Constrained X-ray tensor tomography reconstruction

Jakob Vogel,[1] **Florian Schaff,**[2] **Andreas Fehringer,**[2] **Christoph Jud,**[2] **Matthias Wieczorek,**[1] **Franz Pfeiffer**[2] **and Tobias Lasser**[1,*]

[1] *Computer Aided Medical Procedures, Fakultät für Informatik, Technische Universität München, Germany*

[2] *Lehrstuhl für Biomedizinische Physik, Physik-Department & Institut für Medizintechnik, Technische Universität München, Germany*

[*] *lasser@cs.tum.edu*

Abstract: Quite recently, a method has been presented to reconstruct X-ray scattering tensors from projections obtained in a grating interferometry setup. The original publications present a rather specialised approach, for instance by suggesting a single SART-based solver. In this work, we propose a novel approach to solving the inverse problem, allowing the use of other algorithms than SART (like conjugate gradient), a faster tensor recovery, and an intuitive visualisation. Furthermore, we introduce constraint enforcement for X-ray tensor tomography (cXTT) and demonstrate that this yields visually smoother results in comparison to the state-of-art approach, similar to regularisation.

© 2015 Optical Society of America

OCIS codes: (110.3010) Image reconstruction techniques; (110.6960) Tomography; (110.6955) Tomographic imaging; (340.7440) X-ray imaging; (340.7450) X-ray interferometry.

References and links

1. F. Pfeiffer, T. Weitkamp, O. Bunk, and C. David, "Phase retrieval and differential phase-contrast imaging with low-brilliance X-ray sources," Nat. Phys. **2**, 258–261 (2006).
2. F. Pfeiffer, M. Bech, O. Bunk, P. Kraft, E. F. Eikenberry, C. Brönnimann, C. Grünzweig, and C. David, "Hard-X-ray dark-field imaging using a grating interferometer," Nat. Mater. **7**, 134–137 (2008).
3. A. Yaroshenko, F. G. Meinel, M. Bech, A. Tapfer, A. Velroyen, S. Schleede, S. Auweter, A. Bohla, A. Ö. Yildirim, K. Nikolaou, F. Bamberg, O. Eickelberg, M. F. Reiser, and F. Pfeiffer, "Pulmonary emphysema diagnosis with a preclinical small-animal X-ray dark-field scatter-contrast scanner," Radiol. **269**, 427–433 (2013).
4. M. Ando, K. Yamasaki, F. Toyofuku, H. Sugiyama, C. Ohbayashi, G. Li, L. Pan, X. Jiang, W. Pattanasiriwisawa, D. Shimao, E. Hashimoto, T. Kimura, M. Tsuneyoshi, E. Ueno, K. Tokumori, A. Maksimenko, Y. Higashida, and M. Hirano, "Attempt at visualizing breast cancer with X-ray dark field imaging," Jpn. J. Appl. Phys. **44**, L528–L531 (2005).
5. S. Sidhu, G. Falzon, S. A. Hart, J. G. Fox, R. A. Lewis, and K. K. W. Siu, "Classification of breast tissue using a laboratory system for small-angle x-ray scattering (SAXS)," Phys. Med. Biol. **56**, 6779–6791 (2011).
6. F. Arfelli, L. Rigon, and R. H. Menk, "Microbubbles as X-ray scattering contrast agents using analyzer-based imaging," Phys. Med. Biol. **55**, 1643–1658 (2010).
7. V. Revol, I. Jerjen, C. Kottler, P. Schütz, R. Kaufmann, T. Lüthi, U. Sennhauser, U. Straumann, and C. Urban, "Sub-pixel porosity revealed by X-ray scatter dark field imaging," J. Appl. Phys. **110**, 044912 (2011).
8. A. Malecki, G. Potdevin, T. Biernath, E. Eggl, K. Willer, T. Lasser, J. Maisenbacher, J. Gibmeier, A. Wanner, and F. Pfeiffer, "X-ray tensor tomography," EPL **105** (2014).
9. F. L. Bayer, S. Hu, A. Maier, T. Weber, G. Anton, T. Michel, and C. P. Riess, "Reconstruction of scalar and vectorial components in X-ray dark-field tomography," Proc. Natl. Acad. Sci. U.S.A. **111**, 12699–12704 (2014).
10. A. Filler, "Magnetic resonance neurography and diffusion tensor imaging: Origins, history, and clinical impact of the first 50 000 cases with an assessment of efficacy and utility in a prospective 5000-patient study group," Neurosurg. **65**, A29–A43 (2009).

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

11. J. Vogel, M. Wieczorek, C. Jud, F. Schaff, F. Pfeiffer, and T. Lasser, "X-ray tensor tomography reconstruction," Proc. Fully3D (2015).
12. M. Bech, O. Bunk, T. Donath, R. Feidenhans'l, C. David, and F. Pfeiffer, "Quantitative X-ray dark-field computed tomography," Phys. Med. Biol. **55**, 5529–5539 (2010).
13. V. Revol, C. Kottler, R. Kaufmann, A. Neels, and A. Dommann, "Orientation-selective X-ray dark field imaging of ordered systems," J. Appl. Phys. **112**, 114903 (2012).
14. A. Malecki, G. Potdevin, T. Biernath, E. Eggl, E. Grande Garcia, T. Baum, P. B. Noël, J. S. Bauer, and F. Pfeiffer, "Coherent superposition in grating-based directional dark-field imaging," PLoS ONE **8**, e61268 (2013).
15. A. Fehringer, T. Lasser, I. Zanette, P. B. Noël, and F. Pfeiffer, "A versatile tomographic forward- and backprojection approach on multi-GPUs," Proc. SPIE **9034**, 90344F (2014).
16. A. H. Andersen and A. C. Kak, "Simultaneous algebraic reconstruction technique (SART): A superior implementation of the ART algorithm," Ultrasonic Imaging **6**, 81–94 (1984).
17. H. H. B. Sørensen and P. C. Hansen, "Multicore performance of block algebraic iterative reconstruction methods," SIAM J. Sci. Comp. **36**, C524–C546 (2014).
18. M. R. Hestenes and E. Stiefel, "Methods of conjugate gradients for solving linear systems," J. Res. Nat. Bur. Stand. **49**, 409–436 (1952).
19. J. R. Shewchuk, "An introduction to the conjugate gradient method without the agonizing pain," (1994).
20. Q. Li and J. G. Griffiths, "Least squares ellipsoid specific fitting," Proc. Geom. Model. Process. pp. 335–340 (2004).
21. K. Pearson, "On lines and planes of closes fit to systems of points in space," Philos. Mag. **2**, 559–572 (1901).
22. H. Hotelling, "Analysis of a complex of statistical variables into principal components," J. Educ. Psychol. **24**, 417–441 (1933).
23. C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer, 2006), chap. 12.1.
24. C. Runge, "Über die numerische Auflösung von Differentialgleichungen [About numerically solving differential equations]," Math. Ann. **46**, 167–178 (1895).
25. M. W. Kutta, "Beitrag zur näherungsweisen Integration totaler Differentialgleichungen [Contribution to the approximate integration of total differential equations]," Z. Math. Phys. **46**, 435–453 (1901).
26. C. H. Edwards and D. E. Penney, *Differential Equations* (Pearson, 2007), chap. 4.3.
27. C. R. Tench, P. S. Morgan, M. Wilson, and L. D. Blumhardt, "White matter mapping using diffusion tensor MRI," Magn. Reson. Med. **47**, 967–972 (2002).
28. M. Wieczorek, J. Vogel, and T. Lasser, "CampRecon — a software framework for linear inverse problems," Tech. Rep. TUM-I1401, TU München (2014).
29. R. L. Siddon, "Fast calculation of the exact radiological path for a three-dimensional CT array," Med. Phys. **12**, 252–255 (1985).

1. Introduction and background

Most applications of X-ray imaging limit themselves to attenuation, providing images with strong contrast between highly absorbing structures such as bones or artificially contrasted arteries, and lowly absorbing background such as tissue. From a physical point of view, however, more information could be acquired, for instance by means of X-ray interferometry. For decades, this technique was limited to laboratory setups, as it requires highly coherent X-ray illumination, and thus a synchrotron or an equivalent source.

Only some few years ago, Pfeiffer *et al.* [1] have demonstrated a novel setup for X-ray grating interferometry: As shown in Fig. 1, an 'ordinary' setup of X-ray tube (T), specimen (S) and detector (D) has been extended by inserting a *source grating* (G_0) after the tube, making the incoming illumination partly coherent, and two more gratings (G_1, G_2) forming the *interferometer* itself. Instead of a single image per orientation/angle as would be done in 'classical' computed tomography, multiple images are acquired while interferometry grating G_1 is shifted sideways, thus producing several images of the same situation with different interference patterns. Based on such images, three signal components can be extracted per perspective: An *absorption* image as commonly used in X-ray computed tomography, a *phase-contrast* and a *dark-field* component [2]. Example images are given in Fig. 2.

Despite certain problems such as phase wrapping, tomographic phase-contrast reconstructions have since become quite popular, particularly due to their improved soft-tissue contrast, and improving them is an active field of research.

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

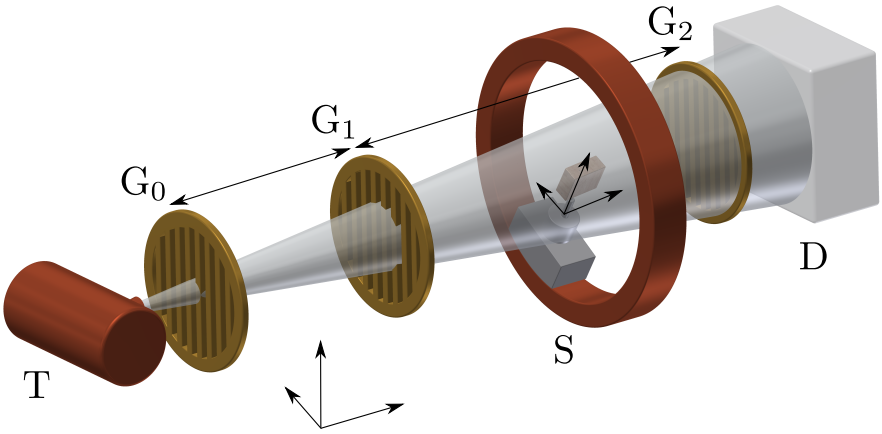

Fig. 1. Sketch of an X-ray grating interferometry setup. X-ray tube (T), source grating (G_0), shifting interferometer grating (G_1), specimen (S), static interferometer grating (G_2), and detector (D).

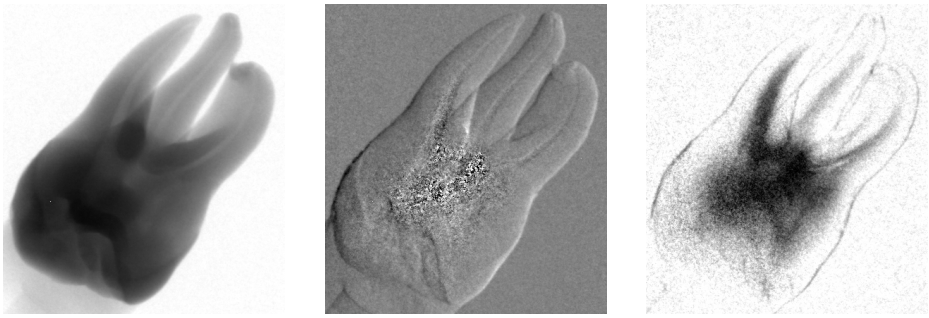

Fig. 2. Three components of a sample projection of a tooth: *Absorption*, *phase contrast* and *dark field*. The contrast of these images has been manually improved for better visibility, and the images have been cropped. The structure in the lower-left quadrant is the sample holder.

A straight-forward three-dimensional tomographic reconstruction of the dark-field signal yields a volumetric representation of the respective X-ray scattering magnitudes. Such information can be of practical value when isotropic scattering is of interest. Previous studies have shown its potential for lung and breast imaging [3–5], micro-bubble contrast agents [6] and material testing [7]. In all these cases, dark-field reconstructions have been shown to provide contrast where the other two signals, phase-contrast and absorption, only yield poor results at best.

Tomographic reconstruction of the anisotropic component of the dark-field signal on the other side has not yet gained much attention. Only recently, Malecki *et al.* [8] have presented *tensor*-valued *tomographic reconstruction*, thus taking directional information into account. In addition to the usual tomographic axis, the sample is also rotated around the other two axes using an Euler cradle. The authors present a first mathematical model and a specialised SART-based reconstruction algorithm, thus obtaining the first scattering tensor reconstructions. Informally speaking, they obtain an ellipsoid at every voxel, and its size and shape hints at the structure of the specimen at this location. In particular, plate-like scattering ellipsoids hint at

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

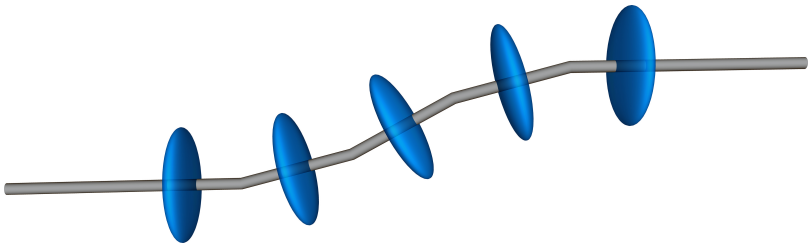

Fig. 3. Expected X-ray scattering (blue ellipsoids) at fibre- or tube-like structures (grey). Vice versa, we will interpret the smallest half-axes of reconstructed scattering ellipsoids as fibre directions for visualisation.

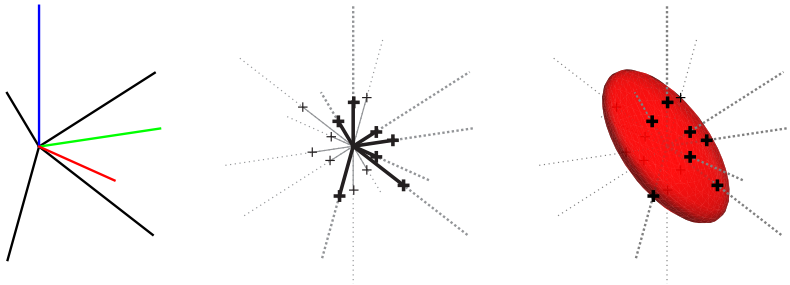

(a) Sampling directions. (b) Coefficients. (c) Ellipsoid fit.

Fig. 4. Sampling directions, and ellipsoid fitting for a single voxel. Usually, we use $K = 13$ directions for reconstruction, but limit ourselves to $K = 7$ in this sketch, for clarity. They consist of the standard base vectors – red, green, blue in (a) for x, y, z, respectively – and diagonals – black. For a given voxel $\mathbf{x}$, reconstruction yields a scattering coefficient $\zeta_k(\mathbf{x})$ for every sampling direction $\hat{\mathbf{\varepsilon}}_k$. These are indicated by bold black marks in (b), and we mirror them along the negative sampling direction, yielding the small black marks. Finally, an ellipsoid can be fitted to that scaled, mirrored 'bouquet' afterwards, see (c).

fibre- or tube-like structures, as sketched in Fig. 3.

The only other work on scattering tensor reconstruction has recently been presented by Bayer *et al.* [9]. In contrast to the aforementioned approach, this group uses the single tomographic axis of rotation only and relies on extended mathematical modelling. In particular, they apparently do not reconstruct full scattering tensors, but projections of fibre directions onto the tomographic plane, and thus a stack of two-dimensional projections rather than a volume of three-dimensional shapes. Apart from X-rays, another well-known tensor-based imaging modality is diffusion tensor magnetic resonance imaging (DTI) [10] – which only shares visualisation techniques.

In this paper, we use the approach of Malecki *et al.*, and present several algorithmic improvements. Among them are, most importantly, a novel, more generic description for solving X-ray tensor tomography (XTT) problems [11], and two ways to enforce tensor-shapes during reconstruction.

2. Methods

Scattering is not a scalar entity (as X-ray attenuation is), but a tensor-valued one, that is a three-dimensional shape per location. This requires a more complicated mathematical model.

Malecki *et al.* [8] propose to consider a finite set of K pre-defined, normalised, well-

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

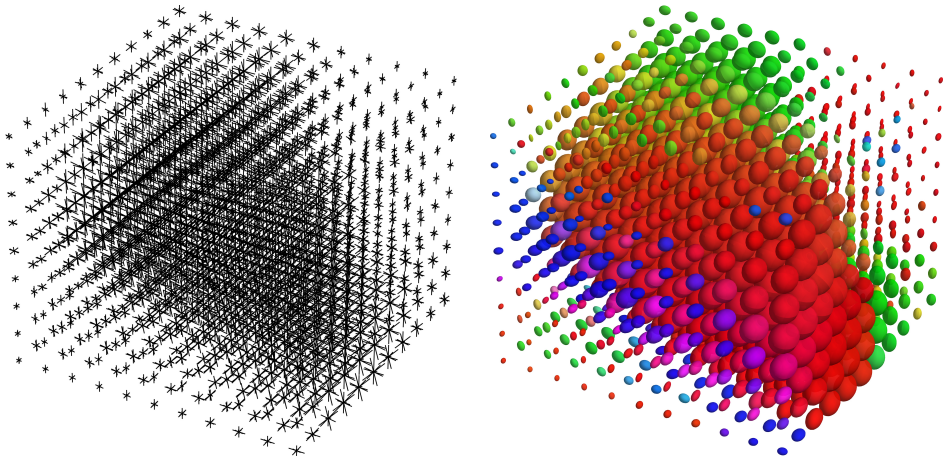

(a) Scaled sampling vectors.

(b) Ellipsoids with axes.

Fig. 5. Different variants of the same data, a downsampled detail from an actual result. Reconstruction is performed in terms of scaling the sampling directions (a), the scattering tensors are obtained by retrospectively fitting ellipsoids in a voxel-wise fashion (b). The colours indicate the directions of the smallest half-axes, see Fig. 8(a).

distributed *sampling directions* $\hat{\boldsymbol{\varepsilon}}_k \in \mathbb{R}^3$ rather than a tensor. Here, and in the remainder of the paper, we use the 'hat' to denote normalised vectors, that is $\|\hat{\boldsymbol{\varepsilon}}_k\|_2 = 1 \ \forall k = 1,\ldots,K$. For each of these $\hat{\boldsymbol{\varepsilon}}_k$ and every voxel $\mathbf{x}_i$, a corresponding *scattering coefficient* $\zeta_k(\mathbf{x}_i) \in \mathbb{R}$ will be reconstructed, and the tensor itself can later be approximated by voxel-wise fitting of tensor descriptions to these 'bouquets' of weighted directions. This entire process is sketched for a single voxel in Fig. 4, a visual account of a larger region is given in Fig. 5. Note that the exact orientation of the sampling directions is rather arbitrary, and the choice does not require any knowledge about the expected reconstruction outcome. The sampling directions are just auxiliary, virtual entities, allowing to avoid an explicit description of the scattering tensor at a given location by considering its projection onto a finite set of well-distributed reference directions instead. They are chosen from a single hemisphere only, as the scattering tensors are supposed to be symmetric.

Based on this approach, Malecki *et al.* derive a forward model for X-ray scattering [12–14], resembling the *Beer-Lambert law* describing X-ray absorption:

$$d_j = \exp\left[-\int_{L_j} \sum_k \left\langle |\hat{\mathbf{l}}_j \times \hat{\boldsymbol{\varepsilon}}_k| (\zeta_k(\mathbf{x})\hat{\boldsymbol{\varepsilon}}_k), \hat{\mathbf{t}}_j \right\rangle^2 \, \mathrm{d}\mathbf{x}\right] \tag{1}$$

Here, $d_j \in \mathbb{R}$ denotes the j'th scalar *dark-field measurement*, and L_j the corresponding ray with normalised direction $\hat{\mathbf{l}}_j \in \mathbb{R}^3$. Note that despite acquiring two-dimensional projections, we treat the individual pixels as independent measurements with corresponding rays. Index $j = 1,\ldots,J$ extends over all pixels of all projection images hence. The normalised *sensitivity direction* $\hat{\mathbf{t}}_j \in \mathbb{R}^3$ is orthogonal to the grating lines and parallel to the surfaces of the – mutually parallel – gratings, and thus depends on the perspective during measurement j.

Considering Eq. (1) in more detail, the true scattering measurement at every voxel $\mathbf{x}$ along the ray is generally modelled as finite sum of K measurements along the pre-defined sampling directions $\hat{\boldsymbol{\varepsilon}}_k$. For every such summand, two things need to be taken into account: Scattering must be possible to occur at all, considering the direction $\hat{\mathbf{l}}_j$ of the incoming X-ray with re-

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

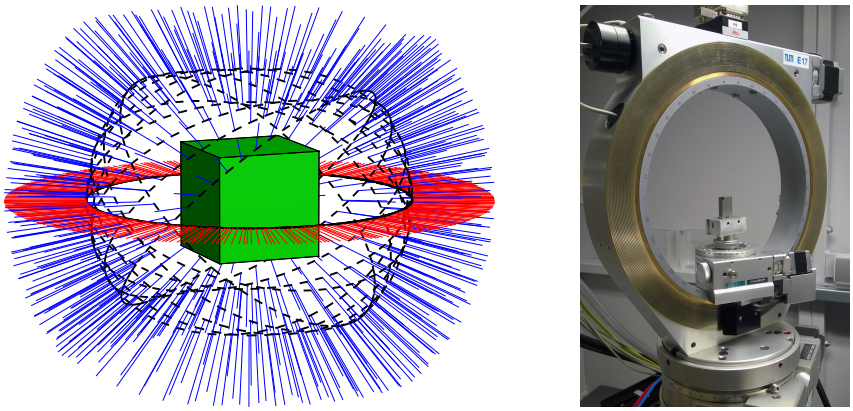

Fig. 6. Viewing directions and Euler cradle. Only the perspectives marked by red lines are in the 'normal' viewing plane as used in standard CT applications. Blue lines mark additional off-plane perspectives. The dashed black line indicates the trajectory of the X-ray source-detector "camera" around the green cube, representing a specimen. The coverage gaps are due to limitations imposed by the Euler cradle.

spect to $\hat{\boldsymbol{\varepsilon}}_k$. This fact is modelled by the magnitude of the cross-product $|\cdot\times\cdot|$, thus excluding head-on views. Furthermore, if scattering occurs, only components along sensitivity direction $\hat{\mathbf{t}}_j$ contribute to the measurement as modelled by the scalar product $\langle\cdot,\cdot\rangle$ [13, 14]. Vice versa, in order to have a good sampling of the scattering tensors, more than just the usual tomographic scanning perspectives are taken into consideration while scanning, see Fig. 6.

Returning to Eq. (1), factoring out the (unknown) *squared scattering coefficients* $\eta_k(\mathbf{x}) := \zeta_k(\mathbf{x})^2$ from the squared scalar product and defining the weight factor

$$v_{kj} := (|\hat{\mathbf{l}}_j \times \hat{\boldsymbol{\varepsilon}}_k|\langle\hat{\boldsymbol{\varepsilon}}_k, \hat{\mathbf{t}}_j\rangle)^2 \tag{2}$$

yields a formulation very similar to the well-known *Radon transform*:

$$-\ln d_j = \int_{L_j} \sum_k v_{kj}\cdot\eta_k(\mathbf{x})\,\mathrm{d}\mathbf{x} \tag{3}$$

$$= \sum_k v_{kj} \int_{L_j} \eta_k(\mathbf{x})\,\mathrm{d}\mathbf{x} \tag{4}$$

Note that the factors defined in Eq. (2) are independent of the unknowns, and can thus be precomputed for more efficient reconstruction.

2.1. *Reconstruction*

Obviously, Eq. (4) is a weighted sum of line integrals. In particular, in a discretised setting, it can be rewritten as scalar product

$$m_j = -\ln d_j = \sum_k v_{kj}\,\langle\mathbf{a}_j, \boldsymbol{\eta}_k\rangle = \sum_k v_{kj}\,\mathbf{a}_j^T\,\boldsymbol{\eta}_k \tag{5}$$

where $\boldsymbol{\eta}_k \in \mathbb{R}^I$ is the vector of the squared coefficients for the kth sampling direction for all voxels (I is the number of voxels). $\mathbf{a}_j$ denotes the *system matrix* row for measurement j, and is thus part of the 'standard' matrix as used for an equivalent attenuation reconstruction. This vector contains geometric information about the arrangement of X-ray source, specimen and

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

sensor during measurement j, an information already included abstractly in Eq. (4) as ray L_j. Let $A = (\mathbf{a}_j)$, $A \in \mathbb{R}^{J \times I}$, denote the entire system matrix describing *all* J line integrals. Furthermore, let $D_k = \mathrm{diag}(v_{k1}, v_{k2}, \dots)$ denote a diagonal scaling matrix containing the weighting factors from Eq. (2) for sampling direction k. Then, using Eq. (5) and defining the measurement vector $\mathbf{m} = (m_j) = (-\ln\, d_j)$, a huge linear system can be derived:

$$\mathbf{m} = \begin{pmatrix} v_{11}\mathbf{a}_1^T \\ v_{12}\mathbf{a}_2^T \\ \vdots \\ v_{1J}\mathbf{a}_J^T \end{pmatrix} \boldsymbol{\eta}_1 + \begin{pmatrix} v_{21}\mathbf{a}_1^T \\ v_{22}\mathbf{a}_2^T \\ \vdots \\ v_{2J}\mathbf{a}_J^T \end{pmatrix} \boldsymbol{\eta}_2 + \cdots + \begin{pmatrix} v_{K1}\mathbf{a}_1^T \\ v_{K2}\mathbf{a}_2^T \\ \vdots \\ v_{KJ}\mathbf{a}_J^T \end{pmatrix} \boldsymbol{\eta}_K \tag{6}$$

$$= \begin{pmatrix} v_{11} & & & \\ & v_{12} & & \\ & & \ddots & \\ & & & v_{1J} \end{pmatrix} \begin{pmatrix} \mathbf{a}_1^T \\ \mathbf{a}_2^T \\ \vdots \\ \mathbf{a}_J^T \end{pmatrix} \boldsymbol{\eta}_1 + \begin{pmatrix} v_{21} & & & \\ & v_{22} & & \\ & & \ddots & \\ & & & v_{2J} \end{pmatrix} \begin{pmatrix} \mathbf{a}_1^T \\ \mathbf{a}_2^T \\ \vdots \\ \mathbf{a}_J^T \end{pmatrix} \boldsymbol{\eta}_2 + \cdots \tag{7}$$

$$= D_1 A\, \boldsymbol{\eta}_1 + D_2 A\, \boldsymbol{\eta}_2 + \cdots + D_K A\, \boldsymbol{\eta}_K = \sum_k D_k A\, \boldsymbol{\eta}_k \tag{8}$$

$$= (D_1 A,\, D_2 A,\, \dots,\, D_K A) \begin{pmatrix} \boldsymbol{\eta}_1 \\ \boldsymbol{\eta}_2 \\ \vdots \\ \boldsymbol{\eta}_K \end{pmatrix} \tag{9}$$

$$\equiv H\mathbf{s} \tag{10}$$

This linear system has K times as many unknowns than the corresponding system for computing a 'traditional' tomographic attenuation reconstruction: The original system matrix A is of size $J \times I$, and H is of size $J \times IK$. In practical settings, that is when reconstructing a sufficiently sized three-dimensional volume, the system matrix A is already far too large to fit into computer memory. Instead of handling the matrix directly, *projector* software is typically used that simulates the ray, thus computing the entries of A on the fly [15].

When reconstructing scattering tensors, it is desirable to use existing software infrastructure that has been put into place for other tomographic reconstruction scenarios rather than solving Eq. (10) directly. Consequently, in their original work, Malecki *et al.* suggest a specially crafted variant of SART [16, 17] for reconstructing the unknown squared coefficients $\boldsymbol{\eta}_k$.

Using Eq. (5), however, a much more generic approach can be taken, supporting arbitrary iterative solvers such as particularly the better-behaving method of Conjugate Gradients (CG) [18, 19]. Let $x_1 = \textbf{approximate}(A, b, x_0)$ denote an auxiliary function running a *single* iteration of an iterative linear solver such as CG or SART, thus very approximately 'solving' the linear system $Ax = b$. Then, an iterative algorithm solving the tensor reconstruction problem (5) can be defined: A single iteration q essentially consists of 'approximately solving' K modified linear systems

$$(D_k \cdot A)\, \mathbf{t}_k^{(q)} = \widetilde{\mathbf{m}}_k^{(q-1)} \tag{11}$$

for each sampling direction k using function **approximate** as previously defined, where right-hand side vector

$$\widetilde{\mathbf{m}}_k^{(q-1)} = \mathbf{m} - \sum_{l \neq k} D_l \cdot A \cdot \boldsymbol{\eta}_l^{(q-1)} \tag{12}$$

denotes a *reduced measurement vector* relevant for sampling direction k, based on the estimates $\boldsymbol{\eta}_l^{(q-1)}$ of the previous iteration, and $\mathbf{t}_k^{(q)}$ an intermediate, temporary vector. The latter is then

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

used to compute the current iterate via relaxation as

$$\boldsymbol{\eta}_k^{(q)} = \frac{K-1}{K}\,\boldsymbol{\eta}_k^{(q-1)} + \frac{1}{K}\,\mathbf{t}_k^{(q)}\ . \tag{13}$$

Altogether, this proposed approach can be thought of as simultaneously solving K linear systems in an interleaved sense. This entire proposed generic tensor reconstruction approach is given in alg. 1. Note that the scattering coefficients need to be extracted after reconstruction of all $\boldsymbol{\eta}_k$ by component-wise application of the square root:

$$\zeta_k(\mathbf{x}) = \sqrt{|\eta_k(\mathbf{x})|} \tag{14}$$

Using the absolute magnitude of η_k is a valid security measure, considering the symmetry of the tensors.

Alg. 1 Generic tomographic X-ray tensor reconstruction. A denotes the system matrix describing the imaging process, D_k a scaling matrix containing weighting factors as defined in Eq. (2). $\mathbf{m}$ is the measurement vector, and $\boldsymbol{\eta}_k^{(q)}$ a vector containing the qth iterate of the voxel-wise squared scattering coefficients corresponding to sampling direction k. **approximate** is a function running a single iteration of an arbitrary iterative linear solver.

$\eta_k^{(0)} = \mathbf{0}\ \forall k$
for iterations q **do**
  // pre-compute weighted forward projections
  for sampling directions $k \in \{1,\ldots,K\}$ **do**
    $WFP_k = D_k \cdot A \cdot \boldsymbol{\eta}_k^{(q-1)}$
  end for

  for sampling directions $k \in \{1,\ldots,K\}$ **do**
    // compute right-hand side
    $RHS = \mathbf{m} - \sum_{l\neq k} WFP_l$

    // run a single solver iteration
    $TEMP = \textbf{approximate}(D_k \cdot A, RHS, \boldsymbol{\eta}_k^{(q-1)})$

    // update coefficients with relaxation
    $\lambda = 1/K$
    $\boldsymbol{\eta}_k^{(q)} = (1-\lambda)\,\boldsymbol{\eta}_k^{(q-1)} + \lambda\ TEMP$
  end for

  // optionally enforce constraints
end for

2.2. *Ellipsoid Fitting*

After reconstruction, that is after executing several iterations as described in section 2.1, voxel-wise scattering coefficients ζ_k for the sampling vectors $\hat{\boldsymbol{\varepsilon}}_k$ have been recovered. Malecki *et al.* propose to fit ellipsoids to these weighted vectors in order to obtain voxel-wise tensors. In particular, they propose to use an iterative ellipsoid fitter [20], apparently intended for rather degenerate cases where ellipsoids need to be matched to just a couple of ill-distributed sample points.

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

However, as shown in Fig. 4, the sampling locations are well-distributed, and we propose to use *principal component analysis* [21–23] (PCA) instead of the iterative approach, thus saving considerable computation time.

In more detail, at every voxel $\mathbf{x}_i$, a set of $2K$ direction vectors

$$S_i := \{\pm\sqrt{|\eta_1(\mathbf{x}_i)|}\cdot\hat{\varepsilon}_1,\ \pm\sqrt{|\eta_2(\mathbf{x}_i)|}\cdot\hat{\varepsilon}_2,\ \ldots\} \quad (15)$$

$$= \{\pm\zeta_1(\mathbf{x}_i)\cdot\hat{\varepsilon}_1,\ \pm\zeta_2(\mathbf{x}_i)\cdot\hat{\varepsilon}_2,\ \ldots\} \quad (16)$$

can be defined by scaling the normalised sampling directions $\hat{\varepsilon}_k$ with the positive and negative reconstructed corresponding coefficients $\zeta_k(\mathbf{x}_i) = |\eta_k(\mathbf{x}_i)|^{1/2}$. Again, this is due to the symmetry of the scattering tensors. Note that the mean of this set trivially equals $\mathbf{0}$. Let $C_i \in \mathbb{R}^{3\times3}$ denote the *covariance* matrix of S_i. Then, its *eigen-decomposition*

$$V_i\cdot\Lambda_i = C_i\cdot V_i \quad (17)$$

yields a diagonal matrix $\Lambda_i \in \mathbb{R}^{3\times3}$ containing the three *eigenvalues* $\lambda_{i,1}$, $\lambda_{i,2}$, $\lambda_{i,3}$ of the set's covariance matrix C_i, and an orthogonal matrix $V_i \in O(3) \subset \mathbb{R}^{3\times3}$ containing the corresponding mutually orthonormal *eigenvectors* $\mathbf{v}_{i,1}$, $\mathbf{v}_{i,2}$, $\mathbf{v}_{i,3}$ as column vectors.

Let $\overline{\eta}_i = \sum_k |\eta_k(\mathbf{x}_i)|/K$ denote the average squared scattering magnitude at voxel $\mathbf{x}_i$, and $\overline{\lambda}_i = (|\lambda_{i,1}| + |\lambda_{i,2}| + |\lambda_{i,3}|)/3$ the corresponding average eigenvalue magnitude. Define a *size correction* factor $\sigma_i := \overline{\eta}_i/\overline{\lambda}_i$ for scaling the statistically defined ellipsoid to the point set S. Then, the scattering ellipsoid at voxel $\mathbf{x}_i$ is defined by half-axis lengths $r_{i,1} = [\sigma_i\lambda_{i,1}]^{1/2}$, $r_{i,2} = [\sigma_i\lambda_{i,2}]^{1/2}$, $r_{i,3} = [\sigma_i\lambda_{i,3}]^{1/2}$ with respect to the orthonormal basis formed by vectors $\mathbf{v}_{i,1}$, $\mathbf{v}_{i,2}$, $\mathbf{v}_{i,3}$.

2.3. Constraint enforcement

As it will turn out in the experiments, unconstrained reconstruction of the scattering coefficients yields useful but noisy results. Instead of just fitting ellipsoids retrospectively, it makes sense to enforce ellipsoidal shapes during reconstruction. Therefore, we propose to post-process the scattering coefficients at the end of *every* iteration as indicated in the last comment of alg. 1, thus forcing them to evolve in the vicinity of the manifold of ellipsoid-shaped tensors.

2.3.1. Hard ellipsoid constraint

An obvious approach to reach this aim is to fit ellipsoids after every iteration, and to project the reconstructed squared coefficients $\boldsymbol{\eta}_k$ onto them. We refer to this method as *hard ellipsoid constraint* as the coefficients will really be forced into ellipsoidal shapes.

In detail, after the computation of all squared coefficients $\boldsymbol{\eta}_k^{(q)}$ in iteration q, we visit every voxel $\mathbf{x}_i$ and fit an ellipsoid as described in section 2.2, thus obtaining a coordinate frame $V_i = [\mathbf{v}_{i,1}, \mathbf{v}_{i,2}, \mathbf{v}_{i,3}] \in \mathbb{R}^{3\times3}$ and corresponding half-axes $r_{i,1}, r_{i,2}, r_{i,3} \in \mathbb{R}$. Then, the reconstructed coefficients $\boldsymbol{\eta}_k^{(q)}$ are replaced with the projections of the respective normalised sampling directions $\hat{\varepsilon}_k$ onto this ellipsoid.

In detail, to do this for every sampling direction k and voxel $\mathbf{x}_i$, we first rotate the (by design normalised) vector into the ellipsoid's coordinate frame, thus obtaining a unit vector $[x_{i,k},\ y_{i,k},\ z_{i,k}]^T = V_i^T\hat{\varepsilon}_k$. By definition, the scaled vector $\sigma_{i,k}\cdot[x_{i,k},\ y_{i,k},\ z_{i,k}]^T$, $\sigma_{i,k} \in \mathbb{R}$, resides on the surface of the ellipsoid if

$$\sigma_{i,k}{}^2\left[\left(\frac{x_{i,k}}{r_{i,1}}\right)^2 + \left(\frac{y_{i,k}}{r_{i,2}}\right)^2 + \left(\frac{z_{i,k}}{r_{i,3}}\right)^2\right] = 1. \quad (18)$$

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

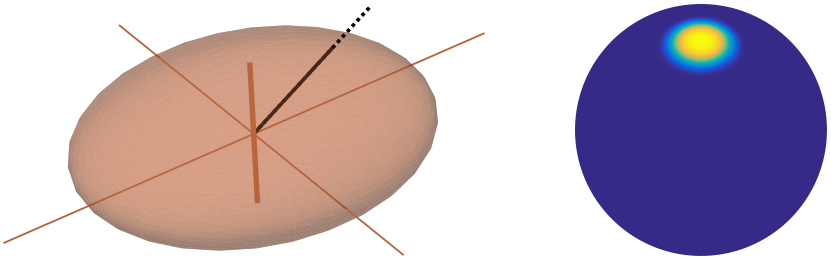

(a) Hard constraint. (b) Soft constraint.

Fig. 7. Constraint enforcement. In both cases, the reconstructed coefficients are forced to be close to the manifold of valid ellipsoids while iterating. The hard constraint (a) projects back onto ellipsoids directly, where a – in this toy example exaggerated – scattering coefficient (dotted black line) is shortened appropriately (solid black line). The soft constraint smoothes the coefficients with respect to the other sampling directions of the *same* voxel respecting the angular relation, thus favouring ellipsoids in a relaxed sense. This is done using a Gaussian smoothing kernel based on the scalar product as shown in (b), ranging between 0 (blue) and 1 (yellow), before normalisation.

This Eq. can simply be solved for ${\sigma_{i,k}}^2$, and considering that squared coefficients are reconstructed, the correctly projected coefficients are

$$\boldsymbol{\eta}_{i,k}^{(q)} = {\sigma_{i,k}}^2 \tag{19}$$

for all voxels $\mathbf{x}_i$ and directions k. See Fig. 7(a) for an illustrative sketch.

2.3.2. Soft ellipsoid constraint

Alternatively, we propose to use a softer variant favouring ellipsoids but allowing some more freedom. Essentially, this is done by smoothing the K direction coefficients per voxel, over the ellipsoid. We refer to this approach as *soft ellipsoid constraint*. Note that we still only consider individual voxels and their K coefficients; we do not take a neighbourhood into account as one would do for smoothness constraints in regularised X-ray attenuation reconstruction.

In detail, we first compute the pair-wise absolute scalar products between all K normalised sampling directions, and obtain a matrix $S \in [0,1]^{K\times K}$, $S_{k,l} = |\langle \hat{\boldsymbol{\varepsilon}}_k, \hat{\boldsymbol{\varepsilon}}_l \rangle|$. In order to increase the influence of more similar sampling directions, we additionally apply a Gaussian and obtain a matrix $G \in \mathbb{R}^{K\times K}$, $G_{k,l} = \exp[-(S_{k,l}-1)^2/2\mu]$ with some selectable mean μ. The larger μ is chosen, the higher will be the influence of the smoothing.

To process coefficient k, a matching convolution kernel can be defined by normalising the kth row of G to a sum of 1, denoted as $\mathbf{g}_k^T$, $\sum_l [\mathbf{g}_k^T]_l = 1$. Collecting all K coefficients into a vector, the regularised kth value for voxel $\mathbf{x}_i$ computes as

$$\boldsymbol{\eta}_{i,k}^{(q)} = \langle \mathbf{g}_k^T, [\boldsymbol{\eta}_{i,1}^{(q)}, \ldots, \boldsymbol{\eta}_{i,K}^{(q)}] \rangle. \tag{20}$$

In other words, this approach can be considered as angle-dependent Gaussian smoothing of the squared coefficients $\eta_k(\mathbf{x}_i)$ over all directions k for every individual voxel $\mathbf{x}_i$. This constraint is illustrated in Fig. 7(b).

2.4. Visualisation

A last, vital point for X-ray tensor tomography is a proper visualisation of the results. This is of particular importance as the voxel-wise tensors can not be accurately shown using rendering

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

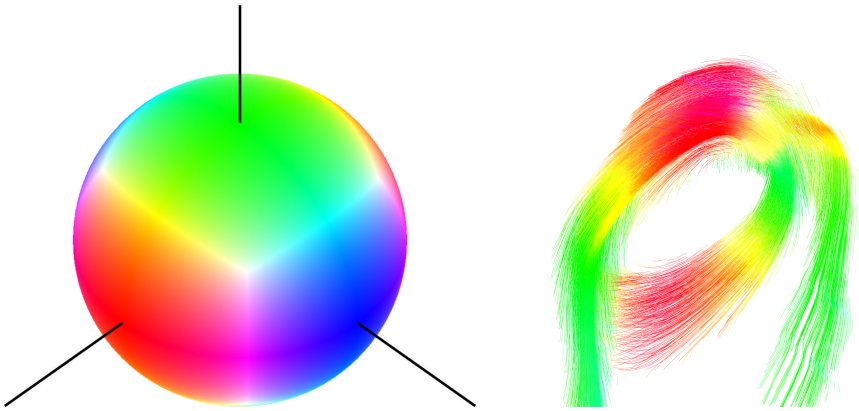

(a) (b)

Fig. 8. Colour coding of streamline visualisation. We use colours sampled from a sphere (a) to give visual cues about the orientation of the 'fibres'. Note that the colour ball is symmetric with respect to the *x*-*y*-, *x*-*z*-, and *y*-*z*-planes. In a sample picture of the carbon knot (b), looking down from a slightly elevated perspective, vertically (green), horizontally (red) and obliquely (orange and yellow) oriented parts can be easily distinguished.

techniques intended for scalar-valued data. Apart from plotting the ellipsoids themselves, a prominent option is to visualise derived shapes.

Interested in fibre- and tube-like structures, we therefore propose a *streamline* visualisation based on the classical Runge-Kutta method [24–26] (RK4). This metaphor is well-known in the field of diffusion-tensor magnetic resonance imaging (DTI) [10, 27], and commonly referred to as *tractography*. As sketched in Fig. 3, plate-like ellipsoids hint at fibres in the direction of the smallest half-axes. We employ RK4 to trace streamlines along these smallest half-axes, adapted to account for the symmetric, 'bidirectional' nature of ellipsoid axes (compared to a vector field). This is done by inverting directions queried from the raw ellipsoid axes whenever the scalar product with the direction of an incoming fibre is negative, i.e. new directions are chosen to continue the streamline in forward direction. Furthermore, we employ colour coding to give cues about the orientation of these fibres. An example is shown in Fig. 8.

Note that the streamlines do not necessarily correspond to real fibres, and there is no intention for this to be the case. The streamlines are to be interpreted as scattering visualisation, however, using a model which is likely similar to the invisible structure.

3. Experiments

We have implemented the proposed reconstruction algorithm within *CampRecon* [28], a C++ framework for solving linear inverse problems. The discrete integration along lines $A\mathbf{x}$ and its adjoint $A^T\mathbf{b}$ are approximated roughly via intersections between ray and pixels as proposed by Siddon [29], using the GPU-based X-ray projector developed by Fehringer *et al.* [15] in OpenCL. All experiments have been run on a computer with dual Intel Xeon E5-2687W processors and a Nvidia Tesla K20 accelerator.

In general, we have computed reconstructions of three datasets, the carbon *'knot'*, a knotted bunch of carbon fibres embedded in hot glue, the tree *'branch'*, a short piece of raw wood, and a *'tooth'*. Photographies of the samples are given in Fig. 9. X-ray images of all samples have been acquired in our experimental setup, see Fig. 10. We use 732 projections of 321×321 pixels for the *knot*, 551 projections of 301×301 pixels for the *branch*, and 902 projections of

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

(a) Knot. (b) Branch. (c) Tooth.

Fig. 9. Photographies of the samples. Knot (a), branch (b), and tooth (c).

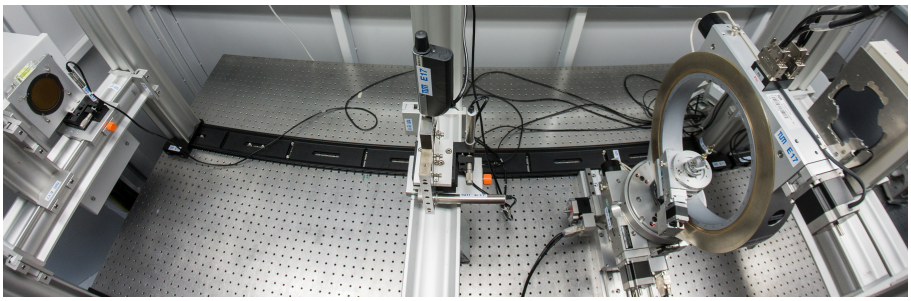

Fig. 10. Wide-angle view of the actual setup. From left to right: X-ray source (T) with source grating (G_0) directly in front of it, at the center the phase grating (G_1), then the Euler cradle with sample (S), and adjacent to it the analyser grating (G_2, in the cross-shaped holder); behind the latter the X-ray detector (D, hidden).

701×701 pixels for the tooth, with trajectories resembling the one in Fig. 6. The projections in Fig. 2 have been taken from the *tooth* dataset.

All samples were measured at an acceleration voltage of 60 kVp. Eight phase steps were recorded per projection with a flat panel detector (Varian) with pixel pitch of 127 μm. The exposure time was 1 s per phase step. A $\pi/2$ phase grating for the design energy of 45 kV composed of 8 μm high Ni lines and period of 5 μm was used. The two other gratings were absorption gratings with 170 μm high Au lines and a period of 10 μm. The interferometer was symmetric with both inter-grating distances being 92.7 cm.

Each of the datasets was reconstructed without constraint enforcement, with hard, and with soft ellipsoid enforcement, each time with $K = 13$ sampling directions. In all cases, we executed 100 iterations of alg. 1 and do not employ any other stopping criterion. We use a single iteration of CG as function **approximate**. Note that this corresponds to a single Landweber step with respect to the sub-problem. For the soft constraint, we generally use an experimentally established smoothing mean of $\mu = 0.1$, unless stated otherwise. Reconstruction times are about 1.75 hours for the *branch* (251^3 voxels), 2 hours for the *knot* (201^3 voxels), and 15.5 hours for the *tooth* ($301 \times 501 \times 291$ voxels). Relating the number of voxels with the input projections, it is clear that the linear system $\mathbf{m} = H\mathbf{s}$ as defined in Eq. (10) is underdetermined in all our experiments. The ratios of measurements to unknowns are 24.3 % for the *branch*, 70.4 % for the *knot*, and 77.7 % for the *tooth*, respectively.

Furthermore, we have obtained a reconstruction of the *knot* produced with the original SART-variant proposed by Malecki *et al.*, for comparison.

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

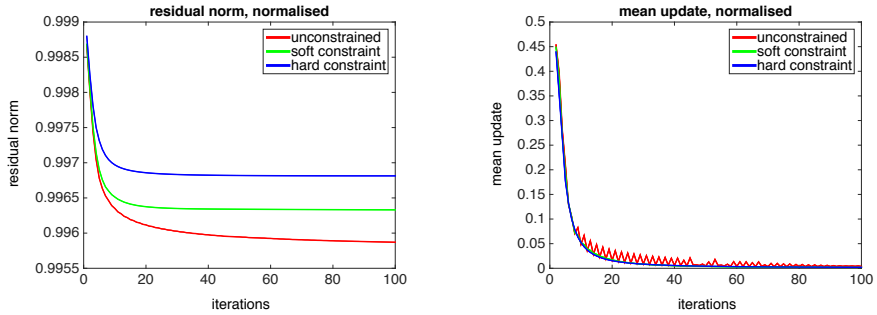

Fig. 11. Behaviour of the algorithm for the *knot*. The plots show the normalised residual norm $r^{(q)}$ as defined in Eq. (21) over iterations q (left) and the normalised mean update $\Delta^{(q)}$ as defined in Eq. (22) over iterations q (right). The unconstrained version yields smaller residuals, but the updates are noisy.

4. Results

4.1. *Numerical behaviour*

In order to check the behaviour of the three algorithm variants (unconstrained, soft, and hard ellipsoid constraint), we have computed the normalised residual norms

$$r^{(q)} := \|\mathbf{m} - \sum_k D_k \cdot A \cdot \boldsymbol{\eta}_k^{(q)}\|_2 \;/\; \|\mathbf{m}\|_2 \tag{21}$$

for iterations $q \in \{1, \ldots, 100\}$, and also the normalised mean updates

$$\Delta^{(q)} := \operatorname{mean}_k \|\boldsymbol{\eta}_k^{(q)} - \boldsymbol{\eta}_k^{(q-1)}\|_2 \;/\; \|\boldsymbol{\eta}_k^{(q)}\|_2. \tag{22}$$

Sample plots of these sequences for the *knot* are given in Fig. 11. There, the curves for both measures flatten out with increasing number of iterations. The unconstrained variant tends to reach a smaller residual norm, while the update norms appear to oscillate. Vice versa, the constrained variants show larger residual norms, but smooth updates. For other datasets, we obtain similar curves.

4.2. *Comparison with state of art*

Next, in order to compare our method with the original SART-based method of Malecki *et al.*, we compare an unconstrained reconstruction of the *knot* with a reference result, see Fig. 12. As Malecki's results can not be considered as ground truth, we restrict ourselves to a visual comparison.

4.3. *Knot*

Figure 13 shows volume renderings of raw scattering coefficients relating to a single sampling direction, but for the three different constraint enforcement schemes. Most importantly, the unconstrained version shows considerable streak artefacts, and the constrained versions do not.

Figure 14 and Media 1 show streamline visualisations of the *knot*. As can be seen, enforcing the two constraints yield *visually* smoother fibres that are more densely packed. Again, note that these fibres visualise the scattering ellipsoids, and are not to be considered reconstructions of the raw carbon fibres.

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

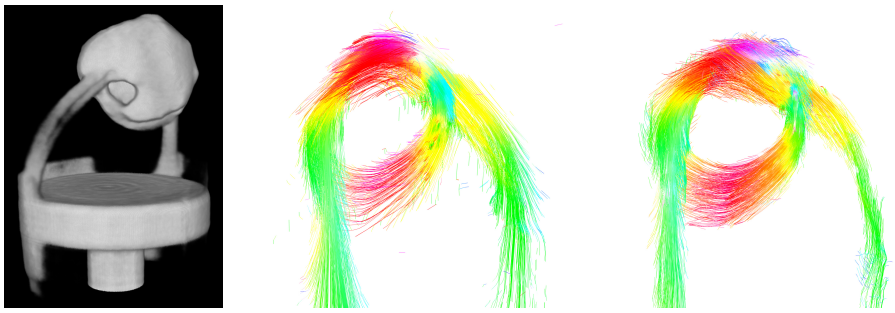

(a) Attenuation. (b) Malecki's reconstruction. (c) Our reconstruction.

Fig. 12. Comparison of attenuation reconstruction (a), Malecki's tensor reconstruction (b), and our (unconstrained) tensor reconstruction (c). The first image is included for reference, to show that scattering data is of a considerably different nature than usual attenuation reconstructions. The two scattering reconstructions are largely equivalent, considering the different algorithms with different parameters.

(a) Unconstrained. (b) Soft constraint. ($\mu = 0.10$) (c) Hard constraint.

Fig. 13. Volume renderings of the *knot*'s scattering coefficients for sampling direction $\hat{\varepsilon}_k = [0,0,1]^T$. The unconstrained reconstruction (a) shows strong streak artefacts, the two constrained versions (b) and (c) are much clearer.

4.4. Branch

Figure 15 and Media 2 show streamline visualisations of the *branch*. For this dataset, enforcing ellipsoid constraints has even more prominent effects. Considering the structure of wood, scattering will be caused primarily by the tiny vessels transporting water towards the leaves. That is, streamlines are supposed to run mainly in parallel to the axes of the branches. In the unconstrained case, reconstruction of scattering along the main branch fails almost entirely, but the constrained versions are able to recover more reasonable scattering streamlines there.

We also modify the smoothing parameter μ for the soft constraint. In particular, a 'combing' effect can be observed for stronger constraint enforcement where streamlines are nicely aligned to each other, but 'combed' away from their expected location, see the side branches in Fig. 15(d). On the other side, a smaller value of μ will increase the amount of noise but cause the reconstruction to stay closer to the data. Especially the side branches appear to benefit from a smaller parameter, as the 'combing' effect is considerably reduced, and the main branch still reconstructs properly.

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

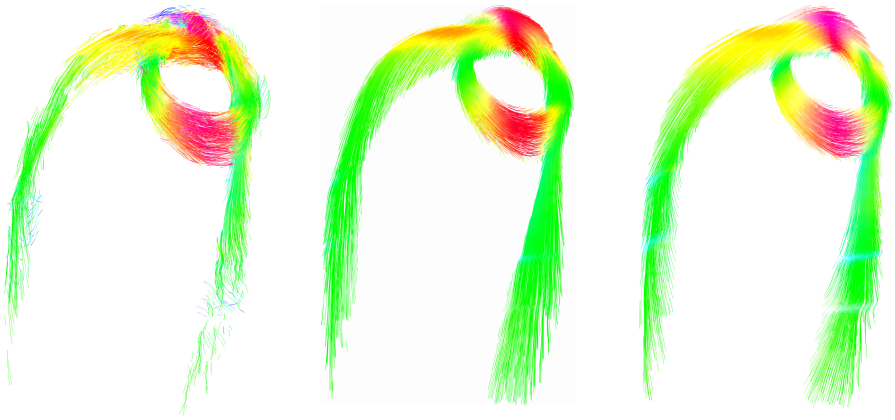

(a) Unconstrained. (b) Soft const. ($\mu = 0.10$) (c) Hard constraint.

Fig. 14. Streamline visualisation of the *knot*'s scattering ellipsoids. (Also see Media 1.) The streamlines are supposed to follow the directions of the carbon fibres, see Fig. 9(a), but are not intended to accurately reconstruct individual fibres. The unconstrained version (a) shows noise, the two constrained versions (b) and (c) are visually considerably smoother. The 'waves' in the lower right appear to be additional scattering caused by the sample holder.

4.5. Tooth

The reconstruction results of the *tooth* are given in Fig. 16 and in Media 3. This sample is particularly intriguing: Teeth generally consist of a hard crown of highly mineralised enamel, and a root covered by cementum. Below these layers, the core of the tooth consists of dentine and, embedded within, the pulp chamber containing living tissue, blood vessels and nerves. Dentine is a fibrous material and less mineralised than enamel. Tiny dentinal tubules (about 2 μm in diameter) are passing through it in radial direction, between pulp chamber and the surface (but not through the enamel). These tubules are generally not visible in X-ray attenuation reconstructions of usual resolution, due to their small size.

Despite this, however, the scattering caused by them can be measured and reconstructed. Note that the results in Fig. 16 have been masked using the reconstruction of the attenuation signal. This is necessary as the contents of the pulp chamber, relicts of the tissue, cause considerable isotropic scattering themselves – clearly visible as distinct black region in the scattering projection in Fig. 2(c).

The influence of constraint enforcement is similar to the one observed for the other samples. However, at the lower parts of the roots, the soft constraint also yields streamlines caused by pulp chamber scattering.

5. Discussion

As shown in the first experiments, the proposed reconstruction algorithm apparently shows reasonable behaviour in our experiments. Plots similar to the ones shown in Fig. 11 can be produced for all experiments that we have conducted so far. Note that these plots only provide hints that the iteration does not become unstable, they neither indicate image quality nor do they prove convergence.

Comparing with the state of art, the unconstrained variant of our algorithm yields a result quite similar to Malecki's original SART-modification, considering the different nature and

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

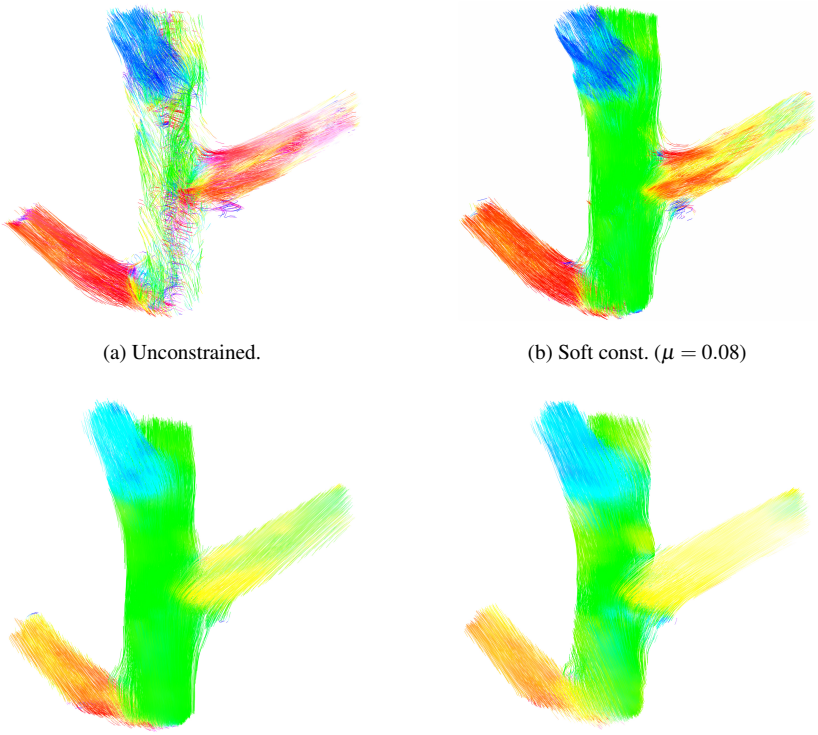

(a) Unconstrained.

(b) Soft const. ($\mu = 0.08$)

(c) Soft const. ($\mu = 0.10$)

(d) Hard constraint.

Fig. 15. Streamline visualisation of the *branch*. (Also see Media 2.) Scattering is supposed to be caused by the tiny tubular vessels embedded in wood that transport water towards the leaves. The streamlines are thus supposed to run mainly in parallel to the individual branches. The unconstrained version (a) fails to recover useful scattering along the main branch. The lightly constrained versions (b) and (c) show more reasonable scattering there. The hard constraint overshoots and produces wavy patterns along the main branch, and introduces a 'combing' effect for the other branches. The latter can already be seen in (c).

parameters of the two methods.

The constrained variants of our method have been shown to produce considerably less artefacts and visually smoother, non-oscillating updates. As most notable example, the main branch of the wood sample inherently requires some kind of constraint enforcement, to push the reconstruction towards a reasonable result. On the other side, exaggerated constraint enforcement may lead to new artefacts, particularly unexpected wavy patterns or 'combed' directions, as can also be seen for the wood sample, for instance. It appears that constraint enforcement is necessary, but its influence must be limited. From our point of view, soft constraint enforcement with a conservatively chosen parameter μ is the most promising approach.

In general, a remarkable feature of X-ray tensor tomography is that scattering effects caused by sub-voxel-sized structures can be recovered, thus giving insight into the microstructure. In our case, this holds for both, the vessels in the branch, and the dentinal tubules in the tooth. Of course, important future work will be an exact correlation of X-ray scattering results with suitable high-resolution pictures of the samples, through μCT or microscopy, for instance.

An important consequence of our work is that the assumption of pure ellipsoids appears not

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

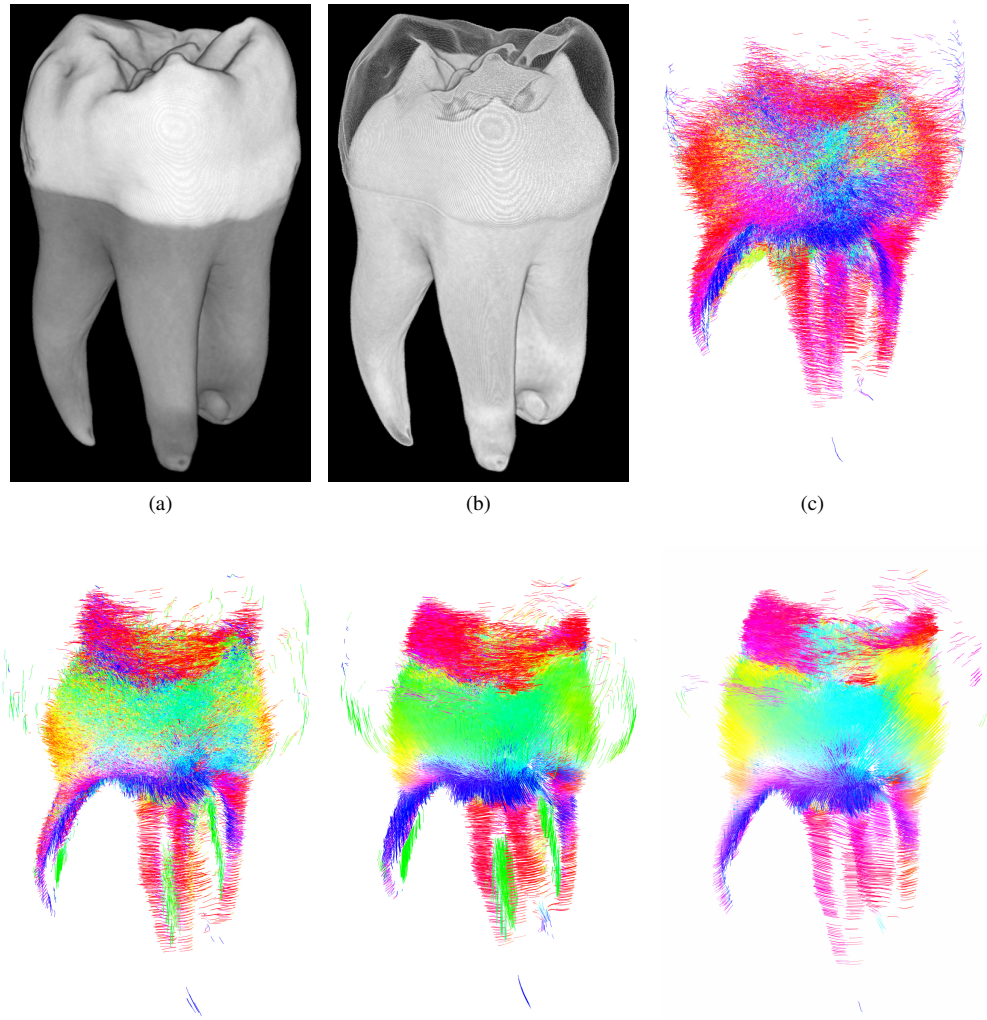

Fig. 16. Reconstruction results of the *tooth*. (Also see Media 3.) Volume rendering of X-ray *attenuation* (a) showing enamel (white) and dentine (gray). Equivalent *attenuation* volume rendering showing dentine only (b). Scattering streamlines obtained without constraint (c), with soft constraint with $\mu = 0.08$ (d) and $\mu = 0.10$ (e), and with hard constraint (f). Scattering is caused by dentinal tubules, tiny structures in radial direction, as indicated by the streamlines. The pulp chamber was obtained by masking.

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA

to be entirely correct, considering the deterioration with stronger constraint enforcement. Based on sampling directions, our reconstruction is agnostic of that assumption, but future work will need to look into more appropriate scattering descriptions. Furthermore, even after flattening out, the residual norm is still rather large. Usually, this would hint at a mathematical problem. Considering the visual quality of the reconstructions and the curve of the update norm, however, this may also be caused by Eq. (1) only partially describing the dark-field signal. Both (highly related) questions will need to be subject to future investigation from a physics perspective.

Considering the proposed algorithm, many interesting mathematical questions are still open and will need to be tackled in future work. As pointed out by a reviewer, it should be possible to position the proposed algorithm in the context of other, well-established solution algorithms, such as column-based block iterative algorithms [17], for the huge system defined in Eq. (10). Particularly interesting will be an investigation of two points: First, our algorithm computes the weighted forward projections WFP_k based on the previous iterates, and thus exclusively uses the previous state to update the solution. This decision has been made from an engineering prespective, and to avoid bias towards the scattering coefficients that are processed earlier. From a mathematical point of view, however, it may be beneficial to only use the previous state where necessary and the already known current state where possible. The second question is the dependence on the exact choice of function **approximate**. In this context, it will also be useful to attempt to solve the original linear system directly. More advanced mathematical topics include an investigation of the convergence, and the influence of the constraint enforcement schemes onto it. Finally, another important open question is inter-voxel regularisation such as total variation minimisation.

6. Conclusion

We have presented a new processing chain for X-ray tensor tomography. In particular, we have detailed a generic formulation for solving the inverse problem using arbitrary linear solvers, firmly based on the forward model and avoiding the tight coupling to SART, along with constraint enforcement schemes improving the reconstructions considerably. Furthermore, we have presented a way for quick recovery of the ellipsoids, and we have introduced streamline visualisation to XTT.

A lot of work still needs to go into XTT. At this point, the mechanical design of the setup requires small samples, and the radiation exposure is considerable. The latter is due to both the large number of acquisitions (more perspectives, more acquisitions per perspective), and the additional absorption caused by the interferometer grating between sample and detector, thus requiring higher radiation at the specimen in order to obtain good detector readings. Further, more theoretical open questions have already been outlined above. Nevertheless, application in ex-vivo imaging or material testing is well within reach today, and improvements are subject of active research.

Acknowledgments

This work was partially funded by the DFG Cluster of Excellence Munich–Centre for Advanced Photonics (MAP) and carried out with the support of the Karlsruhe Nano Micro Facility (KNMF), a Helmholtz Research Infrastructure at Karlsruhe Institute of Technology (KIT). J. V. thanks Andrea Porsche for skilfully recovering the *tooth* sample. F. S. is funded by the TUM Graduate School.

We are indebted to the anonymous reviewer who provided an exceptional report including many valuable comments about mathematical and numerical aspects of the proposed reconstruction approach.

#236868 - $15.00 USD Received 26 Mar 2015; revised 15 May 2015; accepted 25 May 2015; published 1 Jun 2015
© 2015 OSA