

Bayesian Inversion Method for 3-D Dental X-ray Imaging

Ville Kolehmainen, Antti Vanne, Samuli Siltanen, Seppo Järvenpää, Jari P Kaipio, Matti Lassas and Martti Kalke.

Abstract—Diagnostic and operational tasks in dentistry require three-dimensional (3D) information about tissue. A novel type of low dose dental 3D X-ray imaging is considered. Given projection images taken from a few sparsely distributed directions using the dentist’s regular X-ray equipment, the 3D X-ray attenuation function is reconstructed. This is an ill-posed inverse problem, and Bayesian inversion is a well suited framework for reconstruction from such incomplete data. The reconstruction problem is formulated in a well-posed probabilistic form in which *a priori* information is used to compensate for the incomplete data. A parallelized Bayesian method (implemented for a Beowulf cluster computer) for 3D reconstruction in dental radiology is presented (the method was originally presented in (Kolehmainen et al., 2006)). The prior model for dental structures consist of a weighted ℓ^1 and total variation (TV)-prior together with the positivity prior. The inverse problem is stated as finding the maximum a posterior (MAP) estimate. The method is tested with *in vivo* patient data and shown to outperform the reference method (tomosynthesis).

Index Terms—Inverse problems, ill-posed, regularization, tomography, limited angle, X-ray, dental, radiography, Beowulf, parallel computing.

I. INTRODUCTION

Several diagnostic and operative tasks in dentistry require precise three-dimensional (3D) information of dental structures, and often two-dimensional (2D) X-ray images are not sufficient. For example, dental implantology is based on accurate measurements for the optimal depth, size and angle of the screw hole. The hole should be deep enough for firm attachment of the implant while avoiding damage to the mandibular nerve or penetration to the maxillary sinus. See (Brocklebank, 1997; Ekestubbe et al., 1997; Ramesh et al., 2002).

We consider a novel low dose 3D imaging modality for dental implantology. It is intermediate between 2D X-ray imaging and full data tomography: 3D imaging is based on a small number of projection images taken from limited angle of view. The projection data is measured with a modified digital panoramic device; this is a major advantage since

This work was supported by the National Technology Agency of Finland (TEKES projects 2844/31/02 and 1107/401/00) and the Academy of Finland (projects 108299 and 213476, Finnish Programme for Centres of Excellence in Research 2006-2011).

Ville Kolehmainen, Antti Vanne and Jari P Kaipio are with University of Kuopio, P.O.Box 1627, FIN-70211 Kuopio, Finland.

Seppo Järvenpää and Matti Lassas are with Helsinki University of Technology, P.O.Box 1100, FIN-02015 TKK, Finland

Martti Kalke is with PaloDEx Group, P.O.Box 20, FIN-04301 Tuusula, Finland.

Samuli Siltanen is with Tampere University of Technology, P.O.Box 553, FIN-33101 Tampere, Finland.

virtually every dental clinic has a panoramic X-ray device, making the modality well suited to the dentistry workflow (see (<http://www.instrumentariumdental.com/>, 2007)). As opposed to conventional CT data, we refer to the above type of limited projection data as *sparse* projection data.

The reconstruction task is an ill-posed inverse problem because sparse projection data does not contain sufficient information to determine the 3D structure of tissue perfectly. It is well-known that traditional reconstruction methods, such as filtered backprojection (FBP), do not perform well when applied to sparse projection data (Ranggayyan et al., 1985; Natterer, 1986; Hanson, 1987).

Bayesian inversion is a natural framework for reconstruction from sparse projection data. *A priori* information about the tissues is used to compensate for the incomplete data. For example, in dental X-ray imaging we know that the attenuation function is a nonnegative function and that the dental structures consist of a few different tissue types (soft tissue, bone, enamel) separated to a few roughly homogeneous subregions with well-defined boundaries. The unknown X-ray attenuation function and projection data are considered as random variables, and separate statistical models are formulated for (1) the acquisition of the projection data and (2) the *a priori* information. Bayes formula gives the complete solution of the inverse problem in the form of *posterior probability distribution*, from which final images of the tissue are obtained as point estimates. Bayesian methods give improved reconstruction quality over traditional methods, see e.g. (Sauer et al., 1994; Hanson and Wecksung, 1983; Bouman and Sauer, 1993; Frese et al., 2002; Yu and Fessler, 2002; Siltanen et al., 2003; Kolehmainen et al., 2003).

A major practical difficulty in applying Bayesian methods to 3D x-ray imaging is the heavy computational requirements. This is why previous studies on the topic have mostly concentrated on 2D problems. If realistic resolution in a 3D problem is used, the number of unknown voxel values is typically in the range $10^6 - 10^7$, and thus the computation of the posterior statistics leads to large-scale optimization or integration problems. Powerful computers and efficient numerical algorithms are required to compute the 3D reconstruction in clinically acceptable time.

In this study, we present a Bayesian method for 3D reconstruction in dental x-ray imaging with sparse projection data. The method was originally proposed in (Kolehmainen et al., 2006). In the Bayesian model we use a weighted ℓ^1 and total variation (TV) prior, together with the positivity prior, as the prior model for the 3D x-ray attenuation function of the dental structures. A parallelized version of a gradient-based

optimization method is implemented for the computation of the maximum a posterior (MAP) estimate. The results are computed on a 13 node Beowulf cluster constructed for this purpose. The performance of the method is tested with sparse projection data collected using commercial panoramic dental x-ray imaging equipment.

The parallelization of X-ray tomography problem has been previously discussed in (Zheng et al., 2000), where a parallelizable version of an iterative coordinate-descent functional-substitution (ICD/FS) algorithm for Bayesian MAP-estimation was derived.

II. MATHEMATICAL MODEL

A. Forward model

Consider an X-ray source placed on one side of an object, through which radiation passes and is detected on the other side by a digital sensor (2D array of point-like detectors). See Figure 1. The object is modeled by a bounded subset $\Omega \subset \mathbb{R}^3$ with a nonnegative x-ray attenuation function $x : \Omega \rightarrow [0, \infty)$. For the projection measurement, we use the usual pencil-beam model

$$m_j = -\log\left(\frac{I_j}{I_0}\right) = \int_{L_j} x(s) ds, \quad (1)$$

where m_j is the value of the projection measurement for the j^{th} source to detector pixel line L_j in the set of projection data, I_j is the measured x-ray intensity and I_0 is the intensity of the x-ray source.

In the discretization of the attenuation model (1) the domain Ω is divided into a lattice of M disjoint 3D voxels Ω_i and the length of the path L_j inside each voxel Ω_i is computed, see Figure 1. Assuming that the attenuation function $x(s)$ is constant within each voxel Ω_i , the projection measurement m_j can be approximated in the form

$$m_j = \int_{L_j} x(s) ds \approx \sum_{i=1}^M x_i |\Omega_i \cap L_j|. \quad (2)$$

where $|\Omega_i \cap L_j|$ denotes the length of ray L_j through voxel Ω_i . Arranging the whole set of N projection measurements into a vector $m = (m_1, m_2, \dots, m_N)^T \in \mathbb{R}^N$, we obtain

$$m = Ax, \quad (3)$$

where $x = (x_1, x_2, \dots, x_M)^T \in \mathbb{R}^M$ is the vector of attenuation values in the voxels and matrix A implements Eq. (2).

B. Bayesian inversion

We consider the inverse problem as a problem of Bayesian inference. All unknown variables are modeled as random variables. The probabilistic modeling of these variables reflects our uncertainty of their actual values and the degree of uncertainty is coded in their probability distributions.

The posterior distribution

$$p(x | m) = \frac{p_{\text{pr}}(x)p(m | x)}{p(m)}, \quad (4)$$




Fig. 1. Left: Schematic illustration of the pencil-beam attenuation model for x-ray imaging. Right: The domain $\Omega \subset \mathbb{R}^3$ under investigation is discretized into a lattice of M voxels Ω_i .

represents the complete solution of the inverse problem. In Eq. (4), $p(m | x)$ is the so-called likelihood function, $p_{\text{pr}}(x)$ is the prior density and $p(m)$ is the normalization constant. The likelihood function $p(m | x)$ is a statistical model for the observations describing the likelihood that the measured data m would have been observed from a given realization x . The prior density $p_{\text{pr}}(x)$ is statistical model for the unknown x . It is designed based on the *a priori* information about the object.

In this work we compute the MAP estimate

$$p(x_{\text{MAP}} | m) = \max_x p(x | m), \quad (5)$$

which is then shown as the reconstructed image. For further details on Bayesian inversion theory, see (Hanson, 1987; Mosegaard and Sambridge, 2002; Kaipio and Somersalo, 2004; Siltanen et al., 2003).

C. Bayesian model for 3D dental imaging

We use the measurement model

$$m = Ax + \epsilon, \quad (6)$$

for the x-ray projection data. Observation noise is assumed to be independent of x and is modeled by a zero mean Gaussian random vector $\epsilon \sim \mathcal{N}(0, C)$ with an invertible covariance matrix C . By (Bouman and Sauer, 1993; Siltanen et al., 2003) we know that (6) is feasible. In practice, the covariance matrix C can be estimated using repeated phantom measurements. With the model (6), the likelihood function takes the form

$$p(m | x) \propto \exp\left(-\frac{1}{2} \|L(m - Ax)\|^2\right), \quad (7)$$

where $L^T L = C^{-1}$ and $\|\cdot\|$ is the usual Euclidian (ℓ^2) vector norm.

As the prior model for the attenuation function of the dental structures, we write

$$p_{\text{pr}}(x) \propto p_+(x) \exp(-W(x)), \quad (8)$$

where

$$p_+(x) = \prod_{k=1}^M \theta(x_k) \quad (9)$$

is the positivity prior, θ is the Heaviside function

$$\theta(t) = \begin{cases} 1, & t \geq 0 \\ 0, & t < 0 \end{cases} \quad (10)$$

and the functional $W(x)$ is of the form

$$\begin{aligned} W(x) &= \alpha_0 \|x\|_1 + \alpha_1 \text{TV}(x) \\ &= \alpha_0 \sum_{i=1}^M |x_i| + \alpha_1 \sum_{i=1}^M \sum_{j \in \mathcal{N}_i} |x_i - x_j|, \end{aligned} \quad (11)$$

where \mathcal{N}_i denotes the usual six-point neighborhood for voxel i in the 3D lattice, $\text{TV}(x)$ is the discretized total variation functional and $\|x\|_1$ denotes the ℓ^1 -norm of the vector x . The total variation prior is known to favor piecewise regular solutions, in which different tissues are separated to a few subregions with short boundaries and small variation in the attenuation parameter within each subregion. The ℓ^1 -prior is known to favor solutions which consist of a few small high-attenuation targets on low-attenuation background. The parameters α_0 and α_1 play the role of prior parameters and they can be used to tune the relative weighting of the ℓ^1 and TV-functionals. The prior model (8)-(11) corresponds in qualitative sense closely to oral structures, which are expected to consist of a few subregions of different tissues (enamel, bone, soft tissue, air) with crisp boundaries. The features and the use of TV and ℓ^1 priors in inverse problems and image enhancement have been discussed, in (Dobson and Santosa, 1993; Dobson and Santosa, 1996; Donoho and Johnstone, 1992; Kaipio and Somersalo, 2004).

Using Bayes theorem and the models (7) and (8), the posterior density assumes the form

$$p(x|m) \propto p_+(x) \exp\left(-\frac{1}{2} \|L(m - Ax)\|^2 - W(x)\right), \quad (12)$$

where normalization constants are omitted. The computation of the MAP estimate from (12) amounts to solving the constrained optimization problem

$$x_{\text{MAP}} = \arg \min_{x \geq 0} \left\{ \frac{1}{2} \|L(m - Ax)\|^2 + W(x) \right\}$$

D. Computation of the MAP estimate

A major difficulty in Bayesian inversion for 3D X-ray imaging is the scale of the problem. At clinically relevant resolution the number of unknowns is in the range $10^6 - 10^7$ and computing the MAP estimate is computationally intensive.

We compute the MAP estimate by a gradient-based optimization technique introduced by Barzilai and Borwein (Barzilai and Borwein, 1988; Raydan, 1997). This method was chosen because the computation of the gradient for the objective functional is reasonably fast and the method can be parallelized efficiently.

We face two difficulties. First, the functional $W(x)$ is not differentiable because of the absolute value function. To overcome this problem, we use the smooth approximation

$$|t| \approx h_\beta(t) = \frac{1}{\beta} \log(\cosh(\beta t)), \quad (13)$$

where $\beta > 0$ is a parameter. The approximate differentiable prior functional is denoted by $W_\beta(x)$. The second difficulty arises from the positivity constraint. The use of constrained optimization methods is slow in high-dimensional problems. Thus, we take the positivity prior into account by applying the

exterior-point search methods (Fiacco and McCormick, 1990): the constrained problem is approximated by a sequence of unconstrained problems

$$\begin{aligned} x_{\text{MAP}}^{(j)} &= \arg \min \left\{ \frac{1}{2} \|L(m - Ax^{(j)})\|^2 \right. \\ &\quad \left. + W_\beta(x^{(j)}) + E^{(j)}(x^{(j)}) \right\}, \end{aligned} \quad (14)$$

where j is used to denote the j^{th} problem in the sequence and $E^{(j)}(x^{(j)})$ penalizes negative components of the solution $x^{(j)}$:

$$E^{(j)}(x^{(j)}) = \sum_{k=1}^M \phi^{(j)}(x_k^{(j)}), \quad (15)$$

where

$$\phi^{(j)}(x_k^{(j)}) = \begin{cases} \gamma_j (x_k^{(j)})^2, & x_k^{(j)} < 0 \\ 0, & x_k^{(j)} \geq 0 \end{cases}, \quad (16)$$

and $\{\gamma_j, j = 1, 2, \dots, N_s\}$ is a sequence of increasing positive penalty parameters. It should be noted that the exterior point methods guarantee the non-negativity of the solution only in the asymptotic limit $j \rightarrow \infty$.

The Barzilai-Borwein gradient-descent method consists of two steps: the update of the estimate and the computation of the step-length parameter. In the case of solving the problem (15), the update step assumes the form

$$x^{(j,t+1)} = x^{(j,t)} - s_t^{-1} g^{(j,t)} \quad (17)$$

where the update direction is the gradient of the argument in equation (15)

$$g^{(j,t)} = -A^T L^T L(m - Ax^{(j,t)}) + \nabla W_\beta(x^{(j,t)}) + \nabla E^{(j)}(x^{(j,t)}) \quad (18)$$

and the second superindex t denotes the iteration index. The second step, the computation of the step length parameter s_t is of the form

$$s_t = \frac{(x^{(j,t)} - x^{(j,t-1)})^T (g^{(j,t)} - g^{(j,t-1)})}{\|x^{(j,t)} - x^{(j,t-1)}\|^2}. \quad (19)$$

The stopping criteria for the optimization problem (15) is based on the magnitude of the gradient and the decrease in the optimization functional.

III. PARALLELIZED IMPLEMENTATION

A Beowulf cluster (Sterling et al., 1995) was constructed for 3D imaging problems. It consists of thirteen 3.0 GHz Pentium4 desktop-PCs each with 4 gigabytes of memory. The nodes are interconnected by a standard 1 Gb ethernet switch. The optimization algorithm (17)-(19) was implemented using the ANSI-C programming language and PETSc (Balay et al., 1995) application library, which in turn uses MPICH (Gropp et al., 1996) for interprocess communication and ATLAS (Whaley et al., 2001) for linear algebra routines.

In the first stage of the parallel algorithm, the observation matrix A is constructed. PETSc provides a parallel compressed sparse row matrix format, whose rows are distributed across computation nodes. Observation matrix creation is easily parallelized by distributing the x-ray source locations and detector

array geometries among the computation nodes, each of which calculates the respective ray lengths inside the voxel grid and insert the resulting rows in the matrix A independently of other nodes.

The second stage of the parallel algorithm implements the iteration (17)-(19) for the sequence of the exterior point problems (15). At this stage, the value and gradient of the modified prior functional $W_\beta(x)$ need to be evaluated and new update directions calculated. The PETSc library provides a framework for distributing the image vector x across the computation nodes. The framework hides the details of fetching the neighboring voxels that possibly reside on other nodes and hence the implementation of the prior functional $W_\beta(x)$ or other similar type Markov Random Field (MRF) functionals is straightforward. The evaluation of the likelihood gradient, which is the first term in Eq. (18), and the evaluation of residual $\|L(m - Ax)\|^2$ for stopping criteria include matrix-vector products. These are implemented by the PETSc tools that distribute the computation of a matrix-vector product evenly on the computation nodes.

IV. EVALUATION

The proposed method is tested with extraoral projection data acquired from a male patient using a commercial panoramic x-ray device. Tomosynthetic reconstructions (Webber et al., 1997; Webber and Messura, 1999) are shown as reference for the reconstructions with the Bayesian method.

The test case is an example of limited angle 3D imaging in implant planning. In such situations, it is often important to get accurate measurement for the depth and location of the mandibular nerve canal. In the imaging experiment the source and detector are rotated around the head and images of the region of interest (ROI) are taken through the head. The imaging geometry is shown in Figure 2. The projection directions were chosen such that the cross section of the nerve canal would be approximately perpendicular to the projection directions.




Fig. 2. Geometry for limited angle 3D dental imaging in implantology. The source positions are denoted by circles. The thick arc illustrates the jaw bone and dental arc. The region of interest (ROI) is denoted by thin line. For clarity, the detector is depicted only for one source location.

Eleven equispaced projection radiographs of the patient were acquired with a total view angle of 40° . Figure 3 shows

one extraoral projection image of the patient. The size of projection image is 876×876 with pixel size $0.09\text{mm} \times 0.09\text{mm}$. The projection images were downsampled to half the original size before transforming them into tomographic data. Thus, the number of data in the inverse problem was $N = 11 \times 438 \times 438 \approx 2.1 \cdot 10^6$.




Fig. 3. One (mandibular) projection radiograph of the patient.

The results are shown in Figure 4. A vertical slice from tomosynthetic reconstruction is shown on the left, and the respective slice from the 3D MAP estimate on the right. The volume of the domain Ω in the computations was approximately $79\text{mm} \times 79\text{mm} \times 63\text{mm}$ and the voxel size in the reconstruction was $0.38\text{mm} \times 0.38\text{mm} \times 0.38\text{mm}$. With this resolution, the number of unknowns in the inverse problem was $M = 207 \times 207 \times 167 \approx 7.2 \cdot 10^6$. For the prior parameters we used values $\alpha_0 = 10$ and $\alpha_1 = 1$. The smoothing parameter for the approximation of the absolute value function was $\beta = 200$ and the sequence of the exterior point parameters $\{\gamma_j, j = 1, \dots, 5\}$ were in the range from 3000 to 10. Each problem in the sequence of the five exterior-point problems was iterated until convergence. The convergence was verified by monitoring the norm of the gradient $g^{(j,t)}$ and the change in the residual (i.e., the value of the optimization functional in Eq. (15)); the iteration was stopped once either of these figures got below predefined thresholds, or maximum number of iterations was reached. For the norm of the gradient we used threshold value 1, for the change in residual threshold value 0.1 and the maximum number of iterations was 6.

The computation of the MAP estimate took 3 minutes 26 seconds. The location of the mandibular nerve canal in the MAP reconstruction is denoted by an arrow. As can be seen, the nerve canal and the interfaces of solid and spongy bone are seen more clearly in the MAP estimate than in the tomosynthetic slice.

V. CONCLUSION

We introduce a Bayesian model for 3D dental imaging with sparse projection data and implement the computation of the MAP estimate for the 3D attenuation function with parallel computing techniques in a Beowulf cluster computer.




Fig. 4. Reconstructions from patient data (11 projections spanning a view angle of 40°). Left: Vertical slice of tomosynthetic reconstruction. Right: Corresponding slice from the 3D MAP-estimate. The arrow denotes the location of the mandibular nerve channel.

The proposed method shows important anatomical details more clearly than tomosynthesis, which is currently the method of choice for 3D dental imaging with limited data. Also, parallelization speeds the computation up to clinically acceptable level.

REFERENCES

- Brocklebank, L., (1997): *Dental Radiology — Understanding the X-Ray Image*. Oxford University Press. ISBN 0-19-262411-3.
- Ekestubbe, A., Gröndahl, K., Gröndahl, H.-G., (1997): The use of tomography for dental implant planning. *Dentomaxillofacial Radiology* 26: 206–213.
- Ramesh, A., Ludlow, J. B., Webber, R. L., Tyndall, D. A., Paquette, D., (2002): Evaluation of tuned-aperture computed tomography in the detection of simulated periodontal defects. *Oral and Maxillofacial Radiology* 93: 341–349.
- PaloDEX Group (Finland), (2007): *Volumetric Tomography*, [On Line], Available at <http://www.instrumentariumdental.com/>
- Ranggayyan, R. M., Dhawan, A. T., Gordon, R., (1985): Algorithms for limited-view computed tomography: an annotated bibliography and a challenge. *Appl. Optics* 24: 4000–4012.
- Natterer, F., (1986): *The Mathematics of Computerized Tomography*. John Wiley & Sons, Chichester, USA.
- Hanson, K. M., (1987): Bayesian and related methods in image reconstruction from incomplete data. *Image Recovery: Theory and Applications*. Academic, Orlando.
- Sauer, K., James, S. Jr, Klifa, K., (1994): Bayesian estimation of 3-D objects from few radiographs. *IEEE Trans. Nucl. Sci.* 41: 1780–1790.
- Hanson, K. M., Wecksung, G. W., (1983): Bayesian approach to limited-angle reconstruction in computed tomography. *J. Opt. Soc. Am.* 73: 1501–1509.
- Bouman, C., Sauer, K., (1993): A generalized Gaussian image model for edge-preserving MAP estimation. *IEEE Trans. Image Processing* 2: 296–310.
- Frese, T., Bouman, C., Sauer, K., (2002): Adaptive wavelet graph model for Bayesian tomographic reconstruction. *IEEE Trans. Image Processing* 11: 756–770.
- Yu, D. F., Fessler, J. A., (2002): Edge-preserving tomographic reconstruction with nonlocal regularization. *IEEE Trans. Med. Imag.* 21: 159–173.
- Siltanen, S., Kolehmainen, V., Järvenpää, S., Kaipio, J. P., Koistinen, P., Lassas, M., Pirttilä, J., Somersalo, E., (2003): Statistical inversion for medical x-ray tomography with few radiographs I: General theory. *Phys. Med. Biol.* 48: 1437–1463.
- Kolehmainen, V., Siltanen, S., Järvenpää, S., Kaipio, J. P., Koistinen, P., Lassas, M., Pirttilä, J., Somersalo, E., (2003): Statistical inversion for medical x-ray tomography with few radiographs II: Application to dental radiology. *Phys. Med. Biol.* 48: 1465–1490.
- Kolehmainen, V., Vanne, A., Siltanen, S., Järvenpää, S., Kaipio, J. P., Lassas, M. and Kalke, M., (2006): Parallelized Bayesian inversion for three-dimensional dental X-ray imaging. *IEEE Trans. Med. Im.* 25: 218–228.
- Zheng, J., Saquib, S. S., Sauer, K., Bouman, C. A., (2000): Parallelizable Bayesian tomography algorithms with rapid, guaranteed convergence. *IEEE Trans. Image Processing* 9: 1745–1759.
- Mosegaard, M., Sambridge, M., (2002): Monte Carlo analysis of inverse problems. *Inv. Probl.* 18: R29–R54.
- Kaipio, J. P., Somersalo, E., (2004): *Statistical and Computational Methods for Inverse Problems*. Number 160 in Applied Mathematical Sciences. Springer-Verlag, New York. ISBN: 0-387-22073-9.
- Dobson, D. C., Santosa, F., (1994): An image enhancement technique for electrical impedance tomography. *Inv. Probl.* 10: 317–334.
- Dobson, D. C., Santosa, F., (1996): Recovery of blocky images from noisy and blurred data. *SIAM J Appl Math* 56: 1181–1198.
- Donoho, D.L., Johnstone, I. M., Hoch, J. C., Stern, A. S., (1992): Maximum entropy and the nearly black object. *J Roy Statist Ser B* 54: 41–81.
- Barzilai, J., Borwein, J. M., (1988): Two point step size gradient method. *IMA J. Numer. Anal.* 8: 141–148.
- Raydan, M., (1997): The Barzilai and Borwein gradient method for the large scale unconstrained minimization problem. *SIAM J. Optim.* 7: 26–33.
- Fiacco, A. V., McCormick, G. P., (1990): *Nonlinear programming: sequential unconstrained minimization techniques*. SIAM.
- Sterling, T., Savarese, D., Becker, D. J., Dorband, J. E., Ranawake, U. A., Packer, C. V., (1995): BEOWULF: A parallel workstation for scientific computation. In *Proceedings of the 24th International Conference on Parallel Processing*, number I.
- Balay, S., Buchelman, K., Eijkhout, V., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., Zhang, H., (1995): *PETSc Users Manual*. Argonne National Laboratory ANL-95/11, revision 2.1.5 edition.
- Gropp, W., Lusk, E., Doss, N., Skjellum, A., (1996): A high-performance, portable implementation of the MPI message passing interface standard. *Parallel Computing* 22: 789–828.
- Whaley, R. C., Petitet, A., Dongarra, J. J., (2001): Automated empirical optimization of software and the ATLAS project. *Parallel Computing* 27: 3–35.
- Webber, R. L., Horton, R. A., Tyndall, D. A., Ludlow, J. B., (1997): Tuned aperture computed tomography (TACT). Theory and application for three-dimensional dento-alveolar imaging. *Dentomaxillofacial Radiology* 26: 53–62.
- Webber, R. L. and Messura, J. K., (1999): An in vivo comparison of diagnostic information obtained from tuned-aperture computed tomography and conventional dental radiographic imaging modalities. *Oral Surgery Oral Medicine Oral Pathology* 88: 239–247.
- Hanson, K. M., Cunningham, G. S., McKee, R., (1997): Uncertainty assessment for reconstructions based on deformable geometry. *Int. J. Imaging Syst. Technol.* 8: 506–512.