

Mumott – a Python package for the analysis of multi-modal tensor tomography data

Leonard C. Nielsen,¹ Mads Carlsen,² Sici Wang,^{2,3} Arthur Baroni,² Torne Tänzer,^{2,3} Marianne Liebi,^{1,2,3} and Paul Erhart¹

¹*Department of Physics, Chalmers University of Technology, Gothenburg, Sweden*

²*Photon Science Division, Paul Scherrer Institute, Villigen, Switzerland*

³*Institute of Materials, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland*

(Dated: April 24, 2025)

Small and wide angle x-ray scattering tensor tomography are powerful methods for studying anisotropic nanostructures in a volume-resolved manner, and are becoming increasingly available to users of synchrotron facilities. The analysis of such experiments requires, however, advanced procedures and algorithms, which creates a barrier for the wider adoption of these techniques. Here, in response to this challenge, we introduce the MUMOTT package. It is written in Python with computationally demanding tasks handled via just-in-time compilation using both CPU and GPU resources. The package is being developed with a focus on usability and extensibility, while achieving a high computational efficiency. Following a short introduction to the common workflow, we review key features, outline the underlying object-oriented framework, and demonstrate the computational performance. By developing the MUMOTT package and making it generally available, we hope to lower the threshold for the adoption of tensor tomography and to make these techniques accessible to a larger research community.

I. INTRODUCTION

The properties of numerous materials depend on the hierarchical organization of their basic building blocks ranging from the nanometer to the micrometer scale. Examples include plant materials assembled from cellulose and lignin [1], bone constructed of assemblies of mineralized collagen fibrils [2], or polymeric materials such as in the structure of semi-crystalline polymers [3–5], and liquid-crystalline polymers composed of rigid macromolecules [6]. The study of structure-property relationships of hierarchical materials for applications in biology, the biomedical field, or polymer engineering relies on accurate structural characterization from a wide range of techniques. X-ray techniques are of particular interest with respect to providing volume-resolved nanostructural information in macroscopic samples. Due to the high penetration depth and non-destructive nature of the techniques, methods such as x-ray absorption and phase contrast computed tomography (CT) have played an important role in providing high resolution densimetric measurements of 3D samples [7, 8]. In addition to the densimetric fields, the arrangement of nanostructural elements, in particular their direction and degree of alignment, is important for many mechanical and functional properties. This brings an additional challenge in methodologically bridging between the length scales of the nanostructural building blocks and the macroscopic specimen. In addition to methods resolving the nanostructure spatially with high-resolution imaging techniques, the orientation of the nanostructure can also be probed by polarization, scattering, diffraction, or magnetic relaxation methods in a volume averaged way [9]. The first technique for probing the orientation of structures without directly resolving them was polarized light microscopy. It was followed by other techniques that also make use of the polarization of light such as polarized Ra-

man or Fourier transform infrared spectroscopy and polarized second-harmonic-generation imaging. X-ray and neutron diffraction approaches can be used, such as directional dark-field (DDF) imaging[10, 11] which probes the orientation at the micrometer-scale through the integrated scattering signal, or scanning small and wide angle scattering that allow probing of nanoscale structures. Small-angle x-ray scattering (SAXS) probes the spatial variation of the electron density, providing information on microstructural elements with characteristic length scales in the range of tens to hundreds of nanometers, relating to the structural organization and orientation of the materials at the corresponding length scales, while x-ray diffraction (XRD) (in this paper called wide-angle x-ray scattering (WAXS)) probes atomic distances and crystal lattices. Whereas DDF are a family of full-field imaging methods, SAXS and WAXS can be used as scanning imaging techniques in which the sample is raster-scanned with a focused x-ray beam providing an image of the sample consisting of a 2D diffraction pattern in each pixel. Tomographic reconstruction of such measurements using isotropically scattering samples is known as XRD-CT and is frequently used both in the SAXS [12, 13] and WAXS [14–16] regimes at synchrotron x-ray sources.

To access the orientation information of the underlying ultrastructure within a 3D specimen, tomographic methods can be extended from the reconstruction of a scalar fields to tensor fields describing the directionality of the signal, which is in general called tensor tomography (TT). The most established technique in this category is diffusion magnetic resonance imaging (MRI), also called diffusion tensor imaging (DTI), which is widely used to study the 3D arrangement and orientation of neurons. In the case of X-rays, TT has been demonstrated for DDF [17, 18], SAXS [19–23] and WAXS [24]. Other related tomography approaches which can be considered as TT

include probing magnetic field directions with circular polarized x-rays [25] or polarized neutrons [26].

The acquisition and analysis of TT data is a non-trivial undertaking, creating a barrier for the wider adoption of these powerful techniques. In response to this challenge, specifically with regard to the analysis of such data, we here introduce the software package MUMOTT for the reconstruction of TT data. While the current implementation supports the cases of SAXS and WAXS, the framework offers the possibility to include other modalities in the future. In the following we first provide a brief overview of the methodology (Sect. II) before describing the structure and functionality of MUMOTT (Sect. III). Finally, we give a short outlook concerning potential future additions and developments (Sect. IV).

II. METHODOLOGY

SAXS/WAXS-TT is conceptually similar to XRD-CT in that the sample is raster-scanned through a focused beam to produce a number of 2D projections, varying the sample orientation between each projection. Unlike XRD-CT it works with azimuthally re-grouped detector images where the intensity of the scattered x-rays in a number of azimuthal bins is recorded rather than a single azimuthally integrated intensity. The width of the azimuthal bins depends on the desired angular resolution of the reconstruction. The azimuthal re-grouping can be done with a number of freely available software tools such as PYFAI [27] and MATFRAIA [28]. The experimental data is then a five-dimensional dataset consisting of the tomographic rotation, the two directions of the raster scan grid, the scattering angle 2θ , and the azimuthal angle φ . MUMOTT deals with the reconstruction of such a five-dimensional dataset into a six-dimensional reconstruction, consisting of a three-dimensional voxel map containing a three-dimensional reciprocal space map (RSM) in each voxel.

We assume that the data has already been corrected for various experimental errors pertaining to solid angle, geometric distortions, and polarization. To account for the effect of absorption by the sample, the collected data can be normalized by the transmitted intensity as is common practice in XRD-CT. Especially at small scattering angles, this makes it possible to carry out reconstructions even with low sample transmission coefficients ($\approx 1\%$ has been demonstrated) assuming sufficient incident flux. The measurement of the transmitted beam intensity can be done using either a semi-transparent beam stop, a diode mounted on the beam stop, or a fluorescence measurement [29]. Alternatively, synthetic transmission data can be calculated based on an absorption CT reconstruction [30].

The experiment is described in a coordinate system defined by the voxel grid of the sample and the three orthogonal basis vectors $\hat{x}$, $\hat{y}$, and $\hat{z}$. Typically these vectors are chosen to conform to the convention of the beam-

TABLE I. Unit vectors defining the experimental geometry and their values in the standard geometry used in previous publications such as Ref. 21.

Symbol	Standard	Field name
$\hat{p}$	$+\hat{z}$	p_direction_0
$\hat{j}$	$+\hat{y}$	j_direction_0
$\hat{k}$	$+\hat{x}$	k_direction_0
$\hat{q}_0$	$+\hat{x}$	detector_direction_origin
$\hat{q}_{90}$	$+\hat{y}$	detector_direction_positive_90
$\hat{\alpha}$	$+\hat{y}$	inner_axis
$\hat{\beta}$	$+\hat{x}$	outer_axis

line, where the experiments were performed, such that the sample-fixed coordinates correspond to the laboratory coordinates when the goniometer angles are zeroed. The geometry of the instrument is defined by specifying a number of unit vectors in these laboratory coordinates. These vectors include the beam direction $\hat{p}$ (also called the projection vector), the two orthogonal directions of the raster-scan $\hat{j}$ and $\hat{k}$, and two vectors describing the origin and the positive direction of the azimuthal integration $\hat{q}_0$ and $\hat{q}_{90}$ defined by the equation

$$\begin{aligned} \hat{q}(\varphi) &= \cos(2\theta/2)(\cos \varphi \hat{q}_0 + \sin \varphi \hat{q}_{90}) - \sin(2\theta/2)\hat{p} \\ &\approx \cos \varphi \hat{q}_0 + \sin \varphi \hat{q}_{90}. \end{aligned} \quad (1)$$

This equation gives the normalized scattering vector $\hat{q}(\varphi)$ probed by each detector segment as a function of the scattering angle 2θ and the detector azimuth angle φ . The second line gives a useful approximation valid for small scattering angles. The sample can be rotated by a goniometer and the rotation of the sample goniometer at a given setting labeled by s results in a rotation matrix $\mathbf{R}_s$. Typically, the goniometer is constructed by two orthogonal rotation stages, an inner “rotation” and outer “tilt” stage. The full rotation is then defined by a pair of rotation angles α and β with corresponding rotation axes $\hat{\alpha}$ and $\hat{\beta}$, such that $\mathbf{R}_s = \mathbf{R}_{\hat{\beta}}(\beta)\mathbf{R}_{\hat{\alpha}}(\alpha)$. While all these vectors may be chosen freely in MUMOTT (under the restriction that certain vectors are orthogonal to certain other vectors), we work in a standard geometry in this paper given by the choices tabulated in Table I and visualized in Fig. 1.

The scattering from a given voxel (x, y, z) is proportional to a characteristic function $f_{xyz}^{3D}(\mathbf{q})$ called the *3D reciprocal space map (RSM)*. In the context of SAXS-TT, the RSM is the Fourier transform of the auto-correlation function of the electron density taken over a small volume. For the purpose of reconstruction, we consider one “shell” of reciprocal space at a time and the 3D RSM is built up by reconstructing and stacking successive 2D shells (sketched in Fig. 2d). For one such shell we consider the function $f_{xyz}^{2D}(\hat{q})$, which depends only on the *direction* of the scattering vector. This function is mod-




FIG. 1. Illustration of vectors defining the experimental geometry in the laboratory coordinates (i.e., at $\alpha = \beta = 0$).

eled by a sum of basis functions,

$$f_{xyz}^{2D}(\hat{q}) = \sum_i c_{xyzi} B_i(\hat{q}), \quad (2)$$

where $B_i(\hat{q})$ are the basis functions (see Sect. III C below) and c_{xyzi} are the unknown expansion coefficients that we want to find.

$f_{xyz}^{2D}(\hat{q})$ is described in sample-fixed coordinates such that at a given rotation of the sample given by $\mathbf{R}_s$ the detector segment at the angle φ measures the component $f_{xyz}^{2D}(\mathbf{R}_s^\top \hat{q}(\varphi))$, where the superscript $\{\}^\top$ denotes the matrix transpose. In the normal setting, the detector is split into a number of evenly spaced segments indexed by c covering either the full 360° for WAXS or 180° for SAXS. Each detector segment (Fig. 2a) is defined by a start angle $\varphi_{c,\text{start}}$ and an end angle $\varphi_{c,\text{end}}$. As such, the detector segment probes the average of the scattering function within this interval given by the integral

$$I_c \propto \frac{1}{\varphi_{c,\text{end}} - \varphi_{c,\text{start}}} \int_{\varphi_{c,\text{start}}}^{\varphi_{c,\text{end}}} f_{xyz}^{2D}(\mathbf{R}_s^\top \hat{q}(\varphi)) d\varphi$$

by inserting Eq. (2) into the above, we define the constants (Fig. 2a,b)

$$B_{sc,i} = \frac{1}{\varphi_{c,\text{end}} - \varphi_{c,\text{start}}} \int_{\varphi_{c,\text{start}}}^{\varphi_{c,\text{end}}} B_i(\mathbf{R}_s^\top \hat{q}(\varphi)) d\varphi, \quad (3)$$

which describe how much each basis function scatters in the direction measured by a given detector segment, illustrated in Fig. 2a. This is an integral over a single scalar variable, which can be numerically evaluated by standard methods of quadrature.

To complete the forward model, we have to sum up the intensity contributions from all voxels in the path of the incident beam. At a given position of the raster scan and rotation of the goniometer, only the voxels that are illuminated by the beam contribute to the measured




FIG. 2. (a) Layout of vectors and angles on the detector. A single detector segment is marked with a thick black line. (b) Integrated basis function values $B_{sc,i}$ plotted in a stereographic projection. The black solid arc corresponds to the single detector segment marked in (a). (c) Computed probing of each voxel by bilinear interpolation. (d) Splitting of a 3D-RSM into a stack of 2D-RSMs at fixed q -lengths.

scattering. A given voxel is indexed by the three integers x , y and z and at a given setting of the sample goniometer the position of the voxel is

$$\mathbf{r}_{xyz} = a\mathbf{R}_s \begin{bmatrix} x & y & z \end{bmatrix}^\top - b(j - \delta j)\hat{j} - b(k - \delta k)\hat{k}, \quad (4)$$

where j and k are integer indices of the raster scan, a is the step size of the cubic voxel grid, b is the step size of the 2D raster scan, and δj and δk are offsets caused by parasitic movements of the sample stage during rotation. Typically the resolution of the reconstruction is matched with the raster scan such that $a = b$.

Finally, to include the scattering from all probed voxels, we introduce the factor $P_{sjk,xyz}$ which describes how much the xyz -voxel overlaps with the incoming beam at the position j, k of the raster scan at the goniometer setting s . $P_{sjk,xyz}$ takes a value between 0 and 1 with the value 0 for any voxel that does not intersect the x-ray beam. Using this factor we can now write the scattered intensity as a sum over all voxels in the voxel grid:

$$I(\varphi)_{sjk} = \sum_{xyz} P_{sjk,xyz} f_{xyz}^{2D}(\mathbf{R}_s^\top \hat{q}(\varphi)). \quad (5)$$

Figure 2c gives a graphical interpretation of the $P_{sjk,xyz}$ coefficients. By combining Eqs. (2) and (5) we can now write up the full forward model for TT:

$$I_{sjkc} = \sum_{xyz} P_{sjk,xyz} \sum_i B_{sic} c_{xyzi} \quad (6)$$

$$\Leftrightarrow \mathbf{I} = \mathbf{A}\mathbf{c}, \quad (7)$$

where on the second line we have defined the data vector $\mathbf{I}$, the system matrix $\mathbf{A}$, and the coefficient vector $\mathbf{c}$ in order to write the problem in linear algebra terms. The system matrix has the block matrix structure

$$\mathbf{A} = \begin{bmatrix} [P_{0jk,xyz}] \otimes [B_{0c,i}] \\ [P_{1jk,xyz}] \otimes [B_{1c,i}] \\ \vdots \\ [P_{N_s jk,xyz}] \otimes [B_{N_s c,i}] \end{bmatrix}, \quad (8)$$

where $\otimes$ is the Kronecker product. Note that the system matrix does not factorize into a projection part and a reciprocal-space part, as both the projection operator and the basis function matrix depend on the orientation of the sample. This structure highlights the difference between tensor tomography and many other multi-modal tomography techniques such as XRD-CT, x-ray fluorescence tomography[31], time-resolved tomography, and spectral tomography [32], where the real-space projection operation and the mapping of the other modality are decoupled. This prevents the use of many techniques that rely on this factorization such as principal-component-analysis methods [33].

With the forward model defined, we can now formulate the inversion as the solution of a minimization problem:

$$\mathbf{c}^* = \underset{\mathbf{c}}{\operatorname{argmin}} \left[\|\mathbf{I} - \mathbf{A}\mathbf{c}\|_a^a + \mu \|\mathbf{D}\mathbf{c}\|_b^b + \dots \right], \quad (9)$$

where $\|\cdot\|_a$ and $\|\cdot\|_b$ are two, potentially identical, vector norms, μ is a regularization parameter, and $\mathbf{D}$ is a weight matrix. The ellipsis indicates that more regularization terms of the same form as $\mu \|\mathbf{D}\mathbf{c}\|_b^b$ may be added.

III. IMPLEMENTATION

MUMOTT is written in Python with performance-critical parts implemented using the NUMBA package [34] for CPU and GPU acceleration in order to balance computational efficiency, portability, and maintainability. It also depends on NUMPY [35], SCIPY [36], SCIKIT-IMAGE [37], and COLORCET [38]. The package is extensively documented and the documentation is available online at <https://mumott.org> and at <https://doi.org/10.5281/zenodo.7919448>, including various examples in the form of Jupyter notebooks.

A variety of common tasks pertaining to data alignment and reconstruction is accessible via functions that provide a rather simple yet customizable interface. These functions represent so-called “pipelines” (Sect. IIID) and are intended to serve as the primary interface for most users.

The pipeline functions combine a number of individual tasks and components, which are represented via objects and are part of the underlying object-oriented framework (Sect. IIIE). Through the latter, advanced users and developers can customize, adapt, and extend the functionality of MUMOTT. MUMOTT is released under the Mozilla

TABLE II. Minimal example of a reconstruction workflow using the modular iterative tomographic reconstruction algorithm (MITRA) pipeline (Sect. IIID 1).

```
# Load data
data_container = DataContainer('trabecular_bone.h5')

# Perform alignment
shifts, _, _ = run_optical_flow_alignment(
    data_container, use_gpu=True)
data_container.geometry.j_offsets = shifts[:, 0]
data_container.geometry.k_offsets = shifts[:, 1]

# Execute reconstruction pipeline
result = run_mitra(data_container)
```

Public License Version 2.0 and developed as free-and-open-source software, inviting the contributions of other groups and developers.

In the following, we first provide a short demonstration of the workflow (Sect. III A) before addressing basis sets (Sect. IIIC), several common pipelines (Sect. IIID), the underlying object-oriented framework (Sect. IIIE), and computational efficiency (Sect. IIIF).

A. Workflow

Tables II and III show examples of simple workflows in MUMOTT for reconstructing a voxel map of 2D-RSMs from experimental data. In the following sections we explain each of the steps in this process.

1. Loading the data.

The acquired data (after beamline specific preprocessing) is handled using a `DataContainer`, which is created by loading a HDF5 file (Table IV) that contains the azimuthally re-grouped data for one q -bin of the experiment. While a full experimental dataset containing a detector frame for each scan position can be quite large, commonly on the order of 100s of gigabytes, a single q -bin of the azimuthally regrouped data is usually hundreds of megabytes to a few gigabytes. The data files can be prepared containing the geometry data and sample offsets information or only the data.

2. Definition of the geometry.

The geometry is defined by the vectors listed in Table I, which can be given in any consistent coordinate system. In the examples shown, the full geometry data is already contained in the data file (and hence the `DataContainer`), but in general it is possible to override certain parameters after loading.

TABLE III. Extended example for a reconstruction workflow using a Tikhonov (L_2) regularized least-squared model and spherical harmonics as basis functions that uses the object-oriented interface (Sect. III E).

```
# Load data
data_container = DataContainer('trabecular_bone.h5')

# Perform alignment
shifts, _, _ = run_optical_flow_alignment(
    data_container, use_gpu=True)
data_container.geometry.j_offsets = shifts[:, 0]
data_container.geometry.k_offsets = shifts[:, 1]

# Define forward model
projector = SAXSProjectorCUDA(
    data_container.geometry)
basis_set = SphericalHarmonics(ell_max=8)
residual_calculator = GradientResidualCalculator(
    data_container, basis_set, projector)
loss_function = SquaredLoss(residual_calculator)
l2_norm = L2Norm()
loss_function.add_regularizer(
    name='l2norm', regularizer=l2_norm,
    regularization_weight=12e-6)

# Carry out reconstruction
optimizer = LBFGS(loss_function, maxiter=20)
result = optimizer.optimize()
```

3. Aligning the data.

Before a meaningful reconstruction can be carried out the data must be aligned, which means to calculate the offsets defined in Eq. (4). To this end, MUMOTT provides several pipelines that use the transmission measurement or the average scattering to correct misalignment between each projection that occur due to parasitic movements during acquisition. In the examples shown here, we use the function that implements the optical flow alignment procedure [39], which relies on center-of-mass and tomographic consistency techniques. The alignment functions return most importantly the shifts that are needed for aligning the data. These values are then used to override the offsets stored in the `DataContainer` object.

4. Defining the reconstruction model.

The reconstruction model is defined by the choice of basis functions, the form of the cost function as well as the regularization terms. A large number of different algorithms can be constructed by combining these three choices. The simplest approach is to utilize one of the existing pipelines (Sect. III D) as illustrated by the first example (Table II), in which the modular iterative tomographic reconstruction algorithm (MITRA) pipeline (Sect. III D 1) is used. Alternatively one can configure a reconstruction model using the individual objects that

TABLE IV. Outline of the HDF5 file format used by MUMOTT. The spaces indicate the hierarchy of entries; 0 is an entry in the group `projections`, whereas `data` is an entry in the group 0, and so on.

Path	Type
p_direction_0	float(3)
j_direction_0	float(3)
k_direction_0	float(3)
detector_direction_origin	float(3)
detector_direction_positive_90	float(3)
inner_axis	float(3)
outer_axis	float(3)
volume_shape	int(3)
detector_angles	float(n_φ)
projections	Group
0	Group
data	float(n_j, n_k, n_φ)
diode	float(n_j, n_k)
inner_angle	float(1)
j_offset	float(1)
k_offset	float(1)
outer_angle	float(1)
weights	float(n_j, n_k, n_φ)
1	Group
:	

represent the different components. This approach is demonstrated by the second example (Table III), where we choose a basis of spherical harmonics in combination with a squared-difference loss function and Tikhonov (L_2) regularization.

5. Minimizing the loss function.

Once the loss function is defined, the optimization problem can be solved using one of a number of optimization routines. While this step is included in the case of the predefined reconstruction pipeline in the first example (Table II), it needs to be explicitly specified when constructing the workflow as in the second example (Table III), where we use the gradient-based LBFGS optimizer. In the case of regularized models, one should then perform a sweep of the regularization parameter space in order to determine (a) sensible regularization parameter(s).

B. Deriving standard quantities from the output.

The result of a tensor tomography reconstruction is the array of optimized coefficients, $\mathbf{c}^* = [c_{xyzi}^*]$, which are the voxel-by-voxel expansion coefficients of the local 2D-RSM shells in terms of the specific basis functions. In general, the coefficients can be interpreted using the corresponding basis set to compute latitude-longitude maps of the

2D-RSMs shells and reconstructions of several q -bins can be combined to construct 3D-RSMs from these maps. Furthermore, a number of derived quantities are conventionally used for evaluation and visualization of reconstructions and can be calculated efficiently from the coefficients without needing to compute latitude-longitude maps. Note that we define here the derived quantities which are part of the output structure of MUMOTT VX.X, additional quantities can be calculated from the array of optimized coefficient, depending on the basis function.

The mean scattering intensity, also called the isotropic intensity, is a measure of the density of scattering material in all orientations and is defined as:

$$\begin{aligned}\bar{f} &= \langle f_{xyz}^{2D}(\hat{q}) \rangle_{\hat{q}} \\ &= \frac{1}{4\pi} \int_0^\pi \left(\int_0^{2\pi} f_{xyz}^{2D}(\hat{q}(\theta, \phi)) d\phi \right) \sin \theta d\theta,\end{aligned}\quad (10)$$

where θ and ϕ are a pair of polar coordinates for the unit sphere.

A 2D-RSMs can also be expanded in tensor-components. Especially the rank-2 tensor is of interest because it allows easily computing primary directions, given by the eigenvectors of the matrix. The second moment tensor is a 3 by 3 matrix with elements:

$$M_{ij} = \langle q_i q_j f_{xyz}^{2D}(\hat{q}) \rangle_{\hat{q}}, \quad (11)$$

where q_i are the x , y , and z components of $\hat{q}$ for $i = 1, 2, 3$ respectively.

For many samples, a main orientation, such as a fiber-symmetry axis of the nanostructure, can be defined per voxel. The rank-2 tensor provides a means to efficiently compute this direction through its eigendecomposition. However, we point out that the interpretation of the main orientation of the nanostructure depends on its scattering characteristics. Structures where a single direction of strong scattering is expected at two opposite poles, the main orientation is the eigenvector corresponding to the largest eigenvalue. Similarly, for samples where a ring equatorial band with strong scattering is observed (e.g., the structure displayed in Fig. 3) the eigenvector corresponding to the smallest eigenvalue should be chosen. This provides a fast and noise-tolerant approach to finding the main nanostructure orientation except in cases where the rank-2 term vanishes, such as in Bragg scattering from cubic-symmetric materials, where more advanced approaches are needed.

Another quantity frequently used to describe the anisotropy [40] of tensor tomography is the fractional anisotropy (FA) which can be computed from the eigenvalues of the 2nd moment tensor:

$$FA = \frac{\sqrt{(\lambda_1 - \lambda_2)^2 + (\lambda_2 - \lambda_3)^2 + (\lambda_3 - \lambda_1)^2}}{\sqrt{2(\lambda_1^2 + \lambda_2^2 + \lambda_3^2)}}, \quad (12)$$

where $\lambda_1, \lambda_2, \lambda_3$ are the three eigenvalues of the second moment tensor. $FA = 0$ for perfectly isotropic scattering




FIG. 3. Comparison of (a) spherical harmonics (SH), (b) zonal harmonics (ZH), (c) nearest neighbors (NN), and (d) Gaussian kernels (GK) basis sets. The upper row shows a single basis function for each basis set. The lower row shows the RSM at a single q of a single voxel of a reconstruction using each of the four basis sets. The crosses in (c-d) show the grids which are part of the definition of the NN and GK basis sets. The directions used to define the NN model are the face centers of the truncated icosahedron. The directions used in the GK model are given by a modified Kurihara mesh. The circles in (b) indicate the symmetry axes.

and reaches a maximum value of 1 when there is strong scattering in one direction and the scattering goes to zero in the orthogonal directions. Other values can be calculated from the coefficients to describe the anisotropy, also referred to as the degree of orientation, and are described elsewhere [19, 23, 41].

C. Basis sets

The most notable difference between different reconstruction algorithms is the choice of basis functions. Figure 3 shows a comparison of the optimized 2D-RSM shell of a single voxel of the same sample using four different basis set types.

1. Spherical harmonics (SH).

The spherical harmonics are a set of orthogonal polynomials that derive from the solution to the Laplace equation in spherical coordinates. Any function on the unit sphere can be represented by an infinite expansion in spherical harmonics, but in practice the expansion must be truncated at some finite order. Such a finite expansion in spherical harmonics is called a band-limited spherical function, and can be used to represent the 2D-RSM [23]. The SH basis set is fully defined by the band limit, $\ell_{\max}$, at which the expansion in spherical harmonics is truncated. This sets the resolution of the narrowest diffraction features that can be reconstructed to approximately $2\pi/\ell_{\max}$ radians.

2. Nearest neighbors (NN).

A NN basis set uses a set of NN indicator functions. This model is therefore defined by a grid of orientations alone and the resolution is set by the distance between grid points. This basis set can be used to emulate the algorithm first presented in Ref. 20, which splits the TT problem into a set of independent scalar tomography problems.

3. Gaussian kernels (GK).

Like the NN basis set, the GK basis set is defined by a grid of orientations, but instead of indicator functions it uses smooth spherical Gaussian functions, rotated to be centered on the various grid orientations. It therefore needs one extra parameter to define the basis set namely the width of the kernel. The GK basis set has many of the same properties as the NN basis set but unlike the former results in smooth RSMs. The resolution depends both on the distance between grid points and the kernel width. Such spherical kernels are commonly used in texture analysis where the specific function used here is referred to as the Bunge normal distribution[42] to distinguish it from several other Gaussian shaped kernel functions that are frequently used.

4. Zonal harmonics (ZH).

The axially symmetric method established in Ref. 19 uses a ZH basis set and thereby differs from the other methods implemented in MUMOTT by having a non-linear forward model. This requires a separate workflow involving a specialized calculator for the gradients and optimizer. To enable high-order expansions, simplify the code, and ensure interoperability between the ZH and SH workflows, rotations and gradients are calculated in coefficient space using Wigner D-matrices. This allows orders up to $\ell_{\max} = 100$ in the current implementation, although orders higher than $\ell_{\max} \approx 30$ are difficult to handle in practice due to the large number of coefficients. Details of the implementation are given in [43].

The non-linearity of the forward model in the ZH approach makes the loss function non-convex, which renders the optimization problem more challenging. Approaches to overcome this difficulty include regularization of the angle parameters and smoothing of the gradient [21], and the use of an ensemble of randomized starting points [23]. In MUMOTT we use a starting guess provided by a different reconstruction algorithm to determine the symmetry direction.

D. Pipelines

MUMOTT provides various pipelines that implement reconstruction and alignment workflows. The former include both “standard” and asynchronous pipelines. The *standard pipelines* can be run using both CPU and GPU resources and are usually highly customizable. The *asynchronous pipelines* are optimized for GPU resources and thus speed, and usually slightly less adjustable. They employ asynchronous execution on the GPU to avoid the overhead caused by transferring data between CPU and GPU.

1. Standard reconstruction pipelines.

The simultaneous iterative reconstruction technique (SIRT) is a popular reconstruction algorithm thanks to its inherent regularizing properties that result from semi-convergence [44] and the small number of tunable parameters. It has previously been used for tensor tomography by, e.g., Ref. 20 and Ref. 18. In MUMOTT a traditional approach to simultaneous iterative reconstruction technique (SIRT), is implemented in the SIRT pipeline.

While the SIRT algorithm is not conventionally stated as a minimization problem, it has been shown that it is equivalent to a specific preconditioned gradient-descent weighted least-squares optimization [45]. Through this re-formulation, the basic SIRT reconstruction becomes compatible with various regularizers. The weight-preconditioner approach employed in this form of SIRT can also be extended to the RSM given by Eq. (3). This approach is implemented in the modular iterative tomographic reconstruction algorithm (MITRA) pipeline, which permits arbitrary regularizers and basis sets to be used, as well as Nesterov momentum acceleration.

The spherical integral geometric tensor tomography (SIGTT) pipeline sets up the basic reconstruction model using a SH basis set, a squared-difference loss function, and regularization using a finite-difference Laplacian filter [23]. The optimization problem is solved with the LBFGS-B algorithms and uses a stop criterion based on the relative change of the loss function.

The discrete directions (DD) pipeline emulates the reconstruction technique used by Ref. 20, which splits the tensor reconstruction into a set of independent scalar reconstructions, using the NN basis set. The pipeline employs the SIRT algorithm for the individual scalar reconstructions. Discrete directions (DD) has the practical advantage of needing less VRAM than methods which reconstruct the entire RSM at once, as it only loads one scalar component onto the GPU at a time.

2. Asynchronous pipelines.

These pipelines, optimized for GPU execution and speed, include a tensor SIRT pipeline, which is similar to

MITRA without Nesterov momentum. In addition, there is momentum total variation reconstruction (MOTR), which is essentially the default MITRA pipeline with L_1 and two-sided total variation regularization. Finally, robust and denoised tensor tomography (RADTT) optimizes for the Huber norm with two-sided total variation regularization through Nesterov accelerated gradient descent. The latter pipeline requires fine-tuning of the configuration in order to converge reasonably well, but once a proper step size and smoothing terms are found, it is relatively robust against noise.

There are also sparse versions of the asynchronous pipelines, which use a modified version of the John transform that calculates the reciprocal-space and real-space projection operations simultaneously within one kernel, using a sparse approximation to the reciprocal-space mapping. This is not necessarily faster than computing the two mappings separately (unless the representation is very sparse, such as only mapping one basis function to each segment), but it uses less VRAM, as it is not necessary to store the intermediate result between carrying out the John transform and carrying out the reciprocal space mapping.

3. Alignment.

The objective of alignment is to determine the offsets Δj and Δk of Eq. (4) that result from parasitic movement and misalignment of the goniometer and drift during the experiment [46]. The alignment step is essential in tomography as any misalignment will be reflected as an artifact in the tomographic reconstruction. There are many algorithms to solve this problem, leading to sub-pixel alignment accuracy, taking into account various experimental systems and data.

MUMOTT currently provides two alignment pipelines. Both algorithms typically work with the transmitted intensity data stored in the `DataContainer` or another isotropic signal such as the azimuthally integrated intensity. They iteratively update the offsets for each projection by reconstructing the absorption tomogram via a projector. The overall workflow is shown in Fig. 4.

The *phase matching alignment* pipeline is based on cross-correlation and follows Ref. 47. Cross-correlation alignment has been proven for continuous objects in electron microscopy tomography by Ref. 48 and has been widely used since. The principle is to determine the offsets by means of correlation functions formed from image pairs of the projections, comparing the center of mass of the image pair correlation peaks. This method is fast and can provide sub-pixel accuracy for data with small misalignment.

When the data exhibit misalignment of multiple pixels, the cross-correlation alignment alone can struggle to find the appropriate coordinate transformation. For such cases, MUMOTT provides the *optical flow alignment* pipeline, which implements a toolbox algorithm based on




FIG. 4. Alignment pipeline workflow. Steps shown with dashed outlines apply only for the optical flow alignment pipeline.




FIG. 5. Slices from absorption reconstructions (a) before and after alignment with (b) the phase matching method and (c) the optical flow method. The projections have been sorted so that the projection directions of neighboring projections are as close to each other as possible. (d) Rotations and tilts as well as alignment offsets from (e) the phase matching method and (f) the optical flow method. Note how the rotations correlate with changes in Δj , whereas the tilts correlate with changes in Δk . In this case the offsets result from misalignment of the goniometer's rotation axes with the sample center.




FIG. 6. Outline of the object-oriented framework in MUMOTT. Orange boxes show input parameters and data provided by the user, blue ovals show objects, the green box shows the output, and arrows indicate instances of objects interacting with one another.

the work of Ref. 39. This approach uses multiple successive and interconnected alignment procedures, including optical flow projection matching alignment, line vertical alignment, and weight centering. The method is tunable through various parameters and filters and is therefore able to align extremely misaligned data, providing an approach that is usable for a larger variety of experimental data.

Examples of alignment results with the two pipelines are shown in the case of the publicly available experimental dataset of trabecular bone in Fig. 5 [49].

E. Object-oriented framework

The internal architecture of MUMOTT consists of an object-oriented framework with some elements of functional programming. The structure of the framework is described in Fig. 6. Many objects are safely mutable after instantiation and employ hashes of their mutable properties to track the state of linked instances, which means that derived properties can be recomputed when required.

1. Data and geometry.

The **DataContainer** is the owner of the input data, which is stored in HDF5 format. The input (measurements and geometry metadata) is stored as a list of projections, indexed by the direction index s as given in Eq. (5), and the measured tensor tomographic data

can be accessed as a four-dimensional array indexed by $[s, j, k, i]$ as in Eq. (7). The information related to geometry is stored in a **Geometry** object, which is directly linked to the list of projections attached to **DataContainer**. Thus, if a projection is removed from the list, this will be reflected in the corresponding geometry data being removed from the **Geometry** instance. The **Geometry** object stores the basis vectors of the system, i.e., $(\mathbf{p}, \mathbf{j}, \mathbf{k}, \mathbf{q}_0, \mathbf{q}_{90}, \hat{\alpha}, \hat{\beta})$ listed in Table I and shown in Fig. 1. The vectors must be specified in the laboratory coordinate system, which coincides with the sample-fixed coordinates (x, y, z) when $\mathbf{R}(s) = \mathbf{I}$, i.e., the identity transform. A rotation operator is then specified, $\mathbf{R}(s)$, which may be given as a rotation matrix or as an axis-angle quadruplet $(\hat{\alpha}, \alpha(s), \hat{\beta}, \beta(s))$. Using the rotation operator, vectors in the sample-fixed coordinates are then dynamically computed for each s . This information can be specified in the input data file or by the user through direct modification of the **Geometry** object (see, e.g., Table II).

2. Projectors and basis sets.

The **Projector** and **BasisSet** classes contain the methods and properties needed to compute the forward model defined in Eq. (7) and its adjoint. The **Projector** objects depends on a **Geometry** object and employ routines implemented using the NUMBA package [34] to compute the spatial part of the transform, i.e., the matrix elements $P_{sjk,xyz}$ in Eq. (5). This is implemented for both CPU and GPU-based computation, the latter using the NUMBA interface for CUDA. The implementation employs an approach based on Joseph’s method [50] using bilinear interpolation of the field for the forward and the projection for the adjoint computation, respectively, based on the work of Ref. 51 and Ref. 52.

The **BasisSet** evaluates the constants B_{sic} in Eq. (3) for the respective basis $\mathbf{B}$ (Sect. III C) using the provided detector geometry and rotation operator $\mathbf{R}_s^T$. The integral is evaluated using adaptive Newton-Cotes quadrature or approximated using the central angle of each segment. In addition, **BasisSet** provides a routine for computing various properties of reconstructed tensors, such as the orientation as defined by the rank-2 tensor component of the field, the spherical mean, the variance, and relative anisotropy (the spherical standard deviation normalized by the mean).

3. Residual calculation.

The **ResidualCalculator** is a managing object which takes a **DataContainer**, **Projector**, and **BasisSet**, and uses them to compute residuals. It tracks the current reconstruction, i.e., c_{xyzi} . In other words, for the data

D_{sjkc} , and the current reconstruction c_{xyzi} it computes

$$\mathbf{r} = \mathbf{A}\mathbf{c} - \mathbf{I} \quad (13)$$

where $\mathbf{r}$ and $\mathbf{I}$ are flattened vectors of the residual and data matrices respectively using the notation introduced in Eq. (7). It also computes the gradient of a residual norm, which is used by the gradient-based optimization algorithms implemented in MUMOTT.

A special `ZonalHarmonicsGradientCalculator` is defined to be used as a part of the ZH workflow. It is used to map a list of ZH coefficients and two angle coordinates into the space of all spherical harmonics (up to a maximum order) in the sample-coordinate system, and to compute gradients with respect to the ZH coefficients and the angles.

4. Optimization.

The goal of the optimization is to minimize a `LossFunction` (also known as an objective function) by tuning the coefficients of the underlying model by using an `Optimizer`. The `LossFunction` combines a `ResidualCalculator` with one or several `Regularizer` instances and can be given a preconditioner to weight the gradient.

There are currently two types of loss functions that support standard least-squares regression (`SquaredLoss`) and robust regression via the Huber-regressor [53] (`HuberLoss`), respectively.

There are also various regularization options. One can, e.g., smoothen the solution by minimizing the squared L_2 norm of the finite-difference Laplacian operator of the tensor field (`Laplacian`). It is also possible to smoothen the solution in a more robust manner by minimizing the Huber norm of the spatial gradient for each basis-set mode (`TotalVariation`). While it can be more difficult to obtain convergence with more robust terms, it can also be configured to use the Huber approximation for small values to improve convergence.

Other `Regularizer` classes are available to minimize the L_2 and L_1 norms of the tensor field, respectively. While the L_2 norm (`L2Norm`) penalizes large values, which promotes rapid convergence, the L_1 norm (`L1Norm`) encourages sparse solutions and tends to reduce noise in the solution. Finally, one can also use the Huber norm of the tensor field (`HuberNorm`), which acts as an `L1Norm` for large values and an `L2Norm` for small values, converging more easily than `L1Norm`. When applied with the SH basis-set, the `L1Norm` and `HuberNorm` are not rotational invariants, and can bias the solution towards certain directions.

In terms of optimizers MUMOTT provides gradient descent with a fixed step size (`GradientDescent`), with an option to use Nesterov accelerated momentum, as well as the LBFGS-B algorithm for quasi-Newton solution of the optimization (`LBFGS`). For the ZH workflow (Sect. III C 4) there is both a specialized optimizer

(`ZonalHarmonicsOptimizer`) and a gradient calculator (`ZonalHarmonicsGradientCalculator`). It is a basic gradient descent optimizer with a special heuristic rule to determine a safe step size for the angle parameters. Because of the non-convexity of the cost function, it requires a good starting guess for the angles in order to converge to a solution.

F. Computational efficiency and resource requirements

The computational resources required to perform reconstructions in MUMOTT are modest compared to previous implementations due to efficient implementations of the John transform and the use of memory-efficient solvers. The place where a user is most likely to run into problems is the memory requirement for the data set and solution vector, in addition to a few extra similarly sized arrays needed by the solvers. The memory requirement is around a few gigabytes in the most common use cases, but increases with both the size of the voxel grid and the directional resolution. In order to use the GPU-implementation, one requires a CUDA-compatible GPU with sufficient VRAM to store an array the size of the solution vector.

In general, the reconstruction is much less resource-hungry than the preceding data-reduction steps. However, when conducting sweeps of regularization parameters and full q -resolved 3D-RSM reconstructions, the reduced runtime from GPU-acceleration has a considerable impact.

Table V compares the run times for different pipelines, platforms, and configurations. Each configuration was run 10 times on each platform, and the result was obtained by averaging the run times after discarding the first run, to enable on-disk caching to take place. All runs were carried out with a maximum of 20 iterations, although spherical integral geometric tensor tomography (SIGTT) converged in 14 or fewer iterations in all cases.

The runs were carried out in separate, sequentially run processes, which means that just-in-time compiled kernels were not re-used beyond what is automatically cached on disk. This has the largest effect on DD, which creates sub-geometries for each basis function and therefore needs to re-compile code to carry out the John transform for each sub-iteration, which adds approximately a second of overhead per basis function. Therefore, DD can perform substantially better than what is apparent from this table when the same geometry is run multiple times in a single process, as may be done for q -resolved reconstruction. The time required to load data was not included in the timing to eliminate the dependency on the file systems used for benchmarking.

TABLE V. Comparison of reconstruction times in seconds averaged across 10 runs each for a typical single- q dataset consisting of 247 projections, each with 65×55 pixels and 8 detector segments, using different reconstruction pipelines and running on different computers. N is the number of basis functions per voxel. In all cases, relative uncertainties were smaller than 5%, and are omitted to maintain ease of reading. The workstation (WS) data was obtained using an AMD Ryzen 7 3700X processor with 8 physical cores, 64 GB of DDR4 2666 MHz RAM, and for the GPU accelerated calculations an Nvidia GeForce RTX 3060 GPU with 12 GB of VRAM. The high-performance computing (HPC) CPU timings were generated using 8 top-level threads on a 64-core Intel Xeon Platinum 8358 @ 2.0 GHz CPU with some operations utilizing lower-level multithreading. The HPC GPU timings were obtained using an Nvidia A100 GPU with 40 GB of VRAM and 8 threads on 16 cores of a 64-core Intel Xeon Platinum 8358 @ 2.0 GHz.

	CPU			GPU	
	N	WS	HPC	WS	HPC
SIGTT	6	23	18	9	9
	20	45	29	18	14
	72	108	69	60	36
MITRA	18	41	22	13	8
	50	93	45	40	14
	162	271	156	115	37
DD	18	81	72	40	43
	50	156	157	101	111
	162	334	392	290	346
MOTR	18			9	8
	50			12	10
	162			46	22

IV. OUTLOOK

Various additions and improvements to MUMOTT are foreseen for the future. One of the main difficulties of performing tensor-tomography experiments at present in the interfacing with the existing data analysis pipelines at the various synchrotron end stations with the reconstruction pipeline. At present, such an integration relies on two intermediate steps. The first is azimuthal regrouping of detector images, which results in a number of new data files containing the azimuthally regrouped intensities that are organized projection-by-projection mirroring the order in which the experiment was performed. The second step is a slicing of the experimental dataset into the MUMOTT-compatible HDF5 files described in Table IV, which contains the data organized by q -bins. These extra analysis steps are often slow due to requiring many read- and write operations. On-the-fly reconstruc-

tions would require live azimuthal regrouping of detector images and a more efficient data pipeline that allows fast slicing in the q -dimension.

At present, MUMOTT is able to compute various properties of reconstructions and to save the results to HDF5 files. The user then has the responsibility for analysis and visualization of the reconstructed quantities. It will be useful to add the option to write to formats compatible with common visualization software packages.

In Ref. 23, simulated data was used for the purpose of validation and comparison of various reconstruction methods. Being able to easily generate simulated data in MUMOTT would be useful not just for validation, but also to plan experiments and to generate synthetic data for training machine learning models.

The splitting of the tensor-tomography reconstruction into discrete 2D-RSM shells is a useful simplification that reduces the size of individual reconstruction problems. It would however often be advantageous to combine several q -bins into a single reconstruction to enforce certain types of prior knowledge of the nanostructure on the reconstruction for sample systems where an appropriate model is available. An example of this is texture tomography [54] applicable to Bragg-scattering from nanocrystalline materials where the rotational symmetries of the crystal lattice can be imposed on the reconstruction by performing a combined reconstruction of several q -bins at once. Also in the case of fiber-scattering, different q -ranges can contain scattering from different sample orientations and a combined approach to reconstruction is expected to be able to alleviate missing wedge artifacts in the reconstructions.

ACKNOWLEDGMENTS

This work was funded by the European Research Council Starting Grant MUMOTT (ERC-2020-StG 949301), the Swedish research council (VR 2018-041449), and the Chalmers Initiative for Advancement of Neutron and Synchrotron Techniques. MC has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No 884104. Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. The computations and optimizations described in this project were in part enabled by computational resources provided by Chalmers e-Commons at Chalmers University of Technology as well as the National Academic Infrastructure for Supercomputing in Sweden at NSC and C3SE partially funded by the Swedish Research Council through grant agreement No. 2022-06725.

- [1] P. Fratzl and R. Weinkamer, Nature's hierarchical materials, *Progress in Materials Science* **52**, 1263 (2007).
- [2] N. Reznikov, R. Shahar, and S. Weiner, Bone hierarchical structure in three dimensions, *Acta Biomaterialia* **10**, 3815 (2014), biomineralization.
- [3] B. Schrauwen, R. Janssen, L. Govaert, and H. Meijer, Intrinsic deformation behavior of semicrystalline polymers, *Macromolecules* **37**, 6069 (2004).
- [4] N. Stribeck, U. Nöchel, S. S. Funari, and T. Schubert, Tensile tests of polypropylene monitored by saxs. comparing the stretch-hold technique to the dynamic technique, *Journal of Polymer Science Part B: Polymer Physics* **46**, 721 (2008).
- [5] Y. Tang, Z. Jiang, Y. Men, L. An, H.-F. F. Enderle, D. Lilge, S. V. Roth, R. Gehrke, and J. Rieger, Uniaxial deformation of overstretched polyethylene: In-situ synchrotron small angle x-ray scattering study, *Polymer* **48**, 5125–5132 (2007).
- [6] S. Gantenbein, K. Masania, W. Woigk, J. P. W. Sesseg, T. A. Tervoort, and A. R. Studart, Three-dimensional printing of hierarchical liquid-crystal-polymer structures, *Nature* **561**, 226 (2018).
- [7] M. Endrizzi, X-ray phase-contrast imaging, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **878**, 88 (2018).
- [8] X. Ou, X. Chen, X. Xu, L. Xie, X. Chen, Z. Hong, H. Bai, X. Liu, Q. Chen, L. Li, and H. Yang, Recent development in x-ray imaging technology: Future and challenges, *Research* **2021** (2021).
- [9] M. Georgiadis, R. Müller, and P. Schneider, Techniques to assess bone ultrastructure organization: orientation and arrangement of mineralized collagen fibrils, *Journal of The Royal Society Interface* **13**, 20160088 (2016).
- [10] T. H. Jensen, M. Bech, O. Bunk, T. Donath, C. David, R. Feidenhans'l, and F. Pfeiffer, Directional x-ray dark-field imaging, *Physics in Medicine & Biology* **55**, 3317 (2010).
- [11] M. Busi, J. Shen, M. Bacak, M. C. Zdora, J. Čapek, J. Valsecchi, and M. Strobl, Multi-directional neutron dark-field imaging with single absorption grating, *Scientific Reports* 10.1038/s41598-023-42310-y (2023).
- [12] N. Stribeck, A. A. Camarillo, U. Nöchel, C. Schroer, M. Kuhlmann, S. V. Roth, R. Gehrke, and R. K. Bayer, Volume-resolved nanostructure survey of a polymer part by means of saxs microtomography, *Macromolecular Chemistry and Physics* **207**, 1139 (2006).
- [13] C. G. Schroer, M. Kuhlmann, S. V. Roth, R. Gehrke, N. Stribeck, A. Almendarez-Camarillo, and B. Lengeler, Mapping the local nanostructure inside a specimen by tomographic small-angle x-ray scattering, *Applied Physics Letters* **88**, 164102 (2006).
- [14] U. Kleuker, P. Suortti, W. Weyrich, and P. Spanne, Feasibility study of x-ray diffraction computed tomography for medical imaging, *Physics in Medicine & Biology* **43**, 2911 (1998).
- [15] S. R. Stock, F. D. Carlo, and J. D. Almer, High energy x-ray scattering tomography applied to bone, *Journal of Structural Biology* **161**, 144 (2008).
- [16] P. Bleuet, E. Welcomme, E. Dooryhee, J. Susini, J.-L. Hodeau, and P. Walter, Probing the structure of heterogeneous diluted materials by diffraction tomography., *Nature Materials* **7**, 468 (2008).
- [17] A. Malecki, E. Eggl, F. Schaff, G. Potdevin, T. Baum, E. G. Garcia, J. S. Bauer, and F. Pfeiffer, Correlation of x-ray dark-field radiography to mechanical sample properties, *Microscopy and Microanalysis* **20**, 1528–1533 (2014).
- [18] J. Kim, M. Kagias, F. Marone, and M. Stampanoni, X-ray scattering tensor tomography with circular gratings, *Applied Physics Letters* **116**, 134102 (2020).
- [19] M. Liebi, M. Georgiadis, A. Menzel, P. Schneider, J. Kohlbrecher, O. Bunk, and M. Guizar-Sicairos, Nanostructure surveys of macroscopic specimens by small-angle scattering tensor tomography, *Nature* **527**, 349 (2015).
- [20] F. Schaff, M. Bech, P. Zaslansky, C. Jud, M. Liebi, M. Guizar-Sicairos, and F. Pfeiffer, Six-dimensional real and reciprocal space small-angle x-ray scattering tomography, *Nature* **527**, 353 (2015).
- [21] M. Liebi, M. Georgiadis, J. Kohlbrecher, M. Holler, J. Raabe, I. Usov, A. Menzel, P. Schneider, O. Bunk, and M. Guizar-Sicairos, Small-angle X-ray scattering tensor tomography: model of the three-dimensional reciprocal-space map, reconstruction algorithm and angular sampling requirements, *Acta Crystallographica Section A* **74**, 12 (2018).
- [22] Z. Gao, M. Guizar-Sicairos, V. Lutz-Bueno, A. Schröter, M. Liebi, M. Rudin, and M. Georgiadis, High-speed tensor tomography: iterative reconstruction tensor tomography (IRTT) algorithm, *Acta Crystallographica Section A* **75**, 223 (2019).
- [23] L. C. Nielsen, P. Erhart, M. Guizar-Sicairos, and M. Liebi, Small-angle scattering tensor tomography algorithm for robust reconstruction of complex textures, *Acta Crystallographica Section A* **79**, 515 (2023).
- [24] T. A. Grünwald, M. Liebi, N. K. Wittig, A. Johannes, T. Sikjaer, L. Rejnmark, Z. Gao, M. Rosenthal, M. Guizar-Sicairos, H. Birkedal, and M. Burghammer, Mapping the 3d orientation of nanocrystals and nanostructures in human bone: Indications of novel structural features, *Science Advances* **6**, eaba4171 (2020).
- [25] C. Donnelly, M. Guizar-Sicairos, V. Scagnoli, S. Gliga, M. Holler, J. Raabe, and L. J. Heyderman, Three-dimensional magnetization structures revealed with x-ray vector nanotomography, *Nature* 10.1038/nature23006 (2017).
- [26] M. Sales, M. Strobl, T. Shinohara, A. S. Tremsin, L. T. Kuhn, W. R. B. Lionheart, N. M. Desai, A. B. Dahl, and S. Schmidt, Three dimensional polarimetric neutron tomography of magnetic fields, *Scientific Reports* **8** (2017).
- [27] J. Kieffer, V. Valls, N. Blanc, and C. Hennig, New tools for calibrating diffraction setups, *Journal of Synchrotron Radiation* **27**, 558 (2020).
- [28] A. B. Jensen, T. E. K. Christensen, C. Weninger, and H. Birkedal, Very large-scale diffraction investigations enabled by a matrix-multiplication facilitated radial and azimuthal integration algorithm: MatFRAIA, *Journal of Synchrotron Radiation* **29**, 1420 (2022).
- [29] B. R. Pauw, Everything saxs: small-angle scattering pattern collection and correction, *Journal of Physics: Condensed Matter* **25**, 383201 (2013).
- [30] T. A. Grünwald, A. Johannes, N. K. Wittig, J. Palle,

- A. Rack, M. Burghammer, and H. Birkedal, Bone mineral properties and 3D orientation of human lamellar bone around cement lines and the Haversian system, *IUCrJ* **10**, 189 (2023).
- [31] M. D. de Jonge and S. Vogt, Hard x-ray fluorescence tomography—an emerging tool for structural visualization, *Current Opinion in Structural Biology* **20**, 606 (2010).
- [32] P. M. Shikhaliev, Energy-resolved computed tomography: first experimental results, *Physics in Medicine & Biology* **53**, 5595 (2008).
- [33] Z. Gao, M. Odstrčil, S. Böcklein, D. Palagin, M. Holler, D. F. Sanchez, F. Krumeich, A. Menzel, M. Stamparoni, G. Mestl, J. A. van Bokhoven, M. Guizar-Sicairos, and J. Ihli, Sparse ab initio x-ray transmission spectrotomography for nanoscopic compositional analysis of functional materials, *Science Advances* **7**, eabf6971 (2021).
- [34] S. K. Lam, A. Pitrou, and S. Seibert, Numba: A llvm-based python jit compiler, in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC* (2015) pp. 1–6.
- [35] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, Array programming with NumPy, *Nature* **585**, 357 (2020).
- [36] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, *Nature Methods* **17**, 261 (2020).
- [37] S. Van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, and T. Yu, scikit-image: image processing in python, *PeerJ* **2**, e453 (2014).
- [38] P. Kovesi, Good colour maps: How to design them, *arXiv* (2015), 1509.03700.
- [39] M. Odstrčil, M. Holler, J. Raabe, and M. Guizar-Sicairos, Alignment methods for nanotomography with deep sub-pixel accuracy, *Optics Express* **27**, 36637 (2019).
- [40] P. J. Basser and C. Pierpaoli, Microstructural and physiological features of tissues elucidated by quantitative-diffusion-tensor mri, *Journal of Magnetic Resonance, Series B* **111**, 209 (1996).
- [41] L. C. Nielsen, T. Tänzler, I. Rodriguez-Fernandez, P. Erhart, and M. Liebi, Investigating the missing-wedge problem in small-angle X-ray scattering tensor tomography across real and reciprocal space, *Journal of Synchrotron Radiation* **31**, 1327 (2024).
- [42] H. J. Bunge, *Mathematische Methoden der Texturanalyse* (De Gruyter, 1969).
- [43] M. Carlsen, C. Appel, W. Hearn, M. Olsson, A. Menzel, and M. Liebi, X-ray tensor tomography for small-grained polycrystals with strong texture, *Journal of Applied Crystallography* **57** (2024).
- [44] T. Elfving, P. C. Hansen, and T. Nikazad, Semiconvergence and relaxation parameters for projected sirt algorithms, *SIAM Journal on Scientific Computing* **34**, A2000 (2012).
- [45] J. Gregor and J. A. Fessler, Comparison of sirt and sqs for regularized weighted least squares image reconstruction, *IEEE Transactions on Computational Imaging* **1**, 44 (2015).
- [46] J. Frank, P. Penczek, and W. Liu, Alignment, classification, and three-dimensional reconstruction of single particles embedded in ice, *Scanning Microscopy Supplement* **6**, 11 (1992).
- [47] M. Guizar-Sicairos, S. T. Thurman, and J. R. Fienup, Efficient subpixel image registration algorithms, *Optics Letters* **33**, 156 (2008).
- [48] R. Guckenberger, Determination of a common origin in the micrographs of tilt series in three-dimensional electron microscopy, *Ultramicroscopy* **9**, 167 (1982).
- [49] L. Nielsen, M. Liebi, M. Guizar-Sicairos, and M. Georgiadis, Trabecular bone datasets for SAXS Tensor Tomography (2023), doi:10.5281/zenodo.10074598.
- [50] P. M. Joseph, An improved algorithm for reprojecting rays through pixel images, *IEEE Transactions on Medical Imaging* **1**, 192 (1982).
- [51] W. Xu, F. Xu, M. Jones, B. Keszthelyi, J. Sedat, D. Agard, and K. Mueller, High-performance iterative electron tomography reconstruction with long-object compensation using graphics processing units (gpus), *Journal of Structural Biology* **171**, 142 (2010).
- [52] W. Palenstijn, K. Batenburg, and J. Sijbers, Performance improvements for iterative electron tomography reconstruction using graphics processing units (gpus), *Journal of Structural Biology* **176**, 250 (2011).
- [53] P. J. Huber, Robust Estimation of a Location Parameter, *The Annals of Mathematical Statistics* **35**, 73 (1964).
- [54] M. P. K. Frewein, J. Mason, B. Maier, H. Cölfen, A. Medjahed, M. Burghammer, M. Allain, and T. A. Grünwald, Texture tomography, a versatile framework to study crystalline texture in 3D, *IUCrJ* **11** (2024).