf{h}_{,n}. \quad (24)$$

In the implementation of ICM, equations (23) and (24) can be used to compare $\mathcal{J}(\boldsymbol{\zeta}^{(k,\ell)})$ to $\mathcal{J}(\mathbf{x}^{(k,\ell)})$ and to recursively compute $\varepsilon^{(k,\ell)} = \mathbf{y} - \mathbf{H}\mathbf{x}^{(k,\ell)}$, respectively.

By construction, the ICM ensures that the criterion \mathcal{J} never increases. Moreover, convergence occurs after a finite number of iterations, as soon as $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)}$. However, an obvious drawback of this heuristic is that the solution is not invariant with respect to the scanning order in \mathcal{I} .

The latter limitation of ICM can be made by a *single most likely replacement* (SMLR) procedure, as introduced by Kormylo and Mendel [14] in the field of sparse spike train restoration. According to SMLR, each scan yields the swap of only one voxel, the one that produces the largest decrease of the criterion

$$\mathbf{x}^{(k+1)} = \arg \max_{\mathbf{x} \in \mathbb{V}(\mathbf{x}^{(k)})} \mathcal{J}(\mathbf{x}), \quad (25)$$

where $\mathbb{V}(\mathbf{x}^{(k)}) = \{\boldsymbol{\zeta}^{(k,\ell)} : 1 \leq \ell \leq |\mathcal{I}|\}$ is the set of all proposition vectors generated by elementary modifications of $\mathbf{x}^{(k)}$. SMLR shares several features of ICM as follows.

- An efficient implementation is obtained thanks to the recursive expressions (23) and (24).
- The criterion \mathcal{J} never increases, and convergence occurs after a finite number of iterations, as soon as $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)}$.

Although very attractive from the computational side, none of these heuristics provide faithful reconstructions in our context. In particular, we have found that the set $\mathbb{V}(\mathbf{x}^{(k)})$ is not rich enough to prevent ICM and SMLR getting stuck in bad solutions. To gain efficiency, richer sets of propositions should prior allow neighbouring voxels to swap their values, so that the position of the hypothetical defects can adjust more easily along the iterations. A possible solution would be to adapt Chi and Mendel's single-spike-shift detector [22]. We found it more appropriate to develop an original variation of it, which is developed in the following section.

4.3. The BMLR heuristic

Whereas the SMLR only explores the two possible values of a single voxel at each step, the new heuristic explores the $2^8 = 256$ configurations of a basic $2 \times 2 \times 2$ cube of voxels in the ROI. The resulting procedure is called *block most likely replacement* (BMLR).

4.3.1. Formal expression of the heuristic. Let Ω stands for the set of nonempty intersections between the ROI and the $2 \times 2 \times 2$ cubes. Hereafter, such intersections will be called *blocks*. Most of the blocks contain eight voxels, but some of those located at the boundaries of the ROI may contain fewer voxels.

In what follows, the compact notation \mathbf{x}_ω is employed for the subvector $\{x_n, n \in \omega\}$, where ω is a set of sites, i.e., a subset of $\{1, \dots, N\}$. Let us also introduce the notation $\zeta^{(k, \omega, b)}$ for the proposition vector defined by

$$\begin{aligned}\zeta_\omega^{(k, \omega, b)} &= \mathbf{x}_\omega^{(k)} \oplus \mathbf{b}, \\ \zeta_{\mathcal{I} \setminus \omega}^{(k, \omega, b)} &= \mathbf{x}_{\mathcal{I} \setminus \omega}^{(k)},\end{aligned}$$

where $\omega \in \Omega$ and $\mathbf{b} \in \{0, 1\}^{|\omega|}$. Then, the updating equation (25) still holds for the BMLR heuristic, with an extended set of proposition vectors:

$$\mathbb{V}(\mathbf{x}^{(k)}) = \{\zeta^{(k, \omega, b)} : \omega \in \Omega, \mathbf{b} \in \{0, 1\}^{|\omega|}\}.$$

Again, the criterion \mathcal{J} never increases, and convergence occurs after a finite number of iterations, as soon as $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)}$. However, the examination of each block ω requires $2^{|\omega|}$ evaluations of the criterion, where the block size $|\omega|$ varies between 1 and 8. The resulting computational cost can be prohibitive if the implementation issue is not carefully dealt with. The following subsection proposes an efficient solution based on grey codes.

4.3.2. Efficient implementation using grey codes. A grey code (or *reflected code*) is a binary numeral system where two successive values differ in only one digit [23, section 20.2]:

$$(0, 1, 11, 10, 110, 111, 101, 100, 11000, 11001, \dots).$$

Let us remark that the first 2^n elements are obtained by some permutation of the first 2^n natural numbers.

For each block ω , we propose to explore the $2^{|\omega|}$ binary configurations in the grey code order rather than in the usual one. Then, two successive trials \mathbf{b} and \mathbf{b}' , respectively, correspond to two vectors $\zeta^{(k, \omega, b)}$ and $\zeta^{(k, \omega, b')}$ that only differ at a single index. As a consequence, the recursive expressions (23) and (24) can be applied to explore the $2^{|\omega|}$ values of $\mathcal{J}(\zeta^{(k, \omega, b)})$ at an affordable numerical cost, for all $\omega \in \Omega$.

An even more efficient version consists in introducing $f_n = \mathbf{h}_n^t \boldsymbol{\varepsilon}$ and $g_{mn} = \mathbf{h}_m^t \mathbf{h}_n$ for all $m, n \in \omega$, so that (23) can be replaced by

$$\mathcal{J}(\mathbf{x}') = \mathcal{J}(\mathbf{x}) + g_{nn} \mp 2f_n \pm 2\sigma^2\mu,$$

while the $|\omega|$ scalars f_n can be recursively computed using

$$f'_m = \mathbf{h}_{,n}^t \boldsymbol{\varepsilon}' = f_m \mp g_{mn}.$$

While each of the $2^{|\omega|}$ applications of (23) involves a scalar product, none is required anymore in the new recursion. Only $|\omega|(|\omega| + 1)/2$ scalar products are required to compute g_{mn} for all $m, n \in \omega$.

Table 1 depicts the pseudo-code for the central part of the BMLR scheme. For the sake of brevity, the set of blocks Ω is assumed available from previous calculation, as well as the projection matrix \mathbf{H} . For large size problems, \mathbf{H} cannot be stored and its entries must rather be recalculated when required.

In table 1, the initial configuration is the null object $\mathbf{x}^0 = \mathbf{0}$. However, a different initialization point could be easily considered. In particular, a faster but less optimal method

Table 1. Pseudo-code of BMLR.

$K \leftarrow 8; G_1 \leftarrow 1; \lambda \leftarrow 2\sigma^2\mu$	
for $i = 1 : K - 1$ do	
$k \leftarrow 2^i; G_k \leftarrow i + 1$	% $G \in \{0, 7\}$: position of
for $j = 1 : k - 1$ do	% changing bit between
$G_{k+j} \leftarrow G_{k-j}$	% successive elements in
end for	% Gray code of length 2^8
end for	
$\mathbf{x}^0 \leftarrow \mathbf{0}; \boldsymbol{\varepsilon}^0 \leftarrow \mathbf{y}; \mathcal{J}^0 \leftarrow \ \mathbf{y}\ ^2$	% Initialization
repeat	% Main loop
$\mathcal{J}_{\max} \leftarrow \mathcal{J}^0;$	
for all $\boldsymbol{\omega} \in \Omega$ do	% Loop on all blocks
for all $n \in \boldsymbol{\omega}$ do	
$f_n \leftarrow \mathbf{h}_{,n}^t \boldsymbol{\varepsilon}$	
$g_{nn} \leftarrow \mathbf{h}_{,n}^t \mathbf{h}_{,n}$	
for all $m \in \boldsymbol{\omega}, m < n$ do	
$g_{mn} \leftarrow g_{nm} \leftarrow \mathbf{h}_{,m}^t \mathbf{h}_{,n}$	
end for	
end for	
$\mathcal{J} \leftarrow \mathcal{J}^0; \mathbf{x} \leftarrow \mathbf{x}^0$	
for $k = 0 : 2^{ \boldsymbol{\omega} } - 1$ do	% Loop on all block states
$n \leftarrow \boldsymbol{\omega}_{G_k}$	
$s \leftarrow 2x_n - 1$	% $s = 1$ if $x_n = 1 \rightarrow 0$,
$x_n \leftarrow 1 - x_n$	% $s = -1$ if $x_n = 0 \rightarrow 1$
$\mathcal{J} \leftarrow \mathcal{J} + g_{nn} + s(2f_n - \lambda)$	
for all $m \in \boldsymbol{\omega}$ do	
$f_m \leftarrow f_m + sg_{mn}$	
end for	
if $\mathcal{J} < \mathcal{J}_{\max}$ then	% New best configuration
$\mathcal{J}_{\max} \leftarrow \mathcal{J}; n_{\max} \leftarrow n; s_{\max} \leftarrow s$	
end if	
end for	% End of loop on block states
end for	% End of loop on blocks
if $\mathcal{J}_{\max} < \mathcal{J}^0$ then	% New configuration
$\mathcal{J}^0 \leftarrow \mathcal{J}_{\max}; x_{n_{\max}}^0 \leftarrow 1 - x_{n_{\max}}^0; \boldsymbol{\varepsilon}^0 \leftarrow \boldsymbol{\varepsilon}^0 + s_{\max} \mathbf{h}_{,n}$	
end if	
until $\mathcal{J}_{\max} = \mathcal{J}^0$	% End of main loop. The best % found configuration is \mathbf{x}^0 .

could provide an initial solution, e.g., ICM, that would be refined using BMLR. For the sake of computational time, such a two-step approach would be of particular interest to reconstruct objects of larger extent than localized defects.




Figure 3. Left: pairs of voxelized spheres considered in the data generation process, in configuration separated ($L = 26$ mm) and close ($L = 12$ mm), respectively. Right: corresponding noisy projection from source S6 with noise standard deviation $\sigma = 0.01$.

5. Test on synthetic data sets

The 3D synthetic NDE problem presented in section 3.2 is now considered in order to test the proposed reconstruction algorithm. The aim is to illustrate the detection, sizing and separation capabilities of the method on the ground of simulations.

5.1. The simulation context

The considered reconstruction problem involves two superimposed defects of spherical shape in a homogeneous material. The volume of interest contains $64 \times 64 \times 64$ voxels, while each of the two defects spreads over 32 voxels. Two different values have been considered for the distance between the two defects: $L \in \{26 \text{ mm}, 12 \text{ mm}\}$ (see figure 3(left)). The data set \mathbf{y} consists in seven projections that were corrupted by an uncorrelated Gaussian noise with standard deviation $\sigma \in \{0, 0.005, 0.01\}$ (see figure 3(right)). The geometric dimensions of the problem are given in figure 1. They correspond to a realistic situation of in-service pipe inspection in the nuclear industry [1, figure 1].

5.2. Non-binary reconstruction results

Here, a non-binary penalized approach tested in [1] is used as a comparative tool. It consists of considering an optimization problem of the form (1) with $\mathbb{X} = \mathbb{R}_+^N$ and

$$\psi(\mathbf{x}; \boldsymbol{\theta}) = \lambda \sum_{i \sim j} \phi(x_i - x_j; \delta) + \alpha \sum_{n=1}^N |x_n|, \quad (26)$$

where $\lambda, \alpha \geq 0$, $i \sim j$ stands for all pairs of neighbouring sites in the first-order neighbourhood system (hence, each voxel has six neighbours) and $\phi(t; \delta)$ is a Huber function [24]. The role of such a penalty function is to favour piecewise homogeneous solutions. In practice, the

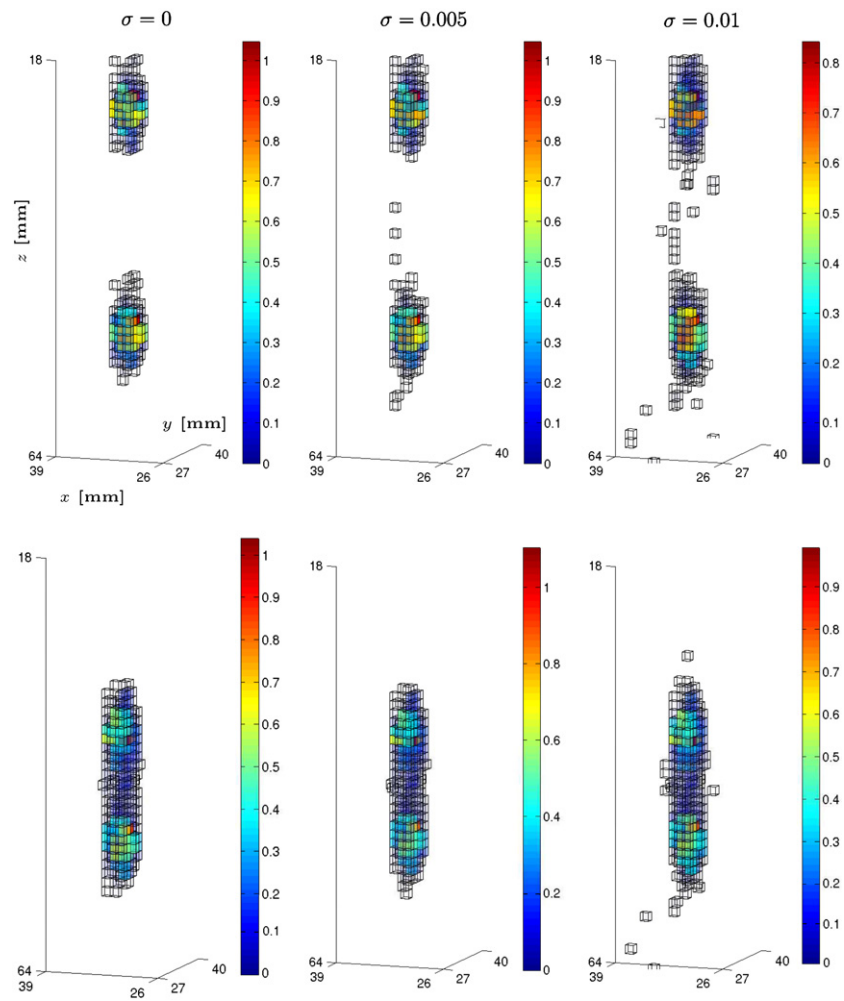


Figure 4. Reconstructions from noisy projections of a pair of voxelized spheres using a penalized approach.

penalized approach requires the tuning of the three hyperparameters $\theta = (\lambda, \delta, \alpha)^t$. Here, they have been empirically adjusted according to qualitative appreciation. The corresponding reconstruction results are depicted in figure 4. As expected, the estimated attenuations are underestimated, while the extension of the defects is overestimated along the vertical axis. In particular, the result does not allow us to clearly discriminate the two spheres in the case $L = 12$ mm.

The penalized least-square criterion being convex, the global minimizer can be computed from an initial guess $\mathbf{x}^{(0)}$ by an appropriate iterative descent algorithm. Here, $\mathbf{x}^{(0)}$ has been chosen as the backprojection of the data, and the minimization step was performed using a projected-gradient method, i.e., a straightforward adaptation of *steepest descent* that handles the positivity constraint in a natural way [25, p 203]. However, this algorithm inherits the slow asymptotic convergence of a gradient method. Table 2 shows that several hours are required to reach the adopted convergence criterion $\psi(\mathbf{x}^{(k-1)}) - \psi(\mathbf{x}^{(k)}) < 10^{-6}$, where $k \in \mathbb{N}$ the

Table 2. Number of iterations and CPU time for the penalized approach and the ICM and BMLR heuristics on a Pentium III 2.4 GHz (586 series) with 1Go RAM.

Algorithm	L (mm)	σ	Number of iterations	CPU time
Penalized approach	26	0	333	3 h 5 min
		5×10^{-3}	97	1 h 25 min
		10^{-2}	198	2 h 55 min
	12	0	325	3 h 42 min
		5×10^{-3}	356	5 h 11 min
		10^{-2}	410	5 h 58 min
ICM	26	0	9	0.408 s
		5×10^{-3}	11	0.452 s
		10^{-2}	11	0.501 s
	12	0	7	0.306 s
		5×10^{-3}	8	0.331 s
		10^{-2}	8	0.333 s
BMLR	26	0	30	1.57 s
		5×10^{-3}	27	1.55 s
		10^{-2}	346	16.94 s
	12	0	19	1.18 s
		5×10^{-3}	19	1.23 s
		10^{-2}	339	17.08 s
BMLR (with isolated voxels discarded from the ROI)	26	0	30	1.57 s
		5×10^{-3}	27	1.53 s
		10^{-2}	173	8.37 s
	12	0	19	1.17 s
		5×10^{-3}	19	1.29 s
		10^{-2}	170	8.43 s

iteration number. Let us remark here that the same constrained minimization problem could probably be addressed using more efficient, but more complex algorithms (see for instance [26]).

5.3. Binary reconstruction results

In this section, the proposed BMLR algorithm is compared to the ICM. In all cases, the initial state is chosen as a defect-free object, i.e., $\mathbf{x}^{(0)} = \mathbf{0}$. Both heuristics require the tuning of the hyperparameter $\mu = \log(p_0/p_1)$. Since the defects are expected to constitute a very small part of the total volume, a strict Bayesian viewpoint would lead to choosing prior probabilities p_0 and $p_1 = 1 - p_0$ such that $p_0 \gg p_1$. For two distinct reasons, we rather set $p_0 = p_1$.

- On the one hand, false positives are preferable to false negatives, which justifies overvaluing p_1/p_0 . This could be reformulated more rigorously according to the Bayesian cost theory, by replacing the MAP estimator by another Bayesian estimator that accounts for distinct costs for false negatives and false positives, respectively c_0 and c_1 . Then our choice corresponds to the assumption that $p_0/c_0 = p_1/c_1$.
- On the other hand, choosing $p_0 = p_1$ leads to a fully unsupervised version of BMLR.

The reconstruction results are shown in figures 5 and 6 for the ICM and the BMLR algorithms, respectively. Even in the noiseless case, the ICM fails to provide a faithful

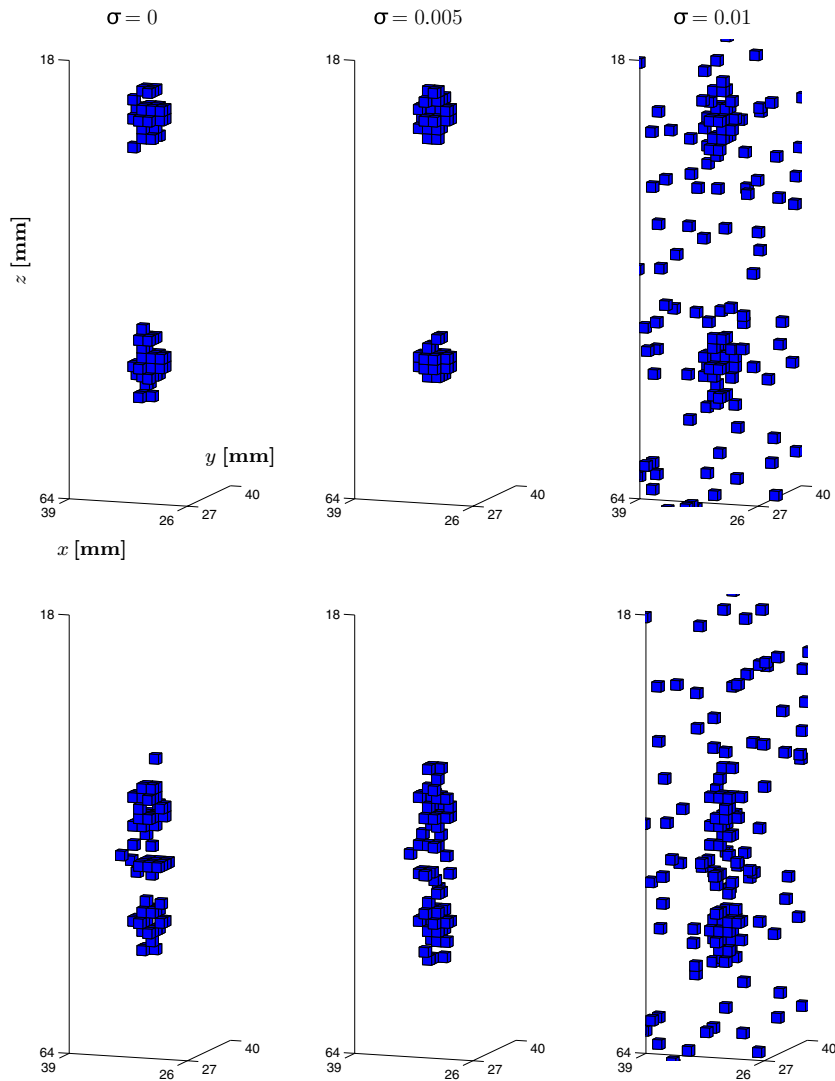


Figure 5. Reconstructions from noisy projections of a pair of voxelized spheres using an ICM approach.

reconstruction in the case of two close spheres. In comparison, the BMLR algorithm achieves a perfect reconstruction in all cases of high or moderate SNR.

On the other hand, the number of false positive voxels is subject to a threshold effect: when the noise standard deviation reaches the magnitude of the projection of a single voxel (whose maximum value is $\sigma_{th} = 1/64 \approx 0.015$ in the present simulation), many spurious, isolated voxels appear in the reconstruction. This is clearly visible in figure 6 in the case $\sigma = 0.01$. In this case, two distinct strategies may be adopted to reduce the number of false positive voxels.

The first one rests upon the fact that the BMLR reconstruction remains acceptable, let alone those spurious voxels. As far as real defects can be assumed to extend over several

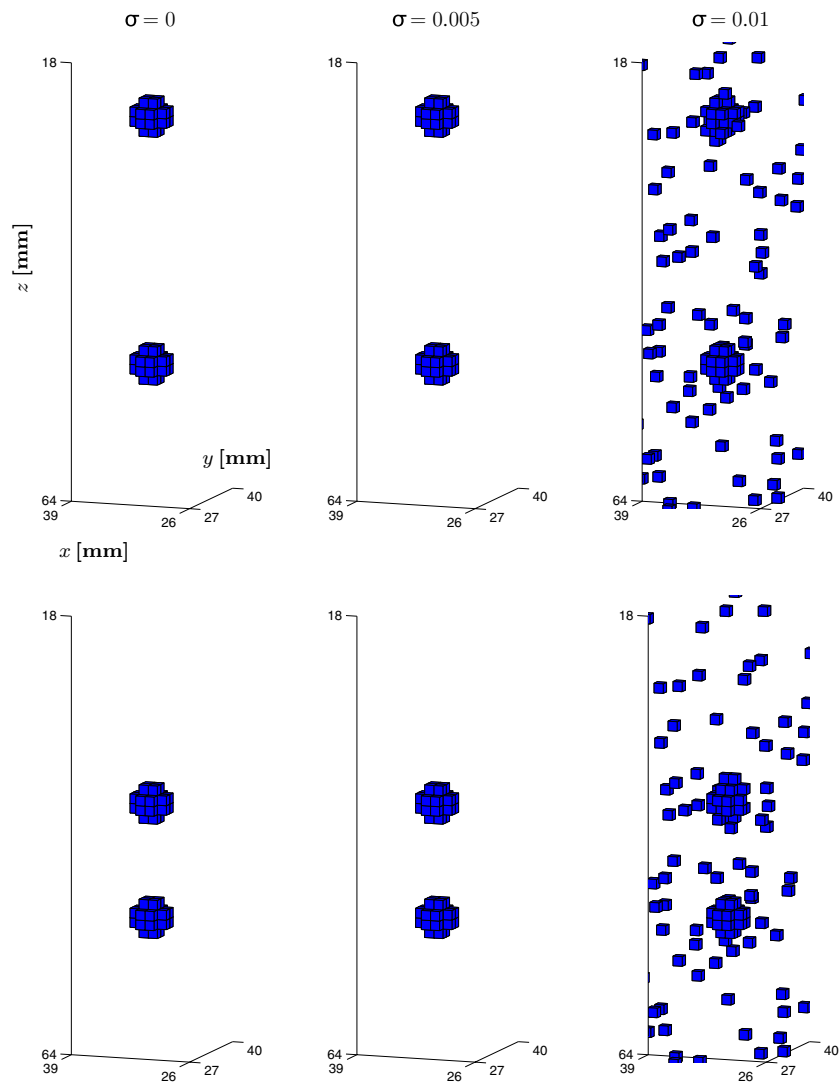


Figure 6. BMLR reconstructions from noisy projections of a pair of discretized defects.

neighbouring voxels, the idea is then to remove isolated voxels by simple post-processing. For the sake of computational burden, it is interesting to also discard isolated voxels from the ROI. In the considered cases where $\sigma = 0.01$, a large number of iterations is spent activating voxels that are isolated in the ROI, and can thus be saved. The resulting reconstruction is depicted in the first column of figure 8.

An alternate strategy is to introduce a penalization of active voxels by assigning a positive value to the regularization parameter $\lambda \stackrel{\text{def}}{=} 2\sigma^2\mu$ in criterion (8). The tuning of λ has a great impact on the final quality of the 3D attenuation map. As illustrated in figure 9, a correct tuning of λ significantly reduces the number of spurious voxels and produces a satisfactory reconstruction of the flaws. However, overestimating λ leads to false negative voxels, so that the tuning of λ must be cautiously handled.

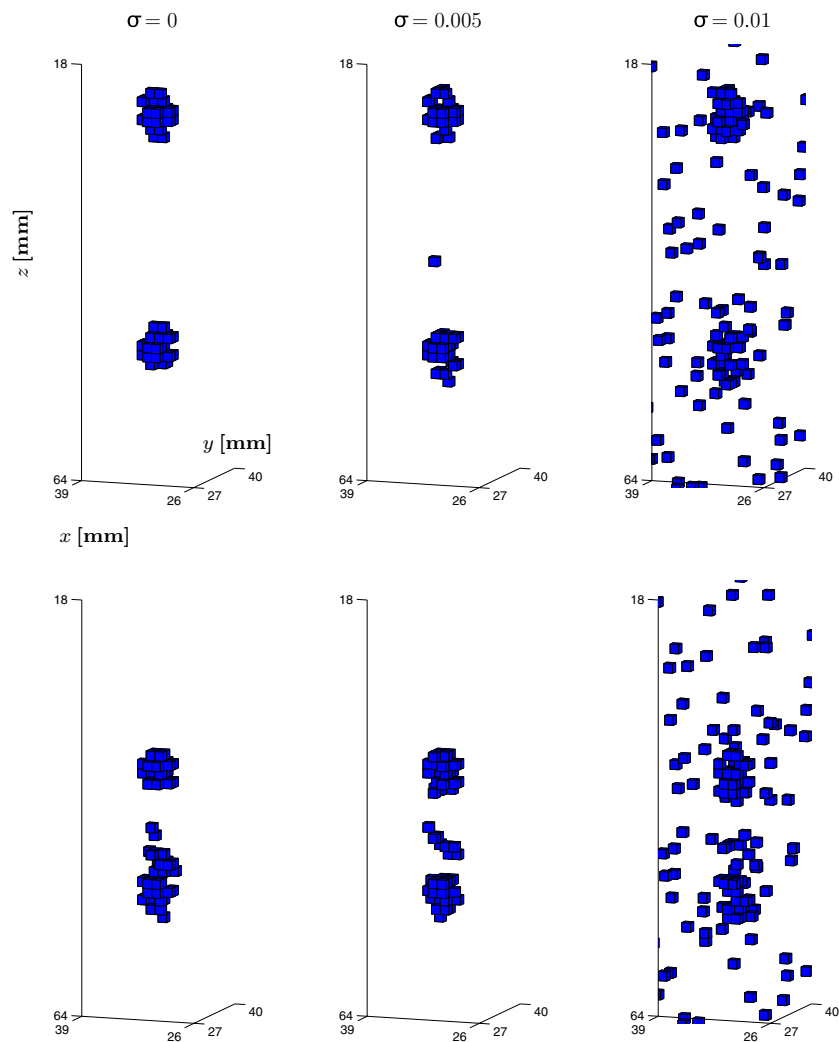


Figure 7. BMLR reconstructions from noisy projections of a pair of continuous defects.

The number of iterations and the CPU time required for convergence for the ICM and BMLR heuristics are shown in table 2. Not surprisingly, ICM is a very fast method. However, BMLR remains very attractive in terms of computing time. The step of building the ROI is not expensive since it results from a simple thresholding of the backprojection map, see (18). Depending on the SNR and on the positions of the spheres, the ROI varies from 1.1% to 1.6% of the total number of voxels in the considered examples. The total number of iterations is roughly proportional to the number of activated voxels in the solution. Obviously, the latter depends on the volume of the defects, which is expected to be small in NDE. However, it is also dependent on the SNR: the BMLR heuristic converges in a dozen iterations if $\sigma < \sigma_{th}$, whereas a few hundreds are required if $\sigma > \sigma_{th}$. In the former case, the convergence is fast and the solution is reached within 2 s. In the latter case, most of the iterations activate only one isolated site of the ROI and the time needed for convergence increases significantly. As

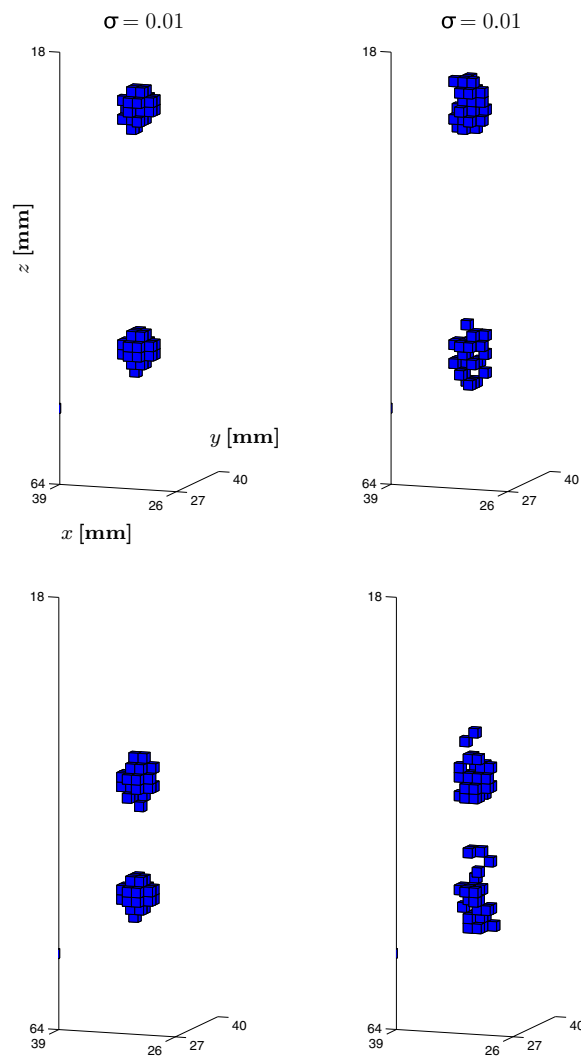


Figure 8. For the lowest SNR values, removing isolated voxels in the final BMLR reconstruction provides a better solution to both discretized (left) and continuous (right) defect problems.

shown in table 2, removing isolated sites from the ROI allows us to speed up the convergence of the BMLR in this case. As stated earlier in this section, this pre-processing of the ROI is natural provided that flaws extend over more than a single voxel in the volume.

In principle, the BMLR heuristic can provide different solutions whether the ROI is considered or not. In order to test the variability of the solution, we also ran the BMLR heuristic with $\mathbb{I} = \mathbb{S}$, i.e., we ignored the ROI. For high-to-moderate SNR values, the results were strictly identical to the previous ones, while only marginal differences appeared for the low SNR case. Finally, let us remark that the BMLR heuristic asks for an unrealistic computation time of several hours when the ROI is ignored, or if all the voxels belong to the ROI. In such an unfavourable situation, the BMLR could still be efficiently implemented using a parallel structure, since the blocks can be explored independently of each other.

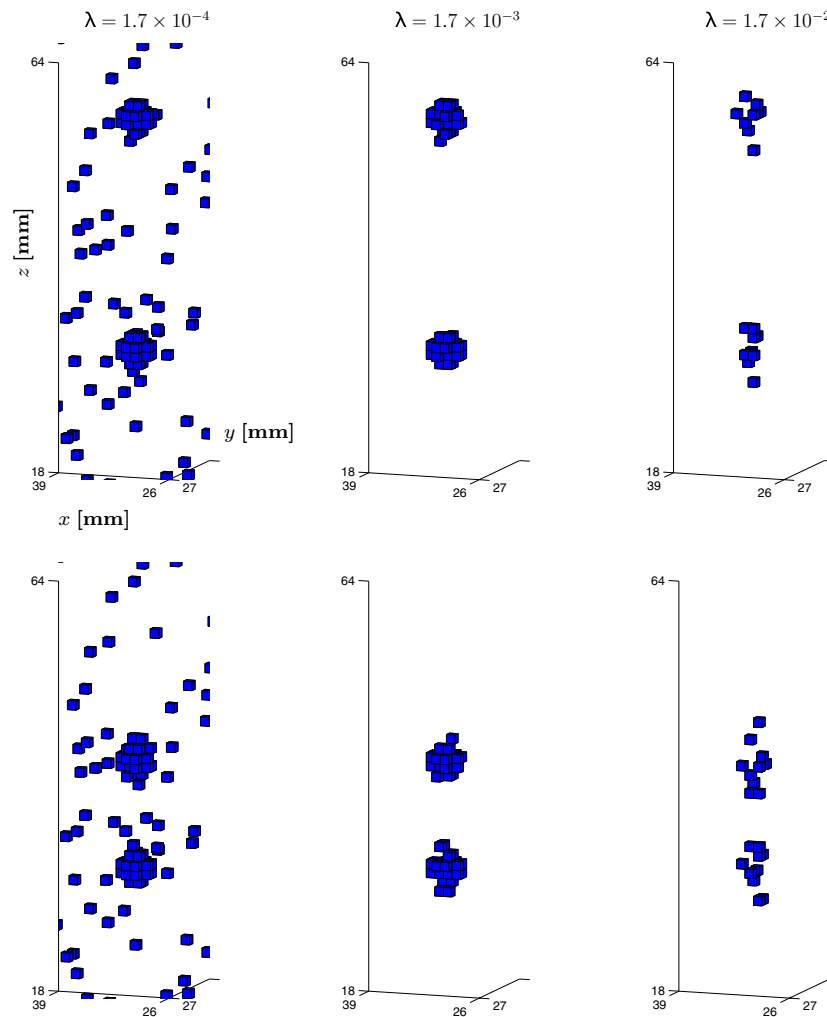


Figure 9. Impact of the regularization parameter $\lambda = 2\sigma^2 \log(p_0/p_1)$ on BMLR reconstruction at the lowest SNR (i.e., $\sigma = 0.01$). Whereas under regularization (left) does not efficiently eliminate spurious voxels and over regularization (right) cancels most of the voxels in the flaws, a correct tuning of λ (centre) produces a satisfactory reconstruction.

5.4. BMLR: robustness with respect to modelling errors

In the previous subsection, the data are noisy projections of discretized objects. Obviously, real objects are continuous and (3) only constitutes an approximate model. Therefore, we found it appropriate to further investigate the robustness of the method by considering projections from perfectly spherical objects, given that closed-form expressions are available for the projection of ellipsoids. The reconstructions provided by the BMLR are shown in figure 7. In the case of distant objects, a good reconstruction is obtained whatever the SNR—up to the fact that spurious, isolated voxels must be removed for $\sigma = 0.01$, as shown in the second column of figure 8. In the case of closer objects, the quality of the reconstruction deteriorates

slightly, but the BMLR still provides meaningful results, since the two objects are still clearly distinguishable one from another.

6. Conclusion

Spatially interacting models such as Markov random fields are usually employed within a Bayesian approach to reconstruction. Here, we rather resorted to a simpler iid model, and we proposed an efficient binary reconstruction algorithm, both in terms of robustness and computational burden. For high-to-moderate SNR, the binary constraint is sufficient to regularize the inversion of the tomographic problem. For low SNR, however, the number of false positive voxels increases dramatically. A simple strategy to deal with this problem is to cancel isolated voxels. Another strategy is to perform a penalization of active voxels by assigning a positive value to the regularization parameter $\lambda = 2\sigma^2\mu$ in criterion (8). However, the supervised tuning of λ must then be handled.

With a view to reconstruction from real data, several comments can be made about the physical model considered here. Firstly, the additive Gaussian noise assumption may not be appropriate for low counting rates τ_m . In such cases, the Poisson model (20) provides a more accurate description of the data. The ROI should then be computed using (21), and straightforward modifications should be brought to the BMLR algorithm. Secondly, it should be noted that the Beer–Lambert relation (2) is only valid for a monochromatic emission of photons under a non-diffractive propagation assumption. In particular, Compton scattering is then negligible, which is not necessarily true for high energy gamma-ray sources (beyond 100 keV). In the latter case, further investigations would be required to define an appropriate ROI, for instance on the basis of the scattering model developed in [27].

Acknowledgment

This work was partially supported by Electricité de France (EDF) under contract number E961/P10AY0/E3.

Appendix A. Proof of proposition 1

First, it can immediately be checked that

$$\bar{\Delta}_n = -\log \underline{R}_n(\mathbf{y}) = \frac{1}{2\sigma^2}(\|\mathbf{y}\|^2 - \|\mathbf{y} - \mathbf{h}_{,n}\|^2) - \mu,$$

which readily gives expression (16).

Now let us show that

$$\underline{R}_n(\mathbf{y}) \leq \frac{f_{Y|X}(\mathbf{y}|\mathbf{x})}{f_{Y|X}(\mathbf{y}|\mathbf{x}')} \quad (\text{A.1})$$

for all couples $(\mathbf{x}, \mathbf{x}') \in \mathbb{S}_n^0 \times \mathbb{S}_n^1$ such that $\mathbf{x}' = \mathbf{x} + \mathbf{e}_n$: given (5), the latter inequality holds, since

$$\frac{1}{\sigma^2}\|\mathbf{y} - \mathbf{H}\mathbf{x}'\|^2 - \frac{1}{\sigma^2}\|\mathbf{y} - \mathbf{H}\mathbf{x}\|^2 + \bar{\Delta}_n = \frac{2}{\sigma^2}\mathbf{h}_{,n}^t \mathbf{H}\mathbf{x} \geq 0.$$

Given (12), inequalities (A.1) jointly imply that (13) holds.

Appendix B. Proof of proposition 3

In the more general setting of proposition 3, let us show that (A.1) is still valid for all couples $(\mathbf{x}, \mathbf{x}') \in \mathbb{S}_n^0 \times \mathbb{S}_n^1$ such that $\mathbf{x}' = \mathbf{x} + \mathbf{e}_n$. Given (19), this amounts to proving that

$$\prod_{m=1}^M \frac{\phi_m(a)\phi_m(d)}{\phi_m(b)\phi_m(c)} \leq 1,$$

with $a = 0$, $b = \mathcal{H}_{mn}$, $c = \mathbf{h}_{m \cdot}^t \mathbf{x}$, and $d = b + c$, where \mathcal{H}_{mn} denotes the (m, n) entry of matrix \mathbf{H} . It is actually true that each term of the above product belongs to $[0, 1]$. This is obvious for those terms for which $\mathcal{H}_{mn} = \mathbf{h}_{m \cdot}^t \mathbf{x} = 0$, since we have then $a = b = c = d$. Otherwise, let us remark that $a \leq b \leq d$ and $a \leq c \leq d$, so that

$$\log \phi_m(b) \geq \theta \log \phi_m(a) + (1 - \theta) \log \phi_m(d) \quad (\text{B.1})$$

$$\log \phi_m(c) \geq \theta' \log \phi_m(a) + (1 - \theta') \log \phi_m(d), \quad (\text{B.2})$$

where

$$\theta = \frac{d - b}{d - a} = \frac{c}{b + c} \in [0, 1], \quad \theta' = \frac{d - c}{d - a} = \frac{b}{b + c} = 1 - \theta.$$

Finally, term-to-term summation of inequalities (B.1) and (B.2) allows us to conclude that $\phi_m(a)\phi_m(d) \leq \phi_m(b)\phi_m(c)$.

References

- [1] Chatellier L, Fournier L, Charbonnier B, Robillard L and Chassignole B 2005 Characterization improvement through signal processing application to radiography inspection: 3D reconstruction *Pressure Vessels and Piping Division Conference* (Denver, CO: ASME)
- [2] Chalmond B, Coldefy F and Lavayssière B 1999 Tomographic reconstruction from non-calibrated noisy projections in non-destructive evaluation *Inverse Problems* **15** 399–411
- [3] Retraint F, Peyrin F and Dinten J-M 1998 Three-dimensional regularized binary image reconstruction from three two-dimensional projections using a randomized ICM algorithm *Int. J. Imaging Syst. Technol.* **9** 139–46
- [4] Hanson K M and Wechsung G W 1983 Bayesian approach to limited-angle reconstruction in computed tomography *J. Opt. Soc. Am.* **73** 1501–9
- [5] Rangayyan R, Dhawan A P and Gordon R 1985 Algorithms for limited-view computed tomography: an annotated bibliography and a challenge *Appl. Opt.* **24** 4000–12
- [6] Quinto E T 1988 Tomographic reconstructions from incomplete data—numerical inversion of the exterior Radon transform *Inverse Problems* **4** 867–76
- [7] Rossi D J and Willsky A S 1984 Reconstruction from projection based on detection and estimation of objects: parts I and II. Performance analysis and robustness analysis *IEEE Trans. Acoust. Speech Signal Process.* **32** 887–906
- [8] Benedetti M, Donelli M and Massa A 2007 Multicrack detection in two-dimensional structures by means of GA-based strategies *IEEE Trans. Antennas Propag.* **55** 205–15
- [9] Hanson K M, Cunningham G S, Jennings G R J and Wolf D R 1994 Tomographic reconstruction based on flexible geometric models *Proc. IEEE ICIP (Austin, TX)* vol 2 pp 145–7
- [10] Soussen C and Mohammad-Djafari A 2004 Polygonal and polyhedral contour reconstruction in computed tomography *IEEE Trans. Image Process.* **13** 1507–23
- [11] Sethian J A 1999 Level set methods and fast marching methods *Cambridge Monograph on Applied and Computational Mathematics* (Cambridge: Cambridge University Press)
- [12] Bruandet J-P, Dinten J-M, Peyrin F and Barlaud M 2002 3D tomographic reconstruction of binary images from cone beam projections: a fast level set approach *Proc. IEEE ISBI (Washington, USA)*
- [13] Besag J E 1986 On the statistical analysis of dirty pictures (with discussion) *J. R. Stat. Soc. B* **48** 259–302
- [14] Kormylo J J and Mendel J M 1982 Maximum-likelihood detection and estimation of Bernoulli–Gaussian processes *IEEE Trans. Inf. Theory* **28** 482–8
- [15] Kak A C and Slaney M 1988 *Principles of Computerized Tomographic Imaging* (New York: IEEE)

- [16] Sauer K D and Bouman C A 1993 A local update strategy for iterative reconstruction from projections *IEEE Trans. Signal Process.* **41** 534–48
- [17] Marroquin J L, Mitter S K and Poggio T A 1987 Probabilistic solution of ill-posed problems in computational vision *J. Am. Stat. Assoc.* **82** 76–89
- [18] Greig D M, Porteous B T and Seheult A H 1989 Exact maximum *a posteriori* estimation for binary images *J. R. Stat. Soc. B* **51** 271–9
- [19] Barahona F 1982 On the computational complexity of Ising spin glass models *J. Phys. A: Math. Gen.* **15** 3241–53
- [20] Istrail S 2000 Statistical mechanics, three-dimensionality and NP-completeness: I. Universality of intractability of the partition functions of the Ising model across non-planar lattices *Proc. 32nd ACM Symp. on Theory of Comp. (Portland, OR)* pp 87–96
- [21] Robert N, Peyrin F and Yaffe M J 1994 Binary vascular reconstruction from limited number of cone beam projections *Med. Phys.* **12** 1839–51
- [22] Chi C Y and Mendel J M 1984 Improved maximum-likelihood detection and estimation of Bernoulli–Gaussian processes *IEEE Trans. Inf. Theory* **30** 429–35
- [23] Press H, William, Flannery B P, Teukolsky S A and Vetterling W T 1989 *Numerical Recipes: the Art of Scientific Computation* (Cambridge: Cambridge University Press)
- [24] Huber P J 1981 *Robust Statistics* (New York: Wiley)
- [25] Bertsekas D P 1995 *Nonlinear Programming* (Belmont, MA: Athena Scientific)
- [26] Johnson C A and Sofer A 2000 A primal-dual method for large-scale image reconstruction in emission tomography *SIAM J. Opt.* **11** 691–715
- [27] Sauer K D, Sachs J J and Klifa C 1994 Bayesian estimation of 3-D objects from few radiographs *IEEE Trans. Nucl. Sci.* **41** 1780–90