

Laue Indexing with Optimal Transport

Tomasz Kacprzak, Stavros Samothrakitis, Camilla Buhl Larsen, Jaromír Kopeček, Markus Strobl, Efthymios Polatidis, Guillaume Obozinski

Abstract—Laue tomography experiments retrieve the positions and orientations of crystal grains in a polycrystalline samples from diffraction patterns recorded at multiple viewing angles. The use of a broad wavelength spectrum beam can greatly reduce the experimental time, but poses a difficult challenge for the indexing of diffraction peaks in polycrystalline samples; the information about the wavelength of these Bragg peaks is absent and the diffraction patterns from multiple grains are superimposed. To date, no algorithms exist capable of indexing samples with more than about 500 grains efficiently. To address this need we present a novel method: *Laue indexing with Optimal Transport* (LAUEOT). We create a probabilistic description of the multi-grain indexing problem and propose a solution based on Sinkhorn Expectation-Maximization method, which allows to efficiently find the maximum of the likelihood thanks to the assignments being calculated using Optimal Transport. This is a non-convex optimization problem, where the orientations and positions of grains are optimized simultaneously with grain-to-spot assignments, while robustly handling the outliers. The selection of initial prototype grains to consider in the optimization problem are also calculated within the Optimal Transport framework. LAUEOT can rapidly and effectively index up to 1000 grains on a single large memory GPU within less than 30 minutes. We demonstrate the performance of LAUEOT on simulations with variable numbers of grains, spot position measurement noise levels, and outlier fractions. The algorithm recovers the correct number of grains even for high noise levels and up to 70% outliers in our experiments. We compare the results of indexing with LAUEOT to existing algorithms both on synthetic and real neutron diffraction data from well-characterized samples. The code and test data are available on <https://github.com/LaueOT/laueotx>.

Index Terms—Optimal Transport, Laue Crystallography, Expectation-Maximization

1 INTRODUCTION

LAUE and Bragg diffraction are the foundation of modern experiments in crystallography and solid state physics [1], [2]. In neutron and X-ray diffraction, the incoming waves are elastically scattered by the atoms in the crystal lattice of the studied sample. Due to the regular arrangement of the atoms in the crystal lattice, constructive interference occurs and produces strong reflections, also known as Bragg peaks, which can be detected with diffraction detectors covering the relevant angular range. Bragg’s law describes the relation between the incidence angles, the lattice spacing, and the wavelength of the incoming and reflected waves, which defines the diffraction condition. This enables modern X-ray and neutron based experimental methods to infer the internal structure of crystalline samples, with numerous applications in material science, chemistry, medicine, biology, electronics, and others.

The process of inferring the orientation of crystal lattices from the recorded Bragg peak positions is a challenging task: multiple methods have been proposed to solve variants of this problem specific to different types of experimental setups (see [3] for review). One of the most difficult problems in Laue diffraction is indexing of polycrystalline samples, which contain multiple crystal grains with different orientations [4], where many overlapping spot patterns are recorded simultaneously. The three-dimensional macroscopic structures of samples can be inferred using tomo-

graphic techniques, where projections from multiple sample rotations are obtained during the experiment. The majority of the crystallographic studies use a monochromatic beam or the time-of-flight approach to resolve the wavelength relative to a diffraction peak measured. This enables immediate identification of the Miller indices corresponding to the recorded diffraction spot, as the Bragg condition is satisfied only for a single incidence angle between the crystal lattice plane and the incoming beam for a specific wavelength. While this simplifies the indexing process, this type of experiment can be time consuming, as monochromatic sources tend have lower overall intensity. The use of white beams enables faster experiments, but leads to a more complicated inference problem, as an observed spot can stem from almost any crystal lattice plane at any orientation. This creates a mixed combinatorial and continuous problem, where the peaks-to-grain assignments have to be found jointly with the corresponding grain orientations and positions. Moreover, the number of grains in the sample is not known a priori and must also be inferred. Thus, the indexing of white beam polycrystalline patterns is one of the most challenging inference problems among crystallographic experiments.

From an experimental perspective, the last two decades have seen progress in three-dimensional (3D) electron- [5]–[11] and X-ray-based [12]–[21] grain indexing and reconstruction methods, significantly improving experimental capabilities. These methods provide valuable insights into the internal structure of samples and have yielded promising results in grain orientation mapping. Despite their impressive spatial resolution, ranging from a few hundred nanometers to micrometers, these techniques are not without limitations. For example, 3D electron backscatter diffraction (3D EBSD) is a destructive method that necessitates serial sectioning

- TK is with Swiss Data Science Center, Paul Scherrer Institute
E-mail: tomasz.kacprzak@psi.ch
- SS, CBL, MS and EP are with the Paul Scherrer Institute
- JK is with the Institute of Physics, Czech Academy of Sciences
- GO is with the Swiss Data Science Center, École Polytechnique Fédérale de Lausanne

of samples. Moreover, the penetration of both electrons and X-rays into metallic specimens is constrained, limiting the volume of the sample that can be probed. On the other hand, neutron-based methods [22]–[30] have emerged as a potential solution owing to their superior penetration ability in various materials. These non-destructive techniques enable investigations into large volumes within the bulk of samples. Among the recently developed neutron diffraction-based tomography techniques for 3D grain mapping, Laue three-dimensional neutron diffraction tomography (Laue 3DNDT) [22] stands out as a particularly promising approach. It employs a polychromatic neutron beam, for multi-grain indexing [23] and morphology reconstruction [30], effectively reducing the experimental time while maximizing the information obtained from the sample.

With multiple solutions proposed for solving these types of problems [22], [31]–[37], only few are directly applicable to wide-beam polycrystalline indexing problem. Recently, [22] proposed a forward-fitting algorithm that achieved good results for samples with up to around 500 grains [23]. However, the method is greedy, incurs in a long runtime, and it is not amenable for parallelization. With upcoming improvements in experimental techniques and instrumentation in white beam diffraction imaging with neutrons [38], there is a clear need for designing novel algorithms for this problem, which allow both for rapid indexing and analyzing samples with orders of magnitude more grains.

The problem of finding orientations and positions of the grains can be approached from the computer vision perspective; the task is similar to the structure-from-motion [39], camera localization in multi-view geometry (see [40] for overview), as well as registration of point clouds [41]. The task of spot-to-grain assignment and selection of prototype grains can both be tackled using recent developments in optimal transport (OT), an area of active development [42]. In particular, the recently proposed *Sinkhorn Expectation Maximization* method [43] demonstrated the deep relationship between the OT and the Expectation-Maximization framework, which is commonly used for solving probabilistic models with latent variables. Regarding methods for finding best-fit parameters, extensive literature exists for convex optimization methods in similar computer vision problems, with coordinate descent enjoying advantage in convergence times in many cases [44].

Utilizing these ideas, we propose a novel technique for indexing of polycrystalline samples from white beam experiments, which we call LAUEOT (Laue indexing with Optimal Transport). We propose a joint inverse problem formulation for the assignment of spots to grains, and the determination of each grain’s position and orientation. Given an assignment of spots to grains, our formulation benefits from closed-form updates for finding the orientation and grain position, and relies on optimal transport for finding the assignments of spots to grains given candidate grain orientations and positions. A critical initial step of selecting prototype grains is also handled using the OT framework [45]. This inverse problem formulation enables very fast optimization using GPU-batch solver with convergence guarantees. The optimal transport methods enable solving for grain parameters and spot assignments *jointly*, which increases the precision of the estimated parameters,

the recall of grains found, and the number of correctly assigned spots.

The algorithm has two pre-processing steps: (i) creating a set of plausible grain candidates by performing a coarse single-grain fitting in the orientation and position space, and (ii) selecting the prototype grains using the optimal transport framework, following [45]. The main solver performs multi-grain fitting, where it jointly optimizes for parameters of the grains and spot-to-grain assignments using the optimal transport framework. The key aspect of the problem is the treatment of outlier spots in the data and unmatched spots in the model. In LAUEOT, we propose a novel probabilistic outlier modeling tailored to the Laue problem. We compare it to other methods in literature for treatment of outliers: unbalanced and partial OT [46], [47].

2 PREVIOUS WORK

Multiple Bragg indexing algorithms have been proposed [3]. The Hough transform approaches [31], [48] are one of the most commonly used in practice. These algorithms use 3D histograms in the orientation space, and each spot “votes” for the plausible voxels. The set of resulting grains corresponds to the voxels with the most votes. This procedure is then typically followed by a minimizer-based refinement to find the spatial positions of the grains inside the sample. This method is particularly effective for monochromatic beams, as each spot votes for relatively small number of rotations. Recently, an extension of this method for pink beams was proposed by [32], where 10-20% divergence in the beam spectrum is handled well by the algorithm. The Hough transform-based methods can be limited by the resolution of voxels in the orientation space, especially that for wider beams the number of votes given by each spot is large, which limits the use of sparse representations of the 3D histogram. Moreover, most of such methods do not include grain position estimation and use post-processing steps to find it.

To address these problems and enable wide-beam analysis, a forward-fitting method LAUE3DNDT was recently proposed [22]. This method first performs an exhaustive search in the space of orientations, keeping the grain positions fixed at the center of the sample. For each candidate, it predicts the model spots and assigns them to the measurements using a nearest neighbor method. Then, the loss is calculated as the median distance between spot pairs. If the loss is smaller than some threshold, the gradient-based optimization is performed to find the position and orientation. If certain criteria are met, the candidate is accepted and the assigned spots are removed from further analysis. The algorithm ends when the number of remaining spots is sufficiently low. The drawback of this method is its long runtime, which is dominated by slow downhill optimizer. The greedy nature of the algorithm excludes efficient parallelization.

Recently, machine learning approaches have been proposed to tackle the problem of indexing. In [37], a fast neural network was created to output grain orientations from X-ray experiments. This method calculates features for each spot and then passes them through a neural network. The feature used is a histogram of angular distances to

neighboring spots within a certain radius. The method is trained before or during the experiments on the simulations of the corresponding crystal lattice in question. It reports good performance and practically instant results. While not tested on polycrystalline data, it could suffer difficulties with interpreting multiple overlapping spot patterns for samples with very large (>1000) number of grains.

There exists rich literature in computer vision (CV) and optimal transport (OT) that is relevant to this problem. The structure-from-motion (SfM) algorithms tackle the problem of reconstructing 3D scene from a series of 2D images, each taken at a different camera position. The inverse-problem formulation with Expectation-Maximization has been successful [49], [50]. In Laue tomography, two-dimensional images are taken after rotating the sample in 3D. While those problems are different, they share many similarities, which suggests that the EM approach to be promising for the Laue reconstruction.

Registration of point clouds is also a classic problem in CV tackled by multiple classic approaches [51]. A major difficulty in this problem is estimation of a rotation matrix without correspondence, which has multiple minima. The Go-ICP algorithm finds a global minimum using a branch-and-bound method. Robust point matching [52] proposes to use a soft assignment scheme, which can be viewed as an optimal transport problem solved by the Sinkhorn algorithm [53]. It uses a deterministic annealing schedule to avoid getting stuck in local minima. As detected and modeled spots can be viewed as point clouds registration problem, this method may also have relevance for Laue diffraction.

As the number of grains is not known a-priori, Laue reconstruction has also similarities to sparse feature selection problems in the context of optimal transport. Recently, sparse optimal transport methods have been proposed [54]–[56]. However, the sparse OT problem is not convex and most of the methods use gradient-based optimization, which can limit the scalability of this approach. Other methods that exploit submodular properties of some OT problems [45], [57], which enable fast approximate optimization with convergence guarantees.

3 LAUE ANALYSIS

In this section we provide a pedagogical introduction to crystallographic lattices, Laue diffraction, and tomography experiments.

3.1 Crystallographic planes

In a single crystal, the atoms are arranged in a specific pattern, described by the crystallographic point group. The lattice of the crystal creates a set of Miller indices, determined by integer triples $[h, k, l]$. They denote a family of parallel lattice planes, sharing the distance between atoms a and normal vector u . Given three lattice vectors a_1, a_2, a_3 that define a unit cell, $[h, k, l]$ denotes planes that intercept the three points $[ha_1, ka_2, la_3]$ or their multiples (for zero indices, the intercept is at infinity and the planes do not intersect the axis). For cubic crystals with lattice constant a , the spacing d between adjacent $[h, k, l]$ lattice planes is




Fig. 1. Bragg back-scatter diffraction from a crystal lattice with cubic symmetry (pure iron). Showing a cross-section through the z -dimension. Example Miller planes with normal vector u and indices $[h, k, l] = [1, 2, 0]$ are shown with the red lines. The incoming beam with wavelength λ is shown with the colored line, with color corresponding to the value of the light sine wave. The diffracted ray r displays constructive interference with aligned phases, shown with the gray line.

$d = a / \sqrt{h^2 + k^2 + l^2}$. In this work we will use the definition of Miller indices as the inverse intercepts along the lattice vectors. It is also common to use Miller indices defined as points in the reciprocal lattice, which we do not use here. See [58] for a comprehensive introduction to crystallography.

3.2 Laue diffraction

Laue-Bragg interference describes the scattering of waves from a crystal lattice. The constructive interference of diffracted waves occurs for specific combinations of the incidence angle θ between the plane and the incoming beam, the wavelength λ , and the distance d between atoms. Bragg's law states that constructive interference occurs when

$$n\lambda = 2d \sin \theta, \quad (1)$$

where the integer n is the diffraction order. For an incoming beam with unit direction e , the wave diffracted by plane with normal unit vector u will have the unit direction r , which can be calculated using the Householder equation, and the wavelength λ at which the scattering occurs

$$r = \pm(\mathbb{I} - 2uu^\top)e, \quad (2)$$

$$\lambda = 2du^\top e, \quad (3)$$

with “-” for transmission (forward-scattering) and “+” for reflection (back-scattering). The intensity of the ray r scales with wavelength λ^4 , the incoming beam intensity at λ and the volume of the crystal. See Figure 1 for an illustration of Bragg diffraction for a single plane with $[h, k, l] = [1, 2, 0]$.

Let's consider a laboratory coordinate system with the crystal at its origin. While each crystallographic plane in 3D has some small offset with respect to the center of the lattice, we will use a common position $x \in \mathbb{R}^3$ in the laboratory coordinate system for all planes belonging to the crystal, which will correspond to its center. We use the incoming beam with direction $e = (1, 0, 0)$, and an example detector

centered at γ , with surface normal vector ν . For a particular $[h, k, l]$ Miller planes defining the diffracted direction r , the ray will create a Bragg peak (also called a *spot*) on the detector, at the position

$$s = x + \frac{\nu^\top(\gamma - x)}{\nu^\top r} r, \quad (4)$$

where $(\nu^\top(\gamma - x))/(\nu^\top r)$ is the distance from the center of the grain to the (3D) position of the spot in the detector.

3.3 Tomography experiments

Laue tomography experiments aim to study the internal structure of polycrystalline samples. A polycrystalline sample is composed of multiple single crystals, or *grains*. Each grain is described by its position x inside the sample and orientation matrix $R \in SO(3)$ with respect to the reference $R^{\text{ref}} = \mathbb{I}$. We consider a sample to consist of a set of grains $\{\mathcal{G}_n\}_{n \in N}$, each characterized by its position and orientation matrix $\mathcal{G}_n = (R_n, x_n)$. Each grain will contribute a set of Miller planes, which we will consider to be located at the grain's center x , as described above. For a given experiment, we can limit the number of considered Miller planes, as higher order planes will not satisfy Bragg's condition because the associated wavelength is too short or too long. We will use a set $\{\mathcal{M}_m\}_{m \in M}$ of plausible Miller planes, each with a corresponding unit direction w and atom spacing d , creating a couple $\mathcal{M}_m = (w_m, d_m)$. This plausible set is provided as an input to the analysis.

The positions and orientations of the crystal are found by recording locations of spots created by diffracted rays on the detector screen after illuminating the sample by a wide-spectrum beam. We will consider the beam to have the wavelength λ in range $\lambda \in [\lambda^{\min}, \lambda^{\max}]$. During a tomographic experiment, the sample is illuminated multiple times after being rotated. Let $\mathcal{T} = \{\Gamma_t\}_{t \in T}$ be a set of rotation steps, where $\Gamma_t \in SO(3)$ is the rotation matrix for step t . The formulations that we will develop are applicable to any rotation, but the experimental setup that we consider allows in practice only rotations along the (0,0,1) axis. Figure 2 shows the scene in question, with two detectors, recording back-scatter or forward-scatter rays. The thick grey arrow shows the incoming beam, the thin dashed line corresponds to the axis of rotation for tomographic projections. A single plane shown in dark grey, while the light grey shows the studied sample in which the plane is located. In general, more detectors can be included in the experiment. In the interest of clarity, we will from now on consider only the backscatter detector and rays. An equivalent analysis can be easily performed for the forward-scattering mode.

At a given sample rotation Γ_t and for grain orientation R_n , the vector u_{mnt} normal to the rotated Miller plane w_m will be

$$u_{mnt} = \Gamma_t R_n w_m, \quad (5)$$

the model ray is emitted in the direction

$$r_{mnt} = (\mathbb{I} - 2\Gamma_t R_n w_m w_m^\top R_n^\top \Gamma_t^\top) e, \quad (6)$$

creates a spot on the detector plane at position s_{mnt} as follows

$$s_{mnt} = \Gamma_t x_n + \frac{\nu^\top(\gamma - \Gamma_t x_n)}{\nu^\top r_{mnt}} r_{mnt}, \quad (7)$$




Fig. 2. A scene of a tomographic diffraction experiment with two detectors configured for recording backscatter and transmission spots. The thick gray arrow corresponds to the incoming beam. The beam direction is set by vector e . The light gray cube is the studied sample. The dark gray plane is an example Miller planes in a grain. The plane is oriented inside the grain with normal vector u . The constructive interference will occur at directions r^- (transmission) and r^+ (backscatter) and create spots s^- and s^+ , respectively, on the detector screens. For clarity, the detector size used here is smaller than in a typical experiment.

and at wavelength λ_{mnt} with

$$\lambda_{mnt} = 2d_m(\Gamma_t R_n w_m)^\top e. \quad (8)$$

3.3.1 Detected spots

We assume that the center γ of a square backscatter detector lies on the x-axis of the laboratory coordinate system, with size l mm per side. We will assume that γ and l are known, as they can be characterized accurately with a proper calibration process [38]. The detection of spots and the measurement of their positions is performed by an external algorithm and is considered as input to this analysis. Next to real Bragg peaks, the set of detected spots will also contain spurious detections, which we discuss in Section 7.2. An experiment will yield a collection of **detected spots** $\{p_i\}_{i \in I}$, in all experimental images obtained for different rotations of the sample. We will denote by t_i the rotation step which produced the i th spot.

We assume a ray noise model, where a noisy unit ray direction follows the von Mises-Fisher distribution in 3 dimensions [59], with mean r and concentration κ

$$\mathbf{vMF}_3(y | r, \kappa) = \frac{\kappa}{4\pi \sinh \kappa} \exp(\kappa \cdot r^\top y).$$

Note that samples from this distribution can be obtained by drawing from a 3D isotropic multivariate normal distribution with covariance $C = \sigma^2 \mathbb{I}$, where $\sigma^2 = 1/\kappa$, and then conditioning $\|x\| = 1$. Therefore the noisy ray \tilde{r}_i is

$$\tilde{r}_i \sim \mathbf{vMF}(r_i, \sigma^{-2}), \quad (9)$$

and the detected spots will be offset from their true positions s_i following the noisy ray \tilde{r}_i

$$s_i = \Gamma_{t_i} x_{n_i} + \frac{\nu^\top(\gamma - \Gamma_{t_i} x_{n_i})}{\nu^\top \tilde{r}_i} \tilde{r}_i. \quad (10)$$

3.3.2 Model spots

In our method, we aim to fit a dataset of detected spots with a set of **model spots** stemming from some candidate model sample in our experimental setup. Not all of the model spots s_{mnt} can be considered in the problem, for three reasons. Firstly, the model spot can lie outside the detector area. Secondly, the brightness of the Bragg peak will be proportional to the intensity of the beam at the corresponding wavelength λ ; if the intensity is close to zero, then the spot will not be detected. Finally, the intensity of the spot will be proportional to the volume of the grain; small grains will yield low-intensity spots that will not be detectable.

We therefore use a sequence of model spots $\mathcal{S}=(s_j)_{j \in J}$, that contains only those spots that simultaneously lie inside the beam wavelength range and the detector screen area

$$\mathcal{S} = \{s_{mnt} : (m, n, t) \in M \times N \times T \\ \wedge \lambda_{mnt} \in [\lambda_{\min}, \lambda_{\max}] \wedge \|s_{mnt} - \gamma\|_{\infty} < l/2\}. \quad (11)$$

We will use a sequence $(r_j)_{j \in J}$ to denote rays corresponding to spots in \mathcal{S} . For clarity, we will use a sequence $(t_j)_{j \in J}$ of indices $t_j \in T$ to label the step at which spot j was generated, itself labeled by index set J_t . Equivalently, let sequence $(n_j)_{j \in J}$ of indices $n_j \in N$ denote the index of the grain at which spot j was created, itself labeled by index set J_n . Thus, $|J_n|$ and $|J_t|$, respectively, will be the number of model spots generated by grain n or step t , meeting the criteria in Equation 11. Finally, let m_j denote the index of the Miller plane which refracted the ray producing spot j and J_m the collection of spot indices associated with a particular Miller plane m .

4 THE LIKELIHOOD

In this section, we propose a probabilistic model for the collection of observed spots detected given a collection of model spots produced by a set of grains N producing each a refraction for a collection of Miller planes M , with latent variables encoding which detected spot matches which model spot. We assume the experimental setup defined above with detector described by (γ, ν, l) , with sample rotation step in T yielding detected spots indexed by I . First, we introduce the problem without outlier spots, followed by introduction of the outliers.

Let's consider a guess grain with position x_n that generates the ray j . For a detected spot s_i is matched with model ray j , we define a ray estimate \hat{r}_{ij} as

$$\hat{r}_{ij} = \frac{s_i - \Gamma_{t_j} x_{n_j}}{\|s_i - \Gamma_{t_j} x_{n_j}\|}. \quad (12)$$

This quantity will be useful in further calculations.

4.1 Spot-to-spot assignment

We assume that there are no outliers among detected spots and no missing model spots (corresponding to unmodeled Miller planes), so that there is a one-to-one correspondence between model and observed spots. To match the detected spots with model rays, we introduce the assignment variables $Z_{ij} \in \{0, 1\} : i \in I, j \in J$, with $Z_{ij}=1$ if the

detected spot s_i is associated to model ray r_j , and $Z_{ij}=0$ otherwise. Only matches between spots occurring at the same sample rotation t are allowed, otherwise $Z_{ij}=0$. Each i is matched exactly to one j , so the joint log-likelihood $\ell := \sum_{i=1}^n \log p(s_i | Z_{i\cdot}, s)$ for all observed spots s_i given $Z_{i\cdot} := (Z_{ij})_{j \in J}$ and r_j as

$$\ell = \frac{1}{\sigma^2} \sum_{i \in I} \sum_{j \in J} Z_{ij} \cdot \hat{r}_{ij}^\top r_j - c. \quad (13)$$

where c is the log normalizing constant. At step t , the number of models spots and detected peaks is $|J_t|$ and $|I_t|$, respectively. Note that if we assume a uniform prior probability for the assignment of i to all j s such that $t_i = t_j$, and with $Z := (Z_{ij})_{i \in I, j \in J}$, we have

$$p(Z) = \prod_t \prod_{(i,j): t_i=t_j=t} |J_t|^{-Z_{ij}} = |J_t|^{-|I_t|}, \quad (14)$$

the joint log-likelihood only differs from ℓ by $-\sum_t |I_t| \log |J_t|$, which just changes the normalizing constant c . For clarity, we denote the spot assignment likelihood matrix as $\tilde{L} \in \mathbb{R}^{|I| \times |J|}$. It contains the cost of assigning observed spot i to model spot j

$$L_{ij} = \frac{1}{\sigma^2} \hat{r}_{ij}^\top r_j \quad i \in I, j \in J. \quad (15)$$

Since the assignment of peaks to models spots is a-priori unknown, treating the variables Z_{ij} as latent variables and maximizing the marginal log-likelihood $\sum_i \log p(s_i | (R_n, x_n)_{n \in N})$ can be done with a classical EM-algorithm, which maximizes the evidence lower bound (ELBO)

$$\mathcal{L}(Q; (R_n, x_n)_n) = \sum_{i \in I} \sum_{j \in J} Q_{ij} L_{ij} - \text{const.}, \quad (16)$$

where $Q \in \mathbb{R}_+^{|I| \times |J|}$ is the responsibility matrix with elements $Q_{ij} = \mathbb{E}[Z_{ij}]$. Instead of the classical EM approach, we propose to use the Sinkhorn-EM (sEM) formulation [43], which was recently introduced for solving mixture model problems. While the classical EM method calculates the cluster memberships independently, sEM computes them using optimal transport, where the responsibilities respect the known proportions. It is shown that sEM displays better global convergence guarantees, while optimizing the lower bound on the log-likelihood and thus maintaining the probabilistic interpretation of the analysis. In sEM, the likelihood is replaced with an entropic-OT likelihood

$$\mathcal{L}_{\text{OT}}(Q; (R_n, x_n)_n) = \sup_{Q \in \Pi(a,b)} \left[\sum_{i \in I} \sum_{j \in J} Q_{ij} L_{ij} + \mathbf{H}(Q) \right], \quad (17)$$

where $\Pi(a, b)$ is a set of all transport plans that have joint distributions with marginals a and b , respectively, and $\mathbf{H}(Q) = -\sum_{i,j} Q_{ij} \log Q_{ij}$ is the entropy of Q ¹. In a discrete setting applicable here, the marginal mass distributions a and b have entries uniformly equal to 1: for each detected and model spot: $a_i = b_j = 1$. This formulation requires solving an optimal transport problem. This likelihood is consistent

1. Here Q refers to the transport plan, not the entropy of the distribution over permutation with mean Q .

in the population limit with the classical likelihood approach while having a better geometrical properties and is proven to be less prone to getting stuck in local optima than classical EM. Using this formulation enables practical solving of the large-scale Laue problem.

Note that Z_{ij} and Q_{ij} can only be non-zero if $t_i=t_j$; in the rest of the paper, we will consider the set \mathcal{Q} of matrices Q such that $Q_{ij}=0$ for all (i, j) with $t_i \neq t_j$.

4.2 Assignments with outlier and missing spots

In this section we consider outlier spots, both in detected and modeled sets. There are two reasons for outliers among spots detected in image data. Firstly, the outliers can be just spurious detections: noise spikes or points corresponding to detection algorithm failures. Blended overlapping spots can also have dramatically wrong position measurement, which lies outside the range allowed by the noise level. Secondly, the sample can contain very small grains that will give rise to only few detectable spots, which can be insufficient to form a diffraction pattern that can be reliably measured by the solver. The number of outliers in the data can be large and it is typically not known by the user.

We also consider unmatched model rays. The area of the screen and wavelength range are limited, which leads to spots moving out and into the image as the grain is rotated. We start with a set of initial (or *prototype*) grains, giving rise to a set of model spots \mathcal{S} , with spots outside the detector screen or outside the beam's wavelength range being excluded (Equation 11). However, as the orientation and position of the grains are optimized, some spots may exit the detector/wavelength range and should be also excluded during the matching process. Conversely, there will be spots that were initially excluded, but would appear for the updated grain parameters. Precise modeling of this property in the likelihood would be complicated. Therefore, we propose a simpler approach, where we treat these cases as unmatched model spots instead. During the optimization we will use the same set of model spots as predicted using for the initial values of the prototypes grains with $(R_n^0, x_n^0)_{n \in N}$. The closer the prototype model parameters are to their true value, the fewer unmatched spots will be encountered. According to our empirical tests for the number of prototypes considered here, the fraction of different between the candidate models and their true counterparts is no more than 10%.

We model the outliers using a full probabilistic description of the problem. We add an extra row to Z , with index 0, reserved for the assignment of spurious detection to the outliers status, i.e. $Z_{i0} = 1$ for outlier spots. Symmetrically, we add an extra column, again with index 0, to assign unmatched model spots, i.e. $Z_{0j} = 1$ if a model spot does not match any detected peak. Note that multiple entries of the 0th row (resp. column) can be non-zero. At the intersection of these rows the entry Z_{00} is ignored. To take into account explicitly the possibility of outliers into the probabilistic model, we modify the likelihood as follows. We assume there is first a fixed probability π_0 for a detected spot to be an outlier, so that $p(Z_{i0}=1)=\pi_0$ and $p(Z_{ij}=1)=(1-\pi_0)/|J|$. Given that it is an outlier, we assume that it appears uniformly over the detector so that $p(s_i|Z_{i0}=1) = \frac{1}{S}$ where

S is the surface of the detector, if it is not an outlier, we use the same model as before. The joint likelihood of a detected spot position s_i and its assignment variable Z_i can then be written as

$$p(s_i, Z_i) = \left(\frac{\pi_0}{S}\right)^{Z_{i0}} \prod_{j=1}^J \left(\frac{1}{c} \exp L_{ij}\right)^{Z_{ij}} \left(\frac{1-\pi_0}{J}\right)^{Z_{ij}}. \quad (18)$$

Each term of the log-likelihood can be written as

$$\ell(s_i, Z_{ij}) = Z_{i0} \log \frac{\pi_0}{S} + (1 - Z_{i0}) \log \frac{1-\pi_0}{Jc} + L_{ij} Z_{ij}.$$

Given that our likelihood is constructed per detected spot, and that we do not model their number, the probability that a model spot is unmatched is modeled implicitly, given the total number of model spots. We do not explicitly model the probability of unmatched model spots in our formulation of the problem, and set it to likelihood value of the von Mises-Fisher distribution with concentration $\kappa=\sigma^{-2}$ that corresponds to the null angle difference. Note that, up to additive constants, the term for a pair (i, j) with $i, j \neq 0$ in the log-likelihood remains as before. We define the extended likelihood matrix

$$\begin{aligned} L_{0j} &= \sigma^{-2} - \log(4\pi\sigma^2 \sinh(\sigma^{-2})) \\ L_{i0} &= \log \left(\frac{\pi_0}{1-\pi_0} \cdot \frac{|J|c}{S} \right) =: L^{\text{out}} \\ L_{ij} &= \frac{1}{\sigma^2} \hat{r}_{ij}^\top r_j \quad \forall i \in I, j \in J. \end{aligned} \quad (19)$$

As in practice it is difficult to estimate π_0 , in our implementation we use a fixed outlier likelihood L^{out} , as defined in Equation 19. We set $L_{00}=L^{\text{out}}$; the corresponding assignment Q_{00} will serve as a sink for unused outlier capacity. The complete log-likelihood of the data and associated assignments take the form $\langle L, Z \rangle - \text{const.} + \log \left(\frac{1-\pi_0}{Jc} \right)$, where now both L and Z have an extra row and column. With constants dropped, the Sinkhorn EM objective is

$$\mathcal{L}_{\text{OT}}(Q; (R_n, x_n)_{n \in N}) = \langle L, Q \rangle + \epsilon \mathbf{H}(Q) \quad (20)$$

where $Q \in \Pi(a, b)$ is the entropic transport plan, a and b are marginal constraints (see Section 4.3), and ϵ is the entropic regularization parameter, which is set to $\epsilon=1$ in standard sEM (also see Section 5.1). In literature, there are other optimal transport problem variants that allow for outlier treatment, such as unbalanced OT [46] and partial OT [60]. We compare the performance of our proposed outlier model with these other variants. See Appendix C for details.

4.3 E-step with Sinkhorn updates

The E-step for finding the assignment Q now requires to solve

$$\begin{aligned} \max_Q \quad & \langle L, Q \rangle + \epsilon \mathbf{H}(Q) \\ \text{s.t.} \quad & Q \geq 0, \quad Q \in \mathcal{Q}, \quad Q \mathbb{1} = a, \quad Q^\top \mathbb{1} = b, \end{aligned} \quad (21)$$

which follows from Equation 20, and ϵ is the entropic regularization parameter (also see Section 5.1). The maximum number of allowed outliers can be encoded in vectors a and b , with T_a and T_b corresponding to the maximum number of spurious detections and unmatched model spots, respectively. We set $a_i=b_j=1$, $i, j > 1$ for matched spots, so

the weight of the outliers is $a_0=T_a$ and $b_0=T_b$. In order for the double-stochastic constraint in Equation 20 to be satisfied, we use T_b as the user-controlled variable and set the $T_a = T_b - b^\top \mathbb{1} + a^\top \mathbb{1}$. The transport plan matrix is updated using the Sinkhorn subroutine $Q \leftarrow \text{BalancedOT}[L, \epsilon, \zeta^{\text{tol}}]$ shown in Algorithm 4, described in Appendix 4. The convergence of this algorithm is controlled by user-specified parameter ζ^{tol} .

4.4 The M-step updates

The M-step finds the maximum likelihood orientations and positions for all grains $(R_n, x_n)_{n \in N}$ by solving a separate problem for each grain

$$\forall n \in N : \max_{x_n, R_n : R_n^\top R_n = \mathbb{I}} \langle L, Q \rangle \quad (22)$$

We solve this problem by coordinate descent in x and R . We propose to find grain position x using Newton iterations. We calculate the update at step k as

$$\Delta x^{k+1} = \left(H(x^k) + \iota \mathbb{I} \right)^{-1} g(x^k), \quad (23)$$

where $H(x^k)$ and $g(x^k)$ are gradient and Hessian matrices evaluated at the current position x^k , and ι is a very small constant added for numerical stability. We show the calculation of these matrices in Appendix B. The update at iteration k is

$$x^{k+1} \leftarrow x^k - \Delta x^k. \quad (24)$$

The iterations are repeated until convergence threshold ζ^x is met: $\|x^{k+1} - x^k\| < \zeta^x$. Subroutine **LauePosNewton** (Algorithm 1, described in Appendix B), shows the details of this process.

Finding the M-step update for the orientation matrix is more complicated. In Appendix B we derive a Majorization-Minimization algorithm for finding the orientation matrix R maximizing the loss \mathcal{L}_{OT} . This sequence starts from an initial rotation guess R^0 . We show that the maximization of the objective is equivalent to the maximization of a concave quadratic, for which a local maximum can be obtained by solving a sequence of linear optimization problems over rotation matrices that each take the form of a *Wahba's problem* (an alternative formulation of the *Orthogonal Procrustes problem*, see Appendix A for details). First, we calculate the matrix M_m^0 using the current values of grain positions x_n and assignments Q . We then shift it so that it becomes PSD by adding a constant diagonal $\iota \mathbb{I}$

$$M_m^0 := - \sum_{j \in J_n \cap J_m} \Gamma_{t_j}^\top \left[\sum_{i \in I} Q_{ij} (\hat{r}_{ij} e^\top + e \hat{r}_{ij}^\top) \right] \Gamma_{t_j}, \quad (25)$$

$$M_m := M_m^0 + \iota \mathbb{I} \quad (26)$$

where w_m is the direction of ray m and M_m , and ι is the absolute value of the most negative eigenvalue of M_m^0 or 0 if M_m^0 does not have negative eigenvalues. The update at iteration k is

$$R^{k+1} \leftarrow \arg \min_{R : R^\top R = \mathbb{I}} \sum_m \|R w_m - M_m R^k w_m\|^2. \quad (27)$$

The iterations terminate when the convergence threshold ζ^R is met: $\|R^{k+1} - R^k\|_F < \zeta^R$. Subroutine **LaueRotMM** (Algorithm 3, described in Appendix B) presents this method in detail.

5 THE LAUEOT SOLVER

To efficiently solve the stated problem at scale, we propose the LAUEOT formulation (Laue indexing with Optimal Transport). We introduce key elements of this method below.

The LAUEOT method uses iterative updates for all variables. This enables very efficient implementation of coordinate ascent steps on GPUs, as all of these operations require only simple tensorized algebra.

As the number of local maxima in this problem can be large, techniques for global optimization can be employed. In this work we use a deterministic annealing scheme to control the “temperature” of the problem [52] [61]. It enables more efficient search for local maxima. During the final steps of the algorithm the temperature control is switched off, which corresponds to finding the local maximum likelihood in Equation 20. This process is described in Section 5.1.

In a typical Laue analysis, the number of spots in the problem can be large, up to several millions. Additionally, only matches of spots at the same sample rotation t are allowed. To enable good performance of the method at scale, we work with a sparse representation of the assignment matrix (see Section 5.2).

The multi-grain fitting step is initialized using a set of prototype grains, which together constitute the final sample. Therefore, finding a good set of prototype grains is a critical step. We create the set of prototype grains using two pre-processing steps: (i) single-grain fitting, where we find local minima for samples consisting only of a single grain, starting from multiple initialization points in rotation and position space, and (ii) prototype selection, where we select a small number of prototypes out of all candidates, that form the final sample. We describe the pre-processing steps in Section 6.

5.1 Deterministic annealing

The problem in Equation 20 is non-convex and has multiple local minima. This is a common problem for EM algorithms. Deterministic annealing variant of EM has been proposed to address this [62]. In this variant, the likelihood function of data x and latent y given parameters Θ is replaced by another function $p(x, y | \Theta) \rightarrow p(x, y | \Theta)^\beta$, where $\epsilon=1/\beta$ is the “temperature”. The temperature starts at a value $\epsilon \gg 1$ and is gradually lowered during the EM iterations. At $\epsilon=\beta=1$, the original EM formulation is recovered. Gradual temperature decrease helps with avoiding local minima. Similar scheme that additionally includes outlier variables was used in [63]. In the Laue problem, using temperature $\epsilon > 1$ is analogous to increasing entropic regularization in optimal transport. Starting with large ϵ allows for many rays to explain a given detected spot. Decreasing ϵ allows to lift gradually the assignment ambiguities. The initial temperature ϵ^{init} is typically chosen to be $\epsilon^{\text{init}} > 100$.

5.2 Working with sparse assignment matrices

Considering only ρ nearest neighbors in the assignment scheme, which is equivalent to setting null assignment weight in the full Q matrix, has the strong advantage to decrease the complexity of Sinkhorn updates from J^2 to ρJ . In the Sinkhorn algorithm, this will be a good approximation




Fig. 3. Illustration of three selected steps of the LAUEOT solver on a toy problem with a simulated ruby-like sample with two grains and noise-free detected spots, for a single projection in the backscatter detector. The left panels shows the initialization of two grains with red and blue model spots, and strength of the initial assignment weights in grey lines. The darker the color of the line, the stronger the assignment weight. The algorithm starts with a high temperature, which enables assignments between model and detected spots that can be far. As the optimization progresses, the temperature is lowered, as shown in the middle panel. Finally, the algorithm assigns the model and detected spots to the grains. Note that some of the detected spots were not matched; those will be considered as outliers. Conversely, some model spots are unmatched.

as long as $\exp(L/\epsilon)$ is close to zero for all points outside the neighbor set. We find this approximation to be stable enough for our implementation, although other approaches could be employed here [55], [64].

5.3 The algorithm

The LAUEOT method is described in Algorithm 9 (Appendix G). Next to the detected spots and algorithm control parameters, the routine takes as input a set of prototype grains $(R_n, x_n)_{n \in N}$, which is calculated in the pre-processing steps. A simple visual demonstration of the algorithm is shown in Figure 3. It shows two candidate grains (blue and red spots) being fit to a set of detected spots (black), at the initial optimization step (left panel), step 50 (middle panel), and the final assignment (right panel). The gray lines show the strength of the assignment Q . Note the outliers (black circles) and unmatched model spots (colored circles) in the final assignment.

6 PRE-PROCESSING

We perform two preprocessing steps:

- (i) single-grain fitting, where we fit multiple candidate single grains from different starting points in rotation and position space,
- (ii) selection of prototypes from the list of optimized candidate grains using the *Selection of Prototypes with Optimal Transport* (SPOT) framework [45].

The goal of the first step is to find single grains that fit the data well and can be included in the final grain set. The second step selects the final set of prototype grains that will be included in the sample and considered for multi-grain optimization.

6.1 Single-grain fitting

The first pre-processing step is a single grain search over the space of its orientations and positions. The necessity for this step stems from the fact that, even for a single grain, the problem in Equation 20 has multiple local minima in the x and R variables.

Generally, the problem of finding orientation matrices to match point sets is multi-modal for data without correspondences [51], [61] and is typically solved by running convex solvers multiple times from different starting points in $SO(3)$. We note that, for many other computer vision problems, a good set of starting points can be pre-computed using simple alignment of principal components axes in source and target sets. Diffraction patterns, however, do not have a specific “shape”, which prohibits this approach. As a pre-processing step to find a large number of candidate grains, we use single-grain fitting runs, initiated a different starting points in the six-dimensional (R, x) space (see Section 6.1). We discretize the search space using a Sobol sequence grid, where the grain orientation is represented using the *modified Rodrigues parameters* [65]. The search over R is restricted according to the crystallographic symmetries of the lattice [66], while the search over x is limited to points within the bounds of sample size. These limits are set by the user depending of the type of lattice corresponding to the composition of the sample, as well as its physical size.

From each initialization point, we run a single grain fitting using a single-grain variant of the problem, described in Appendix D. This formulation will treat most of the detected spots as outliers. The number of runs is variable, depending on the expected number of grains in the sample. For 10 and 1000 grains, we run 5k and 500k initializations, respectively. Once the optimization is finished, we store the converged grain orientations and positions. Subroutine **SingleLaueOT** (Algorithm 8, described in Appendix D) is used for this task.

6.2 Selection of Prototypes with Optimal Transport

The goal of this step is to select the grain candidates that will compose the final sample. As no further grains will be added to the sample later, this is a critical part of the algorithm. Selecting too many or too few grain candidates can lead to errors in the multi-grain fitting solutions. We start with removing badly-fitted models by selecting grains that have the fraction of unmatched model rays greater than f^{out} . A model ray is considered unmatched if its largest weight in Q is assigned to the outlier column. Then to create the final sample, we perform prototype selection in the framework *Selection of Prototypes using Optimal Transport* (SPOT) [45].

The SPOT framework casts the problem of selection prototypes as a problem of learning a sparse source probability distribution supported on k elements that has the minimum OT distance from the target distribution. It poses the prototype selection problem as a maximization version of the OT problem, where the source distribution weights are learned with a cardinality constraint. The proposed loss has a key property of submodularity, which enables the use of efficient greedy method with deterministic approximation guarantees. We use the SPOTGREEDY method from [45], where we pose the problem as grain-to-spot OT problem of model grains to detected spots. Algorithm 7, described in Appendix E, shows the subroutine **LaueSpotGreedy** which is an adaptation of this method for our problem. In this algorithm, we also determine the final number of prototypes N : we find N after which the improvement in the total cost of transport significantly decreases (see Appendix E for details).

We note that this algorithm allows for assignment of a larger number of detected spots to the grain than its number of model spots. Enforcing that constraint leads to a more complicated optimization problem of support selection, recently proposed by [67]. The problem formulation presented there does not include the possibility of target sample outliers, as well as poses a difficult challenge for optimization. As we find that the SPOT method works well in practice, we leave more further investigation in this area to future work.

7 SIMULATIONS AND EXPERIMENTS

We evaluate our proposed LAUEOT algorithm according to the following performance criteria

- 1) accuracy of the solution as a function of number of grains in the sample and the noise level,
- 2) sensitivity of the solutions to the fraction of outliers for different optimal transport formulations, and its dependence on OT outlier control parameters,
- 3) number of successfully indexed grains as a function of the number of initial single-grain optimizations for two types of lattice,
- 4) improvement of the multi-grain solver solution over the single-grain fitting solutions,
- 5) comparison of time-to-solution between LAUEOT and the LAUE3DNDT method from [22],
- 6) comparison of the indexing results on spots obtained from experimental images.

For simulated data, we use two types of lattices, FeNiMn and CoNiGa, which are metal alloys (see [68] for introduction to metal alloys). We test the accuracy of the LAUEOT solutions using simulated data in the following way. We consider a neutron scattering experiment scenario with the wavelength of the beam $\lambda \in [0.6, 6]$ Å and two detectors placed at $\gamma^+ = (-160, 0, 0)$ and $\gamma^- = (160, 0, 0)$ mm. The fiducial experiment configuration for both samples is 12 projections ($\omega = 0^\circ - 360^\circ, \Delta\omega = 30^\circ$). The noise level for the fiducial configuration corresponds to the scatter of $\sigma = 0.4$ mm on the center of the detector screen, which is equivalent to angular spread of $\sigma = 0.143$ deg for rays. The chosen noise level is already conservative, as the size of the point spread function on the scintillator detectors, such as the ones used in neutron scattering studies, is ~ 0.2 mm, and its pixel size is 0.103 mm. We add outliers to the simulated detected spot set. The position of the outliers is drawn from an uniform distribution on the image plane. The number of outliers in the fiducial configuration is set to 10% of the Laue spot count. In the fiducial configuration, the simulated samples contain up to 1000 grains of orientations chosen on a Sobol sequence with additional random scatter.

For synthetic data, we simulated two datasets using a $Fm\bar{3}m$ space group and a $Pm\bar{3}m$ space group with lattice parameters $a=b=c=3.592$ Å and $a=b=c=2.872$ Å, respectively. Hereafter, the simulated samples will be referred to as SynthSampleA and SynthSampleB, respectively. The choice of the above space groups and corresponding lattice parameters was made because they match recent experimental studies [23], [69]. For the experimental data, we used a ruby single crystal, typically used in crystallography for calibration measurements, a CoNiGa oligocrystalline sample, and a FeGa polycrystalline sample. Sample properties, including space groups and lattice parameters, are listed in Appendix F.1.

Use 32 nearest neighbors in the sparse formulation. The initial temperature ϵ^{init} is related to the number of nearest neighbors considered in the sparse formulation: it is chosen so that the farthest neighbor has a non-negligible probability at start. The temperature is lowered by $\epsilon^{\text{cool}}=0.97$ at every iteration. For the likelihood of outliers, we use the log likelihood L^{out} set to the likelihood corresponding to a angular difference of 3σ . To discard badly fitting grains, we require the number of rays matched to the outlier to be less than $f^{\text{out}} \cdot |\mathcal{S}_n|$, with typical value $f^{\text{out}}=0.1$. For the multi-grain fitting, we set the maximum number of unmatched model rays $T_b = \lceil 0.4 \cdot |\mathcal{S}_n| \rceil$ of total number of rays.

7.1 Quality metrics

To assess the quality of the LAUEOT solutions, we compare: (i) the fraction of correctly assigned spots across all grains, (ii) the median difference between true and estimated grain orientation, (iii) the median difference between true and estimated grain position. The results of these comparisons are shown in Figure 4, as a function of number of grains and spot position noise level.

For samples with up to 500 grains, the number of recovered grains was exactly correct. After that, there was at most 1 grain missing. Starting with samples with 700 grains and for higher noise levels, the method begins to




Fig. 4. Quality metrics for indexing with LAUEOT using simulated data with variable number of grains and noise levels, for the constant fiducial outlier fraction of 10%. The left panel shows the fraction of correctly assigned spots in the entire problem. The middle and right panel show the scaled median absolute deviation of the error on the grain orientation and position, respectively. The number of grains was correctly identified for up to 500 grains, and after that there was up to 1 grain missing. The method fails for large noise levels and high grain counts.




Fig. 5. Quality metrics, as in Figure 4, for the multi-grain fitting solution of the full LAUEOT solver, compared to the single-grain fitting initializations. The comparison is done for a fixed number of 500 input grains and the fiducial outlier fraction. The LAUEOT solver finds all grains successfully, while the single-grain fitting results in ~ 940 correctly indexed grains, with slight dependence on the noise level.

fail. This gives an idea for the limitations of LAUEOT. The left panel shows the fraction of wrongly assigned spots as a function of the number of grains and the noise level. LAUEOT mismatches only up to 15% of spots even for the challenging scenario of 500 grains and high noise level $\sigma=0.251^\circ$ (0.7 mm on the image plane). For the low noise level of $\sigma=0.036$ deg, LAUEOT assigns almost all spots correctly. As expected, this fraction decreases with increased noise level. Middle and left panels show the median error of the recovered grain orientation and position, respectively. Predictably, these errors increase with the noise level. They stay almost constant with increased number of grains, which is expected in case of good quality of assignments.

Figure 5 shows the benefit of obtaining solutions by solving the full inverse problem in Equation 20 with OT, compared to using the prototypes from single-grain fitting only. It shows the quality metrics, as in Figure 4, for the SynthSampleA with 500 grains, at varying noise levels and fixed outlier fraction of 10%. To assess the performance for the single-grain fitting, a model spot is hard-assigned to the nearest detected spot. Rays were marked as unmatched

when the probability of the match was lower than the vMF likelihood L^{out} corresponding to the angle of 3σ . The OT formulation decreases the fraction of wrongly assigned spots by a factor of 2. The grain orientation and position errors also improve with the multi-grain fitting, with errors decreasing by up to 25% for high noise levels.

7.2 Robustness to outliers

We test the sensitivity of the LAUEOT method to outliers in the detected spots. We compare the performance of our proposed LAUEOT formulation using probabilistic outlier model (see Section 4.2), as well as two other OT formulations: partial OT and unbalanced OT (see Appendix C). In this ablation study we explore how this choice affects the indexing results. In particular, we test the sensitivity of the results to the choices of outlier control parameters: likelihood associated with outliers, maximum number of unmatched rays for OT with outlier modelling, κ and λ entropic regularization parameters for the unbalanced OT, and fraction of transported mass m_{OT} for the partial OT.




Fig. 6. Performance of the LAUEOT solver for number of outliers in the data, for three outlier treatment methods used. The hyperparameters for each method were found by performing a grid search over multiple combinations and choosing the one that minimizes the average assignment error over outlier fractions.

We choose simulated diffraction images of the SynthSampleA with noise level $\sigma=0.143^\circ$ as the reference. We vary the fraction of outliers up to 70%. For LAUEOT, we use combinations of outlier cost L^{out} and maximum number of unmatched rays T_b . The outlier likelihood corresponds to vMF probability for angular difference of 2–5 σ , while the fraction of unmatched rays between 1–40%. For UNBALANCEDOT, we set $\kappa, \lambda \in [0.001, 1]$. For PARTIALOT, the fraction of transported mass is varied between $m_{\text{OT}} \in [0.4, 0.95]$. For each OT method, we choose the parameter configuration that yields the smallest average assignment error over all outlier fractions.

Figure 6 shows the assignment error (left panel), grain orientation error (middle panel), and position error (right panel), for the three OT outlier treatment methods. The results indicate that the methods are not affected by outliers. The best overall results are obtained with the custom outlier modelling in LAUEOT. The unbalanced and partial OT give slightly worse, but comparable results. Given that the multi-grain fitting is very fast, the user can perform a coarse grid search to find reasonable values for the hyper-parameters.

7.3 Number of initializations and runtime

Figure 7 shows the number of correctly indexed grains as a function of the number of single-grain fitting solutions considered in the prototype selection. Different number of true grains in the sample is considered, using fiducial noise level $\sigma=0.143$ deg. In this test we considered both SynthSampleA and SynthSampleB. As expected, more initializations are needed for high number of grains.

We compare the performance between LAUEOT and the LAUE3DNDT method from [22], which is currently used for wide-beam Laue tomography problems. Table 7.4 shows the comparison between the methods limited to results that can be obtained within 24 hours. This comparison is performed for SynthSampleA using the fiducial noise and outliers, with 32 rotation angles. In this test, LAUEOT finds the exact number of grains within 1h on a single large memory GPU, while LAUE3DNDT outputs 356/500 grains within 24 hours. Further acceleration for LAUEOT can be achieved

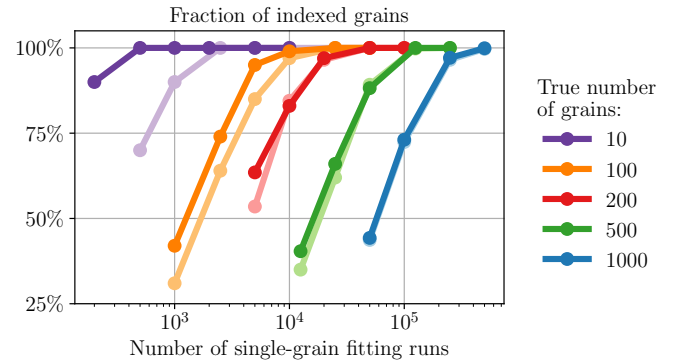


Fig. 7. Fraction of correctly indexed grains as the number of single-grain fitting initializations considered in the prototype selection, for samples SynthSample1 (darker-colored lines) and SynthSample2 (lighter-colored lines). The results are shown for varying number of true grains in the simulation, with spot noise level kept constant.

by running on multiple GPUs, as the single-grain fitting is trivially parallelizable.

7.4 Application to experimental data

We demonstrate the performance of the LAUEOT solver on spot data taken during neutron diffraction experiments. We show results on three samples: (1) a ruby single-grain calibration sample, (2) a CoNiGa alloy oligocrystal, and (3) a FeGa alloy polycrystal. The details of the experiments are described in Appendix F. Table 7.4 describes the samples in terms of the number of projection angles used in the experiment, the number of spots detected from the images, and the wavelength range used in the LAUEOT algorithm. The number of grains found by LAUEOT is 1, 13, and 251 for the ruby, CoNiGa, and FeGa, respectively. The runtime for this analysis was few minutes for the ruby, 5 min for CoNiGa and 30 min for FeGa, on a single A100 GPU. The number of matched spots is 70–80%. The orientation and position of the ruby crystal found by LAUEOT agrees well with analysis by LAUE3DNDT. The ruby crystal measurements was used to calibrate the experimental setup in terms




Fig. 8. Results of LAUEOT indexing for three samples: a ruby single crystal (left panel), an oligocrystalline CoNiGa alloy (middle panel), and a polycrystalline FeGa alloy (right panel). The panels show spots detected at all projection angles stacked together, for the forward-scatter detector. For FeGa, we only show the matches for the first 20 grains. The black spots are detected from experimental images, while the colored spots show the spots for the fitted model grains. Only matched spots are shown; outliers are omitted for clarity. The statistics of the matched spots are described in Table 7.4.

TABLE 1

Crystal grain indexing results for three samples. The numbers in brackets correspond to spot counts in the (backscatter, forward-scatter) detectors.

Sample	Number of projections	Wavelength λ [Å]	Number of spots	Number of grains	Number of matched spots
Ruby	46	[1.0-5.0]	3512 (3093, 419)	1	2852 (2468, 384)
CoNiGa	87	[1.0-5.0]	16822 (11246, 5576)	13	11883 (9039, 2844)
FeGa	86	[1.1-5.0]	269181 (197047, 72134)	251	214471 (169240, 45231)

TABLE 2

Time to solution for the LAUEOT algorithm compared to the solution from LAUE3DNDT [22].

	sample	number of grains found	time-to-solution
LAUE3DNDT	$N_{\text{grains}} = 200$	196	24h
LAUEOT	$N_{\text{grains}} = 200$	200	8 min on a single A100
LAUE3DNDT	$N_{\text{grains}} = 500$	356	24h
LAUEOT	$N_{\text{grains}} = 500$	500	30 min on a single A100, 5m on 8×A100

of distances between the detectors and the sample, as well as tilts of the detector plane [38].

Figure 8 shows the results of indexing of the ruby, CoNiGa, and FeGa samples in the left, middle, and right panels, respectively. The panels show stacked spots for all projection angles in the forward-scatter detector. As in Figure 3, the black points show the spots detected in the experiment, and the colored points show the model fitted by LAUEOT, with different colors corresponding to model grains. For clarity, only the matched spots are shown, the outliers are omitted. For the FeGa sample, we only show spots for the first 20 grains. The empty area in the middle is due to the fact that rays close to direction of the beam can not satisfy the Bragg condition (Equation 1). The agreement is very good for the ruby, and good for the CoNiGa and FeGa samples, up to the precision level of calibrations.

8 CONCLUSIONS

We presented a novel approach to indexing grains in multi-crystal samples from wide-beam diffraction experiments. Our method, called LAUEOT, poses a principled optimization problem formulation for the grain indexing. The problem is solved using a computationally efficient

coordinate descent method, with assignment of grains-to-spots performed in the framework of optimal transport. A crucial part of this model is the treatment of outliers in the experimental data and of redundant model spots. To treat the outliers we propose a model dedicated to the Laue indexing problem, based on a full probabilistic description of outliers. Our new probabilistic model clearly outperforms the two existing formulations in the literature, partial and unbalanced OT.

We show the performance of LAUEOT using simulations for variable number of grains in the sample, as well as different noise levels and outlier fractions. For samples with < 500 , the method recovers all the input grains. For larger samples, the method occasionally misses a single grain. We found the limit of the stable performance of the algorithm is reached for samples with large number of grains at high noise levels. The accuracy of the fitted grain positions and orientations is generally very good, but decreases slightly with increased noise level.

The LAUEOT methods opens the possibility of analyzing samples with large number of grains with short time-to-solution. For samples with ~ 500 grains, the time-to-solution LAUEOT on a Nvidia DGX node (8× A100 GPUs) is of

order few minutes, compared to >24 hours taken by the original indexing method of LAUE3DNDT. This enables almost real-time sample analysis during the experiment, as well as simulation-based setup optimization before starting the measurements.

APPENDIX A WAHBA'S PROBLEM

For two sets of points and their correspondences, the rotation matrix can be found by solving the Orthogonal Procrustes problem. This problem is solved through the singular vector decomposition. *Wahba's* problem [70] generalizes that to the analysis of points with weights and addresses the sign ambiguity of the SVD output. The rotation matrix R for a fixed set of K pair of points a_k and b_k and the assignment weight w_k using the subroutine **Wahba** (Algorithm 2). Similarly to the Orthogonal Procrustes problem, it uses the singular value decomposition, but additionally uses a scaling by the product of determinants to remove the sign ambiguity in the SVD.

APPENDIX B M-STEPS FOR GRAIN POSITION AND ROTATION

In this section we derive the M-step updates for grain position x and rotation R , as introduced in Section 4.4. For both parameters, we find the updates corresponding to the maximum of the likelihood in Equation 20. The M-step requires solving the optimization problem for position x_n and R_n of every grain $n \in N$ separately

$$\forall n \in N : \min_{x_n, R_n} \mathcal{L}(R_n, x_n) = \frac{1}{\sigma^2} \sum_{ij} Q_{ij} \hat{r}_{ij}^\top r_j. \quad (28)$$

We perform coordinate descent alternating between updates for R and x .

B.1 Grain position

We solve this problem for x using the Newton's method. The gradient $\nabla \mathcal{L}(x)$ and Hessian $\nabla^2 \mathcal{L}(x)$ are

$$\nabla \mathcal{L}(x) = \frac{1}{\sigma^2} \sum_{ij} \frac{Q_{ij}}{\tau_{ij}} r_j^\top (\mathbb{I} - \hat{r}_{ij} \hat{r}_{ij}^\top), \quad (29)$$

$$\nabla^2 \mathcal{L}(x) = \frac{1}{\sigma^2} \sum_{ij} \frac{Q_{ij}}{\tau_{ij}^2} \left(3 \cdot r_j^\top \hat{r}_{ij} \cdot \hat{r}_{ij} \hat{r}_{ij}^\top - r_j \hat{r}_{ij}^\top - \hat{r}_{ij} r_j^\top - r_j^\top \hat{r}_{ij} \cdot \mathbb{I} \right), \quad (30)$$

where $\tau_{ij} := \|s_i - \Gamma_{t_j} x_{n_j}\|$ is the length of the ray \hat{r}_{ij} from the grain center to the spot position s_i on the screen. The Newton's method finds the step Δx by solving $\Delta x \leftarrow [\nabla^2 \mathcal{L}(x)]^{-1} \nabla \mathcal{L}(x)$. Subroutine **LauePosNewton** in Algorithm 1 describes the details of this process. The iterations stop when a convergence threshold ζ^{tol} on position is reached.

B.2 Grain rotation

As described in Section 4.4, finding R requires solving a quadratic programming problem with orthogonality constraints. In this section we derive the **LaueRotMM** algorithm, which finds the maximum likelihood solution. The optimization problem in Equation 20 is equivalent to

$$\min_R - \sum_{j \in J_n} \left\langle \sum_{i \in I} Q_{ij} \hat{r}_{ij}, \Gamma_{t_j} R w_{m_j} w_{m_j}^\top R^\top \Gamma_{t_j}^\top e \right\rangle, \quad (31)$$

such that $R^\top R = \mathbb{I}$ and where we dropped the index of the grain from the rotation matrix R . By introducing the matrix

$$M_m^0 := - \sum_{j \in J_n \cap J_m} \Gamma_{t_j}^\top \left[\sum_{i \in I} Q_{ij} (\hat{r}_{ij} e^\top + e \hat{r}_{ij}^\top) \right] \Gamma_{t_j}, \quad (32)$$

this problem can be expressed as

$$\min_{R: R^\top R = \mathbb{I}} - \sum_{m \in M} w_m^\top R^\top M_m^0 R w_m. \quad (33)$$

The matrices M_m^0 are generally not positive semi-definite, which makes the problem not concave. By using a positive semi-definite matrix

$$M_m := M_m^0 + |\iota_m| \mathbb{I}, \quad (34)$$

where $|\iota_m|$ is the smallest eigenvalue or 0 in case M_m^0 has no negative eigenvalue, we can write an equivalent optimization problem

$$\min_{R: R^\top R = \mathbb{I}} - \sum_m w_m^\top R^\top M_m R w_m. \quad (35)$$

This problem is the Quadratic Program with Orthogonality Constraints (QP-OC). There exist multiple solvers for this problem in literature [71], [72]. Here we design a dedicated Majorization Minimization method using an upper bound on the function in Equation 33.

The associated quadratic is an upper bound of its first order Taylor expansion around a rotation matrix R^k and so

$$\begin{aligned} & - \sum_m w_m^\top R^\top M_m R w_m \\ & \leq \sum_m w_m^\top R^{k^\top} M_m R^k w_m - \sum_m 2 w_m^\top R^{k^\top} M_m R w_m. \end{aligned} \quad (36)$$

The minimization of the upper bound is a Wahba's problem (see Appendix A) with unity weights:

$$\min_{R: R^\top R = \mathbb{I}} - \sum_m w_m^\top R^{k^\top} M_m R w_m \quad (37)$$

$$\Leftrightarrow \min_{R: R^\top R = \mathbb{I}} \sum_m \|R w_m - M_m R^k w_m\|^2. \quad (38)$$

The Minimization-Majorization scheme starts with an initial guess R^0 and performs a number of iterative steps according to Equation 38. The routine **LaueRotMM** is shown in Algorithm 3. The iterative steps are performed until a given convergence thresholds on rotation ζ^{tol} is reached.

APPENDIX C

OPTIMAL TRANSPORT VARIANTS

As described in Section 4.2, our main proposed method for finding the transport plan between the model and detected spots is using a custom probabilistic model for outlier treatment. We compare the performance of this model to other OT variants that allow for outlier treatment: unbalanced and partial optimal transport. For all our Sinkhorn methods, the convergence is controlled by a maximum difference to previous iteration ζ^{tol} .

C.1 Partial optimal transport

Partial optimal transport [60] aims to move only a pre-defined fraction m_{OT} of the mass —here the number of spots— that has to be transported in the OT formulation. The optimization problem has a single control parameter, which is specified by the user. The objective function for finding the transportation plan Q is

$$\begin{aligned} \max_Q \quad & \langle L, Q \rangle + \epsilon \mathbf{H}(Q) \\ \text{s.t.} \quad & Q \geq 0, \quad Q \in \mathcal{Q} \\ & Q\mathbb{1} \leq a, \quad Q^\top \mathbb{1} \leq b, \quad \mathbb{1}^\top Q^\top \mathbb{1} = m_{\text{OT}}, \end{aligned} \quad (\text{A})$$

where a and b are vectors corresponding to spot weights (typically equal and set to one), and $m_{\text{OT}} \leq \min\{\|a\|_1, \|b\|_1\}$, and ϵ is the temperature parameter (see Section 5.1), which should be set to $\epsilon=1$ to match Equation 13. In partial OT, the spots to be matched are selected automatically based on their cost. A motivation to consider Partial OT in the presence of outliers and missing spots is that one can hope that only inliers will be matched with their corresponding models spots.

C.2 Unbalanced optimal transport

In this formulation, the constraint that the mass distribution in Q should have marginals that are equal to a and b is relaxed, and they are only encouraged to be close to them by the means of regularization, typically a generalized KL-divergence regularizer due to availability of closed-form solution [46]. The optimization problem for transportation plan Q in Problem (B) is

$$\begin{aligned} \max_Q \quad & \langle L, Q \rangle + \epsilon \mathbf{H}(Q) - \kappa \mathbf{D}(Q\mathbb{1}, a) - \lambda \mathbf{D}(Q^\top \mathbb{1}, b) \\ \text{s.t.} \quad & Q \geq 0, \quad Q \in \mathcal{Q} \end{aligned} \quad (\text{B})$$

where the proximity term $\mathbf{D}(x, y)$ is the generalized Kullback-Leibler divergence between x and y , $\mathbf{D}(x, y) = x^\top \log(x \oslash y) + (y - x)^\top \mathbb{1}$, where \oslash is the element-wise division, and ϵ is the entropic regularization parameter. The strength of these marginal constraint regularization terms controlled by κ and λ .

In practice, due to the entropic term, a consequence of relaxing the equality to the marginal, is that, as in partial OT, the mass transported is smaller than the total mass, which motivates the possible use of this method in the presence of outliers and missing spots. Here, κ and λ influence the outliers in source and target distributions, which will be ignored if the cost of transporting them is too high. In the limiting case of $\kappa \rightarrow \infty, \lambda \rightarrow \infty$, the solution converges to the balanced optimal transport where no outliers are allowed.

APPENDIX D

SINGLE-GRAIN FITTING

The single-grain fitting is performed using the simplified version of LAUEOT. The key difference is that we remove the double-stochastic constraint, which does not make sense for a single grain. The reason for this is that the model spots from a single grain are sufficiently far from each other, that it is extremely unlikely that they would be matched to a common detected spot. Furthermore, removing this constraint enables us to perform assignment updates in a single step, without having to solving Sinkhorn. This speeds up the process significantly. We employ the deterministic annealing scheme here too to obtain better local minima. We maximize the function

$$\begin{aligned} \max_{Q, R, x} \quad & \langle L, Q \rangle + \epsilon \mathbf{H}(Q) \\ \text{s.t.} \quad & Q \geq 0, \quad Q \in \mathcal{Q}, \quad Q\mathbb{1} = a, \quad R^\top R = \mathbb{I}, \end{aligned} \quad (39)$$

where L is the likelihood matrix in Equation 19. With the double-stochastic constraint removed, the assignment matrix is computed in a single update. Algorithm 8 for single-grain fitting is described in detail in Appendix D. We perform the optimization for model spots generated from a single grain model. The annealing schedule is following the same rules as the multi-grain fitting. Here we set the number of iterations to a fixed number k^{max} , which allows us to use perform multiple single-grain fitting runs in a batch mode. This leads to a significant computational performance boost. For grain position and rotation updates, we perform 4 steps in the Newton and MM iterations. This is done to speed up the computation, as the single-grain fitting prioritizes speed over precision.

APPENDIX E

SELECTION OF PROTOTYPES WITH OT

Let's consider a set of candidate grains $\{\mathcal{G}_n^c\}_{n \in N^c}$, indexed by N^c , which are solutions obtained by single-grain fitting, that have the number of outliers $< f^{\text{out}}|\mathcal{S}_n|$. Each candidate will create model spots indexed by J_n^c . The goal is to find a subset of initial grains $N^p \subseteq N^c$, called the *prototypes*, that will constitute the final sample. We start with calculating the grain-to-spot assignment log-likelihood matrix $\bar{L} \in \mathbb{R}^{|I| \times |N^c|}$. The cost of assignment of spot s_i to grain \mathcal{G}_n^c is calculated using the marginal probability of assignment for all model spots in the grain

$$\bar{L}_{in} \leftarrow \log \sum_{j \in J_n^c} p(s_i, Z_{ij} = 1). \quad (40)$$

We then shift the cost matrix to positive values

$$\bar{L}_{in} \leftarrow \|\bar{L}\|_\infty + \bar{L}_{in}. \quad (41)$$

The prototype selection depends on finding a weight vector $w \in \Delta_{N^c}$ corresponding to candidate grains, where simplex $\Delta_{N^c} := \{z \in \mathbb{R}_+^{N^c} \mid z^\top \mathbb{1} = 1\}$. We define a set function

$$\begin{aligned} f(N^p) &:= \max_{w: \text{supp}(w) \subseteq N^p} \max_P \langle \bar{L}, P \rangle \\ \text{s.t.} \quad & P \geq 0, \quad P\mathbb{1} = a, \quad P^\top \mathbb{1} = w, \end{aligned} \quad (42)$$

where P is a OT assignment matrix and $f(\emptyset) = 0$. The objective function finds N^* that maximizes f subject to

cardinality constraint on $|N^*| < k$, enforcing that the size of the set is smaller than k . The SPOT loss is

$$N^* = \operatorname{argmax}_{N^p \subseteq N^c, |N^p| \leq k} f(N^p). \quad (43)$$

Finding the optimum subset N^* is NP complete, but can be solved approximately using greedy incremental methods, which leverage its submodular property. The algorithm proposed in [45] has proven approximation guarantees [73].

We employ the SPOTGREEDY algorithm [45] to find the approximate solution for Equation 43. This algorithm selects prototypes in a greedy fashion, starting with the one that contributes the most to the objective function. The relative contributions of consecutive prototypes to the objective reflect the degree of uniqueness of the prototype. When the increment in the cost functions for a new candidate prototype is small, that indicates that it is similar to a prototype that already has been accepted. Due the greedy nature of the algorithm, for cardinality $k + 1$ it will output the same elements in the same order as for k , extended by one new element.

Practically, the SPOTGREEDY algorithm starts with constant outputs a sequence of indices and their corresponding objective function values $(n_k, f_k)_{k \in N^*}$, indexed by N^* . A subsequence with elements up to k corresponds to the prototypes $N^p \subseteq N^c$, for cardinality constraint $|N^p| \leq k$. To choose the final set of prototypes N^p with size k^* , we employ a criterion based on 2nd discrete difference of f_k

$$k^* \leftarrow \operatorname{argmax}_k f_{k+1} + f_{k-1} - 2f_k \quad (44)$$

for $k > 1$. We use a fractional criterion instead of the commonly used cut-off on the objective increment ϵ (see Algorithm 1 in [45]), as it requires no knowledge of the numerical values of the objective from the user. We call this method **LaueSpotGreedy** (Algorithm 7) and describe it in Appendix E. The prototype grain set is passed on to the next step of the algorithm, which solves problem in Equation 20. This procedure typically finds the correct number of grains, except when the number of candidates obtained from single-grain fitting is not sufficiently large. The number of single-grain fitting candidates required is described in Section 7.

APPENDIX F

EXPERIMENTAL DETAILS

F.1 Samples

Ruby: As a first sample, a \emptyset 6 mm spherical single-crystal ruby, typically used for detector calibrations, was measured. For the indexing of the ruby crystal, a $R\bar{3}c$ space group with lattice parameters of $a=b=4.7606$ Å, $c=12.994$ Å was used.

CoNiGa: A cylindrical \emptyset 4 mm \times 4 mm oligocrystalline CoNiGa sample was selected as an intermediate test case. The sample was prepared via hot-extrusion followed by post-extrusion heat treatment, resulting in millimeter-sized grains. The full details of the sample preparation are given in [69]. For the indexing of the CoNiGa sample a cubic $Pm\bar{3}m$ space group with lattice parameters of $a=b=c=2.8657$ Å was used.

FeGa: As a final test case, a $3 \times 3 \times 3$ mm cube $\text{Fe}_{0.9}\text{Ga}_{0.1}$ polycrystalline sample was measured. The sample was prepared via arc-melting of 99.95% purity Fe and Ga buttons in an argon atmosphere. To obtain a homogenous distribution of elements, the sample was re-melted 4 times. After arc-melting, the sample was annealed in argon at a temperature of 1000 °C with a heating rate of about 20 °C/min and a holding time of 12 h with subsequent quenching to room temperature. The final sample shape was achieved via spark cutting and subsequent polishing with sandpaper, diamond suspension, and colloidal silica. For the indexing of the FeGa data set, a $Im\bar{3}m$ space group with lattice parameters of $a=b=c=2.917$ Å were used.

The noise level was set to $\sigma=1.07^\circ$ and $\sigma=0.27^\circ$ for CoNiGa and FeGa samples, respectively. The CoNiGa sample required higher noise level, as the uncertainty also included systematic errors in the detector setup calibration. For both CoNiGa and FeGa samples, the number of spots detected in each rotation varied by up to 40%. This is not expected, as the number of spots per projection should remain constant. This large variation can be caused by drops in beam flux for that exposure, as well as failures of the spot detection algorithm. Some rotation angles were removed from the analysis, as they had abnormally small number of detected spots compared to other angles. To address this, we set the single-grain selection threshold $f^{\text{out}} < 0.4$ for CoNiGa and $f^{\text{out}} < 0.35$ for FeGa.

F.2 Neutron Laue diffraction tomography

All neutron Laue diffraction tomographies were collected at the Pulse Overlap Diffractometer (POLDI) [74] at the SINQ neutron source of the Paul Scherrer Institute (PSI), Switzerland, using a white beam and the FALCON double-detector system [38]. The FALCON forward- and backscattering detectors were placed upstream and downstream from the sample, respectively. Measurement parameters such as the sample-to-detector distances and exposure times were chosen for each sample to optimize signal quality and overall spot coverage on the detector.

The ruby tomography was carried out with sample-to-detector distances of 16.5 cm and 16.8 cm for the back- and forward-scattering detector, respectively. The sample was rotated a full 360° in steps of 8° , resulting in 46 projections in total. Each projection was measured with an exposure time of 90 s.

The CoNiGa tomography was carried out with sample-to-detector distances of 16.5 cm and 12.5 cm for the back- and forward-scattering detector, respectively. The tomography was performed by rotating the sample a full 360° in steps of 4° , resulting in a total of 91 projections. Each projection was measured with an exposure time of 90 s.

The FeGa tomography was carried out with sample-to-detector distances of 16.5 cm and 12.5 cm for the back- and forward-scattering detector, respectively. The tomography was performed over a full 360° rotation in steps of 4° , resulting in a total of 91 projections. Each projection was measured with an exposure time of 100 s.

APPENDIX G

ALGORITHMS

The symbol \odot stands for element-wise division. Element-wise minimum of elements between two vectors x and y is denoted as $\min(x, y)$. The $\text{diag}(x)$ function creates a diagonal matrix from an input vector x . Vectors a and b correspond to data point weights, which in LAUEOT are set to $a \leftarrow \mathbb{1}, b \leftarrow \mathbb{1}$. LAUEOT employs balanced optimal transport with explicit outlier modeling, the iterative updates are shown in Algorithm 4. For UNBALANCEDOT, the iterative updates are shown in Algorithm 5. For PARTIALOT, the iterative updates are shown in Algorithm 6. These algorithms are described in Appendix C.

Algorithm 1 LauePosNewton: Newton’s method for finding grain positions.

```

1: input:  $(s_i, w_{m_j}, \Gamma_{t_j})_{j \in J_n}, e$ 
2: input:  $x_n^0, R_n, (Q_{ij})_{j \in J_n, i \in I}$ 
3: input:  $\zeta^{\text{tol}}, k^{\text{max}}$ 
4: set:  $k \leftarrow 0, \zeta \leftarrow \inf$ 
5: while  $\zeta > \zeta^{\text{tol}}$  and  $k < k^{\text{max}}$  do
6:    $g \leftarrow \nabla \mathcal{L}(x^k)$  (Equation 29)
7:    $H \leftarrow \nabla^2 \mathcal{L}(x^k)$  (Equation 30)
8:    $\Delta x_n^k \leftarrow H^{-1} g$ 
9:    $x_n^{k+1} \leftarrow x_n^k - \Delta x_n^k$ 
10:   $\zeta \leftarrow \|x_n^{k+1} - x_n^k\|_2$ 
11:   $k \leftarrow k + 1$ 
12: end while
13: output:  $x_n^k$ 
14: a

```

Algorithm 2 Wahba: solution to the Wahba’s problem

```

1: input:  $a, b, w$ 
2:  $B \leftarrow \sum_{k=1}^K w_k a_k b_k^\top$ 
3:  $USV^\top \leftarrow \text{svd}(B)$ 
4:  $M \leftarrow \text{diag}[1, 1, \det(U) \cdot \det(V)]$ 
5:  $R \leftarrow UMV^\top$ 
6: output:  $R$ 

```

Algorithm 3 LaueRotMM: Majorization-Minimization algorithm for finding rotations with reflected points.

```

1: input:  $(s_i, w_{m_j}, \Gamma_{t_j})_{j \in J_n}, e, R_n^0, x_n, (Q_{ij})_{j \in J_n, i \in I}, \zeta^{\text{tol}}, k^{\text{max}}$ 
2: set:  $k \leftarrow 0, \zeta \leftarrow \inf$ 
3: set:  $M_m \leftarrow M_m^0 + |\iota_m| \mathbb{I}$  (Equation 34)
4: while  $\zeta > \zeta^{\text{tol}}$  and  $k < k^{\text{max}}$  do
5:    $R_n^{k+1} \leftarrow \text{Wahba}(M_m R_n^k w_m, w_m, \mathbb{1})$ 
6:    $\zeta \leftarrow \|R_n^{k+1} - R_n^k\|_F$ 
7:    $k \leftarrow k + 1$ 
8: end while
9: output:  $R_n^k$ 

```

Algorithm 4 BalancedOT: Sinkhorn algorithm for balanced optimal transport with entropic regularization.

```

1: input:  $L, \epsilon, \zeta^{\text{tol}}$ 
2: set:  $K \leftarrow \exp(L/\epsilon), x^0 \leftarrow \mathbb{1}, y^0 \leftarrow \mathbb{1}$ 
3: set:  $\zeta \leftarrow \inf, k \leftarrow 0$ 
4: while  $\zeta > \zeta^{\text{tol}}$  do
5:    $x^{k+1} \leftarrow a \odot (K y^k)$ 
6:    $y^{k+1} \leftarrow b \odot (K^\top x^{k+1})$ 
7:    $\zeta \leftarrow \max(\|x^{k+1} - x^k\|, \|y^{k+1} - y^k\|)$ 
8:    $k \leftarrow k + 1$ 
9: end while
10:  $Q \leftarrow \text{diag}(x^k) \cdot K \cdot \text{diag}(y^k)$ 
11: output:  $Q$ 

```

Algorithm 5 UnbalancedOT: Sinkhorn algorithm for unbalanced optimal transport with entropic regularization.

```

1: input:  $L, \epsilon, \kappa, \lambda, \zeta^{\text{tol}}$ 
2: set:  $K \leftarrow \exp(L/\epsilon), x^0 \leftarrow \mathbb{1}, y^0 \leftarrow \mathbb{1}$ 
3: while  $\zeta > \zeta^{\text{tol}}$  do
4:    $x^{k+1} \leftarrow (a \odot (K y^k))^{\frac{\lambda}{\lambda+\epsilon}}$ 
5:    $y^{k+1} \leftarrow (b \odot (K^\top x^{k+1}))^{\frac{\kappa}{\kappa+\epsilon}}$ 
6:    $\zeta \leftarrow \max(\|x^{k+1} - x^k\|, \|y^{k+1} - y^k\|)$ 
7:    $k \leftarrow k + 1$ 
8: end while
9:  $Q \leftarrow \text{diag}(x^k) \cdot K \cdot \text{diag}(y^k)$ 
10: output:  $Q$ 

```

Algorithm 6 PartialOT: Sinkhorn algorithm for partial optimal transport with entropic regularization.

```

1: input:  $L, \epsilon, m_{\text{OT}}, \zeta^{\text{tol}}$ 
2: set:  $K \leftarrow \exp(L/\epsilon), Q_1, Q_2, Q_3 \leftarrow \mathbb{I}_{IJ}$ 
3: set:  $x \leftarrow \mathbb{1}, y \leftarrow \mathbb{1}, K^0 \leftarrow K$ 
4: while  $\zeta > \zeta^{\text{tol}}$  do
5:    $K_1 \leftarrow \text{diag}[\min(a \odot K^k \mathbb{1}, x)][K^k \cdot Q_1]$ 
6:    $Q_1 \leftarrow Q_1 \cdot K^k \odot K_1$ 
7:    $K_2 \leftarrow [K_1 \cdot Q_2] \text{diag}[\min(b \odot K_1^\top \mathbb{1}, y)]$ 
8:    $Q_2 \leftarrow Q_2 \cdot K_1 \odot K_2$ 
9:    $K^{k+1} \leftarrow K_2 \cdot Q_3 \cdot m_{\text{OT}} / (\mathbb{1}^\top K_2 \mathbb{1})$ 
10:   $Q_3 \leftarrow Q_3 \cdot K_2 \odot K^k$ 
11:   $\zeta \leftarrow \|K - K^{\text{prev}}\|_F$ 
12:   $k \leftarrow k + 1$ 
13: end while
14:  $Q \leftarrow K$ 
15: output:  $Q$ 

```

Algorithm 7 LaueSpotGreedy: Selection of grain prototypes using optimal transport in the Laue indexing problem.

```

1: input data:  $N^c, \bar{L}_{in}$ 
2: initialize:  $N^0 = \emptyset$ 
3: for  $1 \leq k \leq |N^c|$  do
4:   Define vector  $f_n \leftarrow f(N^{k-1} \cup N_n^c) - f(N^{k-1})$ 
5:   Find largest increment  $n_k \leftarrow \max_n f_n$ 
6:   Create a grain set for  $k$ :  $N^k \leftarrow N^{k-1} \cup N_{n_k}^c$ 
7:   Store the increment function  $f_k \leftarrow f_{n_k}$ 
8: end for
9: Get final number of grain prototypes  $k^*$  (Equation 44)
10: Get the final prototype grain set  $N^* \leftarrow N^{k^*}$ 
11: output:  $N^*$ 

```

Algorithm 8 SingleLaueOT: Single-grain fitting for Laue indexing.

```

1: input data:  $(s_i)_{i \in I}, \sigma$ 
2: input number of starting grains:  $N^s$ 
3: input outlier parameters:  $L^{\text{out}}$ 
4: input control parameters:  $\epsilon^{\text{init}}, \epsilon^{\text{cool}}, \zeta^{\text{tol}}, k^{\text{max}}$ 
5: set: Calculate Sobol grid of starting grains  $(R_n, x_n)_{n=1}^{N^s}$ 
6: set:  $L_0 \leftarrow L^{\text{out}}$ 
7: set:  $L_0 \leftarrow \sigma^{-2} - \log(4\pi\sigma^2 \sinh(\sigma^{-2}))$ 
8: for  $0 < n < N^s$  do
9:    $\triangleright$  Use a sample with a single grain
10:   $\epsilon \leftarrow \epsilon^{\text{init}}$ 
11:  for  $0 < k < k^{\text{max}}$  do
12:    Calculate model rays,  $\forall j \in J$ :
13:     $r_j \leftarrow (\mathbb{I} - 2\Gamma_{t_j} R_{n,j}^k w_{m_j} w_{m_j}^\top R_{n,j}^{k\top} \Gamma_{t_j}^\top) e$ 
14:    Calculate detected ray estimates,  $\forall i \in I, j \in J$ :
15:     $\hat{r}_{ij} \leftarrow \frac{s_i - \Gamma_{t_j} x_{n,j}^k}{\|s_i - \Gamma_{t_j} x_{n,j}^k\|}$ 
16:    Log likelihood matrix,  $\forall i \in I, j \in J$ :
17:     $L_{ij} \leftarrow \frac{1}{\sigma^2} \hat{r}_{ij}^\top r_j$ 
18:     $\triangleright$  E-step
19:    Update assignments:
20:     $K \leftarrow \exp(L/\epsilon)$ 
21:     $Q_n^k \leftarrow K \cdot \text{diag}(a \oslash K \mathbb{1})$ 
22:     $\triangleright$  M-step
23:    Update positions and orientations,  $\forall n \in N$ :
24:     $x_n^{k+1} \leftarrow \text{LauePosNewton}[x_n^k, Q_n^k, \dots]$ 
25:     $R_n^{k+1} \leftarrow \text{LaueRotMM}[R_n^k, Q_n^k, \dots]$ 
26:    Decrease the temperature:  $\epsilon \leftarrow \epsilon \cdot \epsilon^{\text{cool}}$ 
27:     $k \leftarrow k + 1$ 
28:  end for
29: end for
30: output:  $(R_n, x_n)_{n=1}^{N^s}$ 

```

ACKNOWLEDGMENTS

The authors thank Marc Caubet Serrabou and administrators of the Merlin cluster at PSI for excellent technical support. We thank Andreas Adelman for the access to the Gwendolen node. TK thanks Luis Barba for helpful discussions and Ilnura Usmanova for valuable comments on the manuscript. The authors thank Benjamin Bejar for his initial work conceptual works and contributions to proposals that led to this project. The project was supported by grant C20-14 from Swiss Data Science Center.

REFERENCES

- [1] M. Lax, "Multiple scattering of waves," *Rev. Mod. Phys.*, vol. 23, pp. 287–310, Oct 1951. [Online]. Available: <https://link.aps.org/doi/10.1103/RevModPhys.23.287>
- [2] T. Welberry and T. Weber, "One hundred years of diffuse scattering," *Crystallography Reviews*, vol. 22, no. 1, pp. 2–78, 2016. [Online]. Available: <https://doi.org/10.1080/0889311X.2015.1046853>
- [3] A. Morawiec, *Indexing of Crystal Diffraction Patterns*, 2022.
- [4] H. O. Sørensen, S. Schmidt, J. P. Wright, G. B. M. Vaughan, S. Techert, E. F. Garman, J. Oddershede, J. Davaasambuu, K. S. Paithankar, C. Gundlach, and H. F. Poulsen, "Multigrain crystallography," vol. 227, no. 1, pp. 63–78, 2012. [Online]. Available: <https://doi.org/10.1524/zkri.2012.1438>
- [5] M. A. Groeber, B. Haley, M. D. Uchic, D. M. Dimiduk, and S. Ghosh, "3D reconstruction and characterization of polycrystalline microstructures using a FIB-SEM system," *Materials Characterization*, vol. 57, no. 4-5, pp. 259–273, 2006.

Algorithm 9 LaueOT: Laue indexing with optimal transport, multi-grain fitting

```

1: input data:  $(s_i)_{i \in I}, \sigma$ 
2: input prototype grain parameters:  $(R_n^0, x_n^0)_{n \in N}$ 
3: input OT parameters:  $L^{\text{out}}, T_a$ 
4: input control parameters:  $\epsilon^{\text{init}}, \epsilon^{\text{cool}}, \zeta^{\text{tol}}$ 
5: set:  $L_0 \leftarrow \sigma^{-2} - \log(4\pi\sigma^2 \sinh(\sigma^{-2}))$ 
6: set:  $L_0 \leftarrow L^{\text{out}}$ 
7: set:  $\epsilon \leftarrow \epsilon^{\text{init}}, \zeta \leftarrow \inf$ 
8: while  $\epsilon \geq 1$  or  $\zeta > \zeta^{\text{tol}}$  do
9:   Calculate model rays,  $\forall j \in J$ :
10:   $r_j \leftarrow (\mathbb{I} - 2\Gamma_{t_j} R_{n,j}^k w_{m_j} w_{m_j}^\top R_{n,j}^{k\top} \Gamma_{t_j}^\top) e$ 
11:   Calculate detected ray estimates,  $\forall i \in I, j \in J$ :
12:   $\hat{r}_{ij} \leftarrow \frac{s_i - \Gamma_{t_j} x_{n,j}^k}{\|s_i - \Gamma_{t_j} x_{n,j}^k\|}$ 
13:   Calculate log-likelihood matrix,  $\forall i \in I, j \in J$ :
14:   $L_{ij} \leftarrow \frac{1}{\sigma^2} \hat{r}_{ij}^\top r_j$ 
15:   $\triangleright$  E-step
16:  Update assignments by solving OT:
17:   $Q^{k+1} \leftarrow \text{BalancedOT}[L, \epsilon, \dots]$ 
18:   $\triangleright$  M-step
19:  Update orientations and positions,  $\forall n \in N$ :
20:   $x_n^{k+1} \leftarrow \text{LauePosNewton}[x_n^k, Q^{k+1}, \dots]$ 
21:   $R_n^{k+1} \leftarrow \text{LaueRotMM}[R_n^k, Q^{k+1}, \dots]$ 
22:   $\triangleright$  Deterministic annealing temperature control
23:  if  $\epsilon > 1$  then
24:    Decrease temperature:  $\epsilon \leftarrow \epsilon \cdot \epsilon^{\text{cool}}$ 
25:  end if
26:   $\zeta \leftarrow \|Q^{k+1} - Q^k\|_F$ 
27:   $k \leftarrow k + 1$ 
28: end while
29: output:  $Q^k, (R_n^k, x_n^k)_{n \in N}$ 

```

- [6] S. Zaefferer, S. Wright, and D. Raabe, "Three-dimensional orientation microscopy in a focused ion beam-scanning electron microscope: A new dimension of microstructure characterization," *Metallurgical and Materials Transactions A*, vol. 39, no. 2, pp. 374–389, 2008.
- [7] H. Liu, S. Schmidt, H. F. Poulsen, A. Godfrey, Z. Liu, J. Sharon, and X. Huang, "Three-dimensional orientation mapping in the transmission electron microscope," *Science*, vol. 332, no. 6031, pp. 833–834, 2011.
- [8] P. J. Konijnenberg, S. Zaefferer, and D. Raabe, "Assessment of geometrically necessary dislocation levels derived by 3D EBSD," *Acta Materialia*, vol. 99, pp. 402–414, 2015.
- [9] G. Stechmann, S. Zaefferer, P. Konijnenberg, D. Raabe, C. Gretener, L. Kranz, J. Perrenoud, S. Buecheler, and A. N. Tiwari, "3-Dimensional microstructural characterization of CdTe absorber layers from CdTe/CdS thin film solar cells," *Solar Energy Materials and Solar Cells*, vol. 151, pp. 66–80, 2016.
- [10] S. Smeets, X. Zou, and W. Wan, "Serial electron crystallography for structure determination and phase analysis of nanocrystalline materials," *Journal of Applied Crystallography*, vol. 51, no. 5, pp. 1262–1273, 2018.
- [11] W. Zhu, G. Wu, A. Godfrey, S. Schmidt, Q. He, Z. Feng, T. Huang, L. Zhang, and X. Huang, "Five-parameter grain boundary character distribution of gold nanoparticles based on three dimensional orientation mapping in the tem," *Scripta Materialia*, vol. 214, p. 114677, 2022.
- [12] H. F. Poulsen, S. F. Nielsen, E. M. Lauridsen, S. Schmidt, R. Suter, U. Lienert, L. Margulies, T. Lorentzen, and D. Juul Jensen, "Three-dimensional maps of grain boundaries and the stress state of individual grains in polycrystals and powders," *Journal of Applied Crystallography*, vol. 34, no. 6, pp. 751–756, 2001.
- [13] B. Larson, W. Yang, G. Ice, J. Budai, and J. Tischler, "Three-dimensional X-ray structural microscopy with submicrometre resolution," *Nature*, vol. 415, no. 6874, pp. 887–890, 2002.
- [14] D. J. Jensen, E. Lauridsen, L. Margulies, H. Poulsen, S. Schmidt, H. Sørensen, and G. Vaughan, "X-ray microscopy in four dimensions," *Materials Today*, vol. 9, no. 1-2, pp. 18–25, 2006.

- [15] W. Ludwig, S. Schmidt, E. M. Lauridsen, and H. F. Poulsen, "X-ray diffraction contrast tomography: a novel technique for three-dimensional grain mapping of polycrystals. I. direct beam case," *Journal of Applied Crystallography*, vol. 41, no. 2, pp. 302–309, 2008.
- [16] G. Johnson, A. King, M. G. Honnicke, J. Marrow, and W. Ludwig, "X-ray diffraction contrast tomography: a novel technique for three-dimensional grain mapping of polycrystals. II. the combined case," *Journal of Applied Crystallography*, vol. 41, no. 2, pp. 310–318, 2008.
- [17] W. Ludwig, P. Reischig, A. King, M. Herbig, E. Lauridsen, G. Johnson, T. Marrow, and J.-Y. Buffiere, "Three-dimensional grain mapping by x-ray diffraction contrast tomography and the use of Friedel pairs in diffraction data analysis," *Review of Scientific Instruments*, vol. 80, no. 3, p. 033905, 2009.
- [18] W. Ludwig, A. King, P. Reischig, M. Herbig, E. M. Lauridsen, S. Schmidt, H. Proudhon, S. Forest, P. Cloetens, S. R. Du Roscoat *et al.*, "New opportunities for 3d materials science of polycrystalline materials at the micrometre lengthscale by combined use of X-ray diffraction and X-ray imaging," *Materials Science and Engineering: A*, vol. 524, no. 1-2, pp. 69–76, 2009.
- [19] U. Lienert, S. Li, C. Hefferan, J. Lind, R. Suter, J. Bernier, N. Barton, M. Brandes, M. Mills, M. Miller *et al.*, "High-energy diffraction microscopy at the advanced photon source," *JOM*, vol. 63, no. 7, pp. 70–77, 2011.
- [20] D. J. Jensen and H. F. Poulsen, "The three dimensional X-ray diffraction technique," *Materials Characterization*, vol. 72, pp. 1–7, 2012.
- [21] S. A. McDonald, P. Reischig, C. Holzner, E. M. Lauridsen, P. J. Withers, A. P. Merkle, and M. Feser, "Non-destructive mapping of grain orientations in 3D by laboratory X-ray microscopy," *Scientific Reports*, vol. 5, no. 1, pp. 1–11, 2015.
- [22] M. Raventós, M. Tovar, M. Medarde, T. Shang, M. Strobl, S. Samothrakitis, E. Pomjakushina, C. Grünzweig, and S. Schmidt, "Laue three dimensional neutron diffraction," *Scientific Reports*, vol. 9, no. 1, p. 4798, 2019. [Online]. Available: <https://doi.org/10.1038/s41598-019-41071-x>
- [23] S. Samothrakitis, C. B. Larsen, J. Čapek, E. Polatidis, M. Raventos, M. Tovar, S. Schmidt, and M. Strobl, "Microstructural characterization through grain orientation mapping with laue three-dimensional neutron diffraction tomography," *Materials Today Advances*, vol. 15, p. 100258, 2022.
- [24] S. Peetermans and E. H. Lehmann, "Simultaneous neutron transmission and diffraction contrast tomography as a non-destructive 3D method for bulk single crystal quality investigations," *Journal of Applied Physics*, vol. 114, no. 12, p. 124905, 2013.
- [25] R. Woracek, D. Penumadu, N. Kardjilov, A. Hilger, M. Boin, J. Banhart, and I. Manke, "3D mapping of crystallographic phase distribution using energy-selective neutron tomography," *Advanced Materials*, vol. 26, no. 24, pp. 4069–4073, 2014.
- [26] S. Peetermans, A. King, W. Ludwig, P. Reischig, and E. Lehmann, "Cold neutron diffraction contrast tomography of polycrystalline material," *Analyst*, vol. 139, no. 22, pp. 5765–5771, 2014.
- [27] H. Sato, Y. Shiota, S. Morooka, Y. Todaka, N. Adachi, S. Sadamatsu, K. Oikawa, M. Harada, S. Zhang, Y. Su *et al.*, "Inverse pole figure mapping of bulk crystalline grains in a polycrystalline steel plate by pulsed neutron Bragg-dip transmission imaging," *Journal of Applied Crystallography*, vol. 50, no. 6, pp. 1601–1610, 2017.
- [28] A. Cereser, M. Strobl, S. A. Hall, A. Steuwer, R. Kiyonagi, A. S. Tremsin, E. B. Knudsen, T. Shinohara, P. K. Willendrup, A. B. da Silva Fanta *et al.*, "Time-of-flight three dimensional neutron diffraction in transmission mode for mapping crystal grain structures," *Scientific Reports*, vol. 7, no. 1, pp. 1–11, 2017.
- [29] R. Woracek, J. Santisteban, A. Fedrigo, and M. Strobl, "Diffraction in neutron imaging-A review," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 878, pp. 141–158, 2018.
- [30] S. Samothrakitis, M. Raventós, J. Čapek, C. B. Larsen, C. Grünzweig, M. Tovar, M. Garcia-Gonzalez, J. Kopeček, S. Schmidt, and M. Strobl, "Grain morphology reconstruction of crystalline materials from Laue three-dimensional neutron diffraction tomography," *Scientific Reports*, vol. 10, no. 1, pp. 1–7, 2020.
- [31] S. Schmidt, "GrainSpotter: a fast and robust polycrystalline indexing algorithm," *Journal of Applied Crystallography*, vol. 47, no. 1, pp. 276–284, Feb 2014. [Online]. Available: <https://doi.org/10.1107/S1600576713030185>
- [32] Y. Gevorkov, A. Barty, W. Brehm, T. A. White, A. Tolstikova, M. O. Wiedorn, A. Meents, R.-R. Grigat, H. N. Chapman, and O. Yefanov, "pinkIndexer – a universal indexer for pink-beam X-ray and electron diffraction snapshots," *Acta Crystallographica Section A*, vol. 76, no. 2, pp. 121–131, Mar 2020. [Online]. Available: <https://doi.org/10.1107/S2053273319015559>
- [33] Y. Gevorkov, O. Yefanov, A. Barty, T. A. White, V. Mariani, W. Brehm, A. Tolstikova, R.-R. Grigat, and H. N. Chapman, "XGANDALF – extended gradient descent algorithm for lattice finding," *Acta Crystallographica Section A*, vol. 75, no. 5, pp. 694–704, Sep 2019. [Online]. Available: <https://doi.org/10.1107/S2053273319010593>
- [34] J. A. Kalinowski, A. Makal, and P. Coppens, "The LaueUtil toolkit for Laue photocrystallography. I. Rapid orientation matrix determination for intermediate-size-unit-cell Laue data," *Journal of Applied Crystallography*, vol. 44, no. 6, pp. 1182–1189, Dec 2011. [Online]. Available: <https://doi.org/10.1107/S0021889811038143>
- [35] S. Schlitt, T. E. Gorelik, A. A. Stewart, E. Schömer, T. Raasch, and U. Kolb, "Application of clustering techniques to electron-diffraction data: determination of unit-cell parameters," *Acta Crystallographica Section A*, vol. 68, no. 5, pp. 536–546, Sep 2012. [Online]. Available: <https://doi.org/10.1107/S0108767312026438>
- [36] C. Wejdemann and H. F. Poulsen, "Multigrain indexing of unknown multiphase materials," *Journal of Applied Crystallography*, vol. 49, no. 2, pp. 616–621, Apr 2016. [Online]. Available: <https://doi.org/10.1107/S1600576716003691>
- [37] R. R. P. Purushottam Raj Purohit, S. Tardif, O. Castelnau, J. Eymery, R. Guinebreière, O. Robach, T. Ors, and J.-S. Micha, "LaueNN: neural-network-based *hkl* recognition of Laue spots and its application to polycrystalline materials," *Journal of Applied Crystallography*, vol. 55, no. 4, pp. 737–750, Aug 2022. [Online]. Available: <https://doi.org/10.1107/S1600576722004198>
- [38] S. Samothrakitis, C. B. Larsen, A. Kaestner, E. Polatidis, J. Čapek, J. Hovind, A. Fazan, J. Allibon, M. Busi, S. Schmidt, J. Kopeček, and M. Strobl, "The FALCON double-detector Laue diffractometer add-on for grain mapping at POLDI," *Journal of Applied Crystallography*, vol. 56, no. 6, pp. 1792–1801, Dec 2023. [Online]. Available: <https://doi.org/10.1107/S1600576723009640>
- [39] J. L. Schonberger and J.-M. Frahm, "Structure-from-motion revisited," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [40] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge University Press, 2004.
- [41] X. Huang, G. Mei, J. Zhang, and R. Abbas, "A comprehensive survey on point cloud registration," *arXiv preprint arXiv:2103.02690*, 2021.
- [42] G. Peyré, M. Cuturi *et al.*, "Computational optimal transport: With applications to data science," *Foundations and Trends® in Machine Learning*, vol. 11, no. 5-6, pp. 355–607, 2019.
- [43] G. Mena, A. Nejatbakhsh, E. Varol, and J. Niles-Weed, "Sinkhorn em: an expectation-maximization algorithm based on entropic optimal transport," *arXiv preprint arXiv:2006.16548*, 2020.
- [44] S. J. Wright, "Coordinate descent algorithms," *Mathematical programming*, vol. 151, no. 1, pp. 3–34, 2015.
- [45] K. S. Gurumoorthy, P. Jawanpuria, and B. Mishra, "Spot: A framework for selection of prototypes using optimal transport," in *Machine Learning and Knowledge Discovery in Databases. Applied Data Science Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings, Part IV 21*. Springer, 2021, pp. 535–551.
- [46] L. Chizat, G. Peyré, B. Schmitzer, and F.-X. Vialard, "Scaling algorithms for unbalanced optimal transport problems," *Mathematics of Computation*, vol. 87, no. 314, pp. 2563–2609, 2018.
- [47] J.-D. Benamou, G. Carlier, M. Cuturi, L. Nenna, and G. Peyré, "Iterative bregman projections for regularized transportation problems," *SIAM Journal on Scientific Computing*, vol. 37, no. 2, pp. A1111–A1138, 2015.
- [48] H. Sharma, R. M. Huizenga, and S. E. Offerman, "A fast methodology to determine the characteristics of thousands of grains using three-dimensional x-ray diffraction. ii. volume, centre-of-mass position, crystallographic orientation and strain state of grains," *Journal of Applied Crystallography*, vol. 45, no. 4, pp. 705–718, 2012. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1107/S0021889812025599>
- [49] F. Dellaert, S. Seitz, C. Thorpe, and S. Thrun, "Structure from motion without correspondence," in *Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000 (Cat. No. PR00662)*, vol. 2, 2000, pp. 557–564 vol.2.

- [50] J. L. Schonberger and J.-M. Frahm, "Structure-from-motion revisited," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [51] P. J. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 14, pp. 239–256, 1992.
- [52] S. Gold, A. Rangarajan, C.-P. Lu, S. Pappu, and E. Mjolsness, "New algorithms for 2d and 3d point matching: pose estimation and correspondence," *Pattern Recognition*, vol. 31, no. 8, pp. 1019–1031, 1998.
- [53] R. Sinkhorn, "A relationship between arbitrary positive matrices and doubly stochastic matrices," *The Annals of Mathematical Statistics*, vol. 35, no. 2, pp. 876–879, 1964. [Online]. Available: <http://www.jstor.org/stable/2238545>
- [54] M. Blondel, V. Seguy, and A. Rolet, "Smooth and sparse optimal transport," in *International conference on artificial intelligence and statistics*. PMLR, 2018, pp. 880–889.
- [55] B. Schmitzer, "Stabilized sparse scaling algorithms for entropy regularized transport problems," *SIAM Journal on Scientific Computing*, vol. 41, no. 3, pp. A1443–A1481, 2019.
- [56] T. Liu, J. Puigcerver, and M. Blondel, "Sparsity-constrained optimal transport," 2023.
- [57] D. Alvarez-Melis, T. Jaakkola, and S. Jegelka, "Structured optimal transport," in *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, A. Storkey and F. Perez-Cruz, Eds., vol. 84. PMLR, 09–11 Apr 2018, pp. 1771–1780. [Online]. Available: <https://proceedings.mlr.press/v84/alvarez-melis18a.html>
- [58] C. Hammond, *The basics of crystallography and diffraction*. International Union of Crystallography texts on crystallography, 2015, vol. 21.
- [59] R. Fisher, "Dispersion on a Sphere," *Proceedings of the Royal Society of London Series A*, vol. 217, no. 1130, pp. 295–305, May 1953.
- [60] A. Figalli, "The optimal partial transport problem," *Archive for Rational Mechanics and Analysis*, vol. 195, no. 2, pp. 533–560, 2010.
- [61] A. Rangarajan, H. Chui, and F. L. Bookstein, "The softassign procrustes matching algorithm," in *Information Processing in Medical Imaging*, J. Duncan and G. Gindi, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 1997, pp. 29–42.
- [62] N. Ueda and R. Nakano, "Deterministic annealing variant of the em algorithm," in *Advances in Neural Information Processing Systems*, G. Tesauero, D. Touretzky, and T. Leen, Eds., vol. 7. MIT Press, 1994. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/1994/file/92262bf907af914b95a0fc33c3f33bf6-Paper.pdf
- [63] R. Frühwirth and W. Waltenberger, "Redescending m-estimators and deterministic annealing, with applications to robust regression and tail index estimation," *Austrian Journal of Statistics*, vol. 37, no. 34, p. 301–317, Apr. 2016. [Online]. Available: <https://www.ajs.or.at/index.php/ajs/article/view/vol37%2C%20no3%264%20-%207>
- [64] B. Schmitzer, "A sparse multiscale algorithm for dense optimal transport," *Journal of Mathematical Imaging and Vision*, vol. 56, no. 2, pp. 238–259, 2016.
- [65] G. Terzakis, M. Lourakis, and D. Ait-Boudaoud, "Modified rodrigues parameters: An efficient representation of orientation in 3d vision and graphics," *Journal of Mathematical Imaging and Vision*, vol. 60, no. 3, pp. 422–442, 2018.
- [66] Y. He and J. J. Jonas, "Representation of orientation relationships in Rodrigues–Frank space for any two classes of lattice," *Journal of Applied Crystallography*, vol. 40, no. 3, pp. 559–569, Jun 2007.
- [67] B. Riaz, Y. Karahan, and A. J. Brockmeier, "Partial optimal transport for support subset selection," *Transactions on Machine Learning Research*, 2023. [Online]. Available: <https://openreview.net/forum?id=75CcOPxlr>
- [68] H. E. Boyer and T. L. Gall, "Metals handbook; desk edition," 1985.
- [69] S. Samothrakitis, C. B. Larsen, R. Woracek, L. Heller, J. Kopeček, G. Gerstein, H. J. Maier, M. Rameš, M. Tovar, P. Šittner *et al.*, "A multiscale study of hot-extruded coniga ferromagnetic shape-memory alloys," *Materials & Design*, vol. 196, p. 109118, November 2020.
- [70] G. Wahba, "A least squares estimate of satellite attitude," *SIAM Review*, vol. 7, no. 3, pp. 409–409, 1965. [Online]. Available: <https://doi.org/10.1137/1007077>
- [71] A. M.-C. So, *Improved Approximation Bound for Quadratic Optimization Problems with Orthogonality Constraints*, pp. 1201–1209. [Online]. Available: <https://epubs.siam.org/doi/abs/10.1137/1.9781611973068.130>
- [72] H. Liu, W. Wu, and A. M.-C. So, "Quadratic optimization with orthogonality constraints: Explicit lojasiewicz exponent and linear convergence of line-search methods," in *International Conference on Machine Learning*. PMLR, 2016, pp. 1158–1167.
- [73] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, "An analysis of approximations for maximizing submodular set functions—i," *Mathematical Programming*, vol. 14, no. 1, pp. 265–294, 1978.
- [74] U. Stuhr, M. Grosse, and W. Wagner, "The tof-strain scanner poldi with multiple frame overlap—concept and performance," *Materials Science and Engineering: A*, vol. 437, no. 1, pp. 134–138, 2006.