

Atomic Super-Resolution Tomography

Poulami Somanya Ganguly^{1,2}, Felix Lucka^{1,3}, Hermen Jan Hupkes², and Kees Joost Batenburg^{1,2}

¹ Centrum Wiskunde & Informatica, Science Park, Amsterdam, The Netherlands

² The Mathematical Institute, Leiden University, Leiden, The Netherlands

³ Centre for Medical Image Computing, University College London, London, United Kingdom

{poulami.ganguly, felix.lucka, joost.batenburg}@cwi.nl
hhupkes@math.leidenuniv.nl

Abstract. We consider the problem of reconstructing a nanocrystal at atomic resolution from electron microscopy images taken at a few tilt angles. A popular reconstruction approach called discrete tomography confines the atom locations to a coarse spatial grid, which is inspired by the physical *a priori* knowledge that atoms in a crystalline solid tend to form regular lattices. Although this constraint has proven to be powerful for solving this very under-determined inverse problem in many cases, its key limitation is that, in practice, defects may occur that cause atoms to deviate from regular lattice positions. Here we propose a grid-free discrete tomography algorithm that allows for continuous deviations of the atom locations similar to super-resolution approaches for microscopy. The new formulation allows us to define atomic interaction potentials explicitly, which results in a both meaningful and powerful incorporation of the available physical *a priori* knowledge about the crystal's properties. In computational experiments, we compare the proposed grid-free method to established grid-based approaches and show that our approach can indeed recover the atom positions more accurately for common lattice defects.

Keywords: electron tomography · discrete tomography · mathematical super-resolution · molecular dynamics · crystallographic defects.

1 Introduction

Electron tomography is a powerful technique for resolving the interior of nanomaterials. After preparing a microscopic sample, a series of projection images (so called tilt-series) is acquired by rotating the specimen in the electron microscope, acquiring data from a range of angles. In recent years, electron tomography has been successfully applied to reconstruct the 3D positions of the individual atoms in nanocrystalline materials [14, 28, 30].

Since the first demonstration of atomic resolution tomography of nanocrystals in 2010 by discrete tomography [29], a range of tomographic acquisition techniques and reconstruction algorithms have been applied to reconstruct nanocrystals of increasing complexity. In the discrete tomography approach, atoms are

assumed to lie on a regular lattice and the measured projections can be considered as atom counts along lattice lines. A key advantage of this approach is its ability to exploit the constraints induced by the discrete domain and range of the image. As a consequence, a small number of projection angles (typically less than 5) can already lead to an accurate reconstruction [7, 8]. The theoretical properties of the discrete reconstruction problem have been studied extensively with results on algorithm complexity, uniqueness, and stability [4, 6, 18]. A key drawback of the discrete lattice assumption when considering real-world applications to nanocrystal data is that in many interesting cases the atoms do not lie on a perfect lattice due to defects in the crystal structure or interfaces between different crystal lattices. In such cases the atoms do not project perfectly into columns, forming a mismatch with the discrete tomography model.

As an alternative, it has been demonstrated that a more conventional tomographic series consisting of hundreds of projections of a nanocrystal can be acquired in certain cases. An image of the nanocrystal is then reconstructed using sparsity based reconstruction techniques on a continuous model of the tomography problem. This approach does not depend on the lattice structure and allows one to reconstruct defects and interfaces [20]. As a downside, the number of required projections is large and to accurately model the atom positions the reconstruction must be represented on a high-resolution pixel grid resulting in a large-scale computational problem. This raises the question if a reconstruction problem can be defined that fills the gap between these two extremes and can exploit the discrete nature of the lattice structure while at the same time allowing for continuous deviations of atom positions from the perfect lattice.

In this paper we propose a model for the atomic resolution tomography problem that combines these two characteristics. Inspired by the algorithm proposed in [11], our model is based on representing the crystal image as a superposition of delta functions with continuous coordinates and exploiting sparsity of the image to reduce the number of required projections. We show that by incorporating a physical model for the potential energy of the atomic configuration, the reconstruction results can be further improved.

2 Problem setting

In this section we formulate a mathematical model of the atomic resolution tomography problem and discuss several approaches to solve it. Some of these approaches assume that the atom locations are restricted to a perfect grid, the *crystal lattice*, which corresponds to only one possible local minimum of the potential energy of the atomic configuration. To overcome certain limitations of this assumption, we propose an alternative formulation where the atom locations are allowed to vary continuously and an explicit model of the potential energy of their configuration is used to regularise the image reconstruction.

An atomic configuration is characterised by a positive measure μ on a bounded subset X of \mathbb{R}^d . We denote the space of such measures by $\mathcal{M}(X)$. The measure represents the *electron density*, which is the probability that an electron is present

at a given location. The electron density around an atomic configuration is highest in regions where atoms are present. In electron tomography, electron density is probed by irradiating a sample with a beam of electrons. The beam undergoes absorption and scattering due to its interactions with the electrons of the atomic configuration. The transmitted or scattered signal can then be used to form an image. The Radon transform provides a simplified mathematical model of this ray-based image formation process. For $d = 2$, the Radon transform $\mathcal{R}\mu$ can be expressed as integrals taken over straight rays

$$\mathcal{R}[\mu](r, \theta) := \int_{l(r, \theta)} d\mu, \quad (1)$$

$$l(r, \theta) = \{(x_1, x_2) \in \mathbb{R}^2 \mid x_1 \cos \theta + x_2 \sin \theta = r\}, \quad (2)$$

where we parametrised the rays by the projection angle θ and the distance on the detector r . The corresponding inverse problem is to recover μ from noisy observations of $y = \mathcal{R}\mu + \varepsilon$. One way to formulate a solution to this problem is via the following optimisation over the space of measures:

$$\underset{\mu \in \mathcal{M}(X)}{\text{minimise}} \quad \|\mathcal{R}\mu - y\|_2^2, \quad (3)$$

which is an infinite dimensional non-negative linear least-squares problem. In the following, we will introduce a series of discretisations of this optimisation problem. Numerical schemes to solve them will be discussed in Section 3.

In situations where we only have access to data from a few projection angles, introducing a suitable discretisation of (3) is essential for obtaining a stable reconstruction. One way to achieve this is to restrict the atom locations to a spatial grid with n nodes, $\mathbf{x}_{i=1}^n$, and model their interaction zone with the electron beam by a Gaussian with known shape G . The atom centres are then delta peaks $\delta_{\mathbf{x}_i}$ on the gridded image domain. The Gaussian convolution of atom centres can be viewed as the “blurring” produced by thermal motion of atoms. In fact, it is known from lattice vibration theory that, for large configurations, the probability density function of an atom around its equilibrium position is a Gaussian, whose width depends on temperature, dimensionality and interatomic forces [23]. The discretised measure μ can then be written as

$$\mu_{\text{grid}} = \sum_{i=1}^n w_i (G * \delta_{\mathbf{x}_i}), \quad (4)$$

where n is the total number of grid points and weights $w_i \geq 0$ were introduced to indicate confidence in the presence or absence of an atom at grid location i . If we insert (4) in (3) and introduce the forward projection of a single atom as $\psi_i := \mathcal{R}(G * \delta_{\mathbf{x}_i})$ we get

$$\begin{aligned} \|\mathcal{R}\mu_{\text{grid}} - y\|_2^2 &= \|\mathcal{R} \sum_{i=1}^n w_i (G * \delta_{\mathbf{x}_i}) - y\|_2^2 = \\ &= \left\| \sum_{i=1}^n w_i \mathcal{R}(G * \delta_{\mathbf{x}_i}) - y \right\|_2^2 =: \left\| \sum_{i=1}^n \psi_i w_i - y \right\|_2^2 =: \|\Psi w - y\|_2^2 \end{aligned}$$

The corresponding optimisation problem is given as

$$\min_{w \in \mathbb{R}_+^n} \|\Psi w - y\|_2^2, \quad (5)$$

which is a finite dimensional linear non-negative least squares problem.

The choice of the computational grid in (4) is unfortunately not trivial. Only in certain situations, one can assume that all atoms lie on a lattice of known grid size and orientation, and directly match this lattice with the computational grid. In general, one needs to pick a computational grid of much smaller grid size. With the data y given, the grid admits multiple solutions of (5) and most efficient computational schemes tend to pick a blurred, artefact-ridden solution with many non-zero weights far from the true, underlying μ , as we will demonstrate in Section 4. To obtain a better reconstruction, one can choose to add *sparsity constraints* which embed our physical *a priori* knowledge that μ originates from a discrete configuration of atoms. In our model (4), this corresponds to a $w \in \mathbb{R}_+^n$ with few non-zero entries. To obtain such a sparse solution we can add a constraint on ℓ^0 norm of the weights to the optimisation problem:

$$\begin{aligned} & \underset{w \in \mathbb{R}_+^n}{\text{minimise}} && \|\Psi w - y\|_2^2 \\ & \text{subject to} && |w|_0 \leq K. \end{aligned}$$

However, this problem is NP-hard [16]. A approximate solution can be found by replacing the ℓ^0 norm with the ℓ^1 norm and adding it to the objective function:

$$\underset{w \in \mathbb{R}_+^n}{\text{minimise}} \|\Psi w - y\|_2^2 + \lambda \|w\|_1, \quad (6)$$

where λ is the relative weight of the sparsity-inducing term. This particular choice of formulation is not always best and alternative formulations of the same problem exist [16].

For atomic configurations where only one type of atom is present, the weights can be considered to be one where an atom is present and zero everywhere else. This corresponds to discretising the range of the reconstructed image. The fully discrete optimisation problem then becomes:

$$\underset{w \in \{0,1\}^n}{\text{minimise}} \|\Psi w - y\|_2^2. \quad (7)$$

With image range discretisation, a constraint on the number of atoms is typically no longer needed because adding an additional atom with weight 1 after all atoms have been found leads to an increase in the objective function.

Although the optimisation problems (5), (6) and (7) allow for the recovery of atomic configurations without solving (3), all of them rely crucially on discretisation of the domain of the reconstructed image, i.e. the assumption that atoms lie on a grid. However, this assumption is not always true. In particular, atomic configurations often contain defects where atom positions deviate from the perfect lattice. Fig. 1 shows examples of common lattice defects. In




Fig. 1. Atomic configurations with (a) an interstitial point defect, (b) a vacancy and (c) an edge dislocation.

order to resolve these defects correctly, the image domain must be discretised to higher resolutions, i.e. the grid of possible atom positions must be made finer. This introduces two main problems: First, making the grid finer for the same data makes the inverse problem more ill-posed. Second, the computational time increases significantly even for modestly sized configurations.

In order to overcome these difficulties, we revisit (4) and remove the requirement for \mathbf{x}_i to lie on a grid. The projection of a single atom now becomes a function of its location $\mathbf{x} \in \mathbb{R}^d$, $\psi(\mathbf{x}) := \mathcal{R}(G * \delta_{\mathbf{x}})$. We keep the image range discretisation introduced above by requiring $w_i \in \{0, 1\}$. Now, (7) becomes

$$\underset{\mathbf{x} \in X^n, w \in \{0, 1\}^n}{\text{minimise}} \quad \left\| \sum_{i=1}^n w_i \psi(\mathbf{x}_i) - y \right\|_2^2. \quad (8)$$

The minimisation over \mathbf{x} is a non-linear, non-convex least-squares problem which has been studied extensively in the context of mathematical super-resolution [11–13]. In these works, efficient algorithms are derived from relating it back to the infinite dimensional linear least-squares problem on the space of measures (3). For instance, for applications such as fluorescence microscopy [11] and ultrasound imaging [3], an alternating descent conditional gradient (ADCG) algorithm has been proposed, which we will revisit in the next section. Compared to these works, we have a more complicated non-local and under-determined inverse problem and the minimisation over w adds a combinatorial, discrete flavor to (8). To further tailor it to our specific application, we will incorporate physical *a priori* knowledge about atomic configurations of crystalline solids by adding a functional formed by the atomic interaction potentials. This will act as a regularisation of the underlying under-determined inverse problem.

2.1 Potential energy of the atomic configuration

The total energy of an atomic configuration is the sum of its potential energy and kinetic energy. As we consider only static configurations, the kinetic energy




Fig. 2. The normalised Lennard-Jones pair potential as a function of normalised interatomic separation.

of the configuration is zero and the total energy is equal to the potential energy. In order to compute the potential energy of the atomic configuration, we must prescribe the interaction between atoms. In this paper, we use the Lennard-Jones pair potential, which is a simplified model of interatomic interactions. The Lennard-Jones potential V_{LJ} as a function of interatomic separation r is given by

$$V_{LJ}(r) = \begin{cases} 4\epsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^6 \right], & r < r_{\text{cut}} \\ 0, & r \geq r_{\text{cut}} \end{cases} \quad (9)$$

where ϵ is the depth of the potential well and σ is the interatomic separation at which the potential is zero. The separation at which the potential reaches its minimum is given by $r_m = 2^{1/6}\sigma$. The parameter r_{cut} denotes a cut-off separation beyond which the potential is inactive. Fig. 2 shows the form of the the Lennard-Jones pair potential as a function of interatomic separation. The potential energy of the atomic configuration is computed by summing over the pairwise interaction between all pairs of atoms

$$V_{\text{tot}}(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N) = \sum_{i>j} V_{LJ}(\mathbf{x}_i - \mathbf{x}_j). \quad (10)$$

Adding this energy to the objective in (8) leads to

$$\underset{\mathbf{x} \in \mathcal{C}, w \in \{0,1\}^n}{\text{minimise}} \quad \left\| \sum_{i=1}^n w_i \psi(\mathbf{x}_i) - y \right\|_2^2 + \alpha V_{\text{tot}}(\mathbf{x}). \quad (11)$$

The regularisation parameter, α , adjusts the relative weight of the energy term, so that by tuning it we are able to move between atomic configurations that are data-optimal and those that are energy-optimal. The constraint set $\mathcal{C} \subset X^n$ is defined by a minimum distance r_{\min} , such that $|\mathbf{x}_i - \mathbf{x}_j| > r_{\min}, \forall i > j$. The minimum distance, r_{\min} , is chosen to be smaller than the optimal interatomic separation r_m and allows us to set α to 0 and still avoid configurations where atoms are placed exactly at the same position. For small separations, the energy is dominated by the $(\frac{\sigma}{r})^{12}$ term and increases sharply for separations less than r_m . Thus, for non-zero α , configurations where atoms have a separation less than r_m are highly unlikely.

3 Algorithms

In this section we discuss several algorithms to solve the optimisation problems introduced in Section 2.

Projected gradient descent The non-negative least-squares problem (5) can be solved with a simple iterative first-order optimisation scheme. At each step of the algorithm, the next iterate is computed by moving in the direction of the negative gradient of the objective function. Non-negativity of the weights is enforced by projecting negative iterates to zero. Mathematically, each iterate is given by

$$w^{k+1} = \prod_+ \left(w^k + t\Psi^T(\Psi w^k - y) \right), \quad (12)$$

where t is the step size and the projection operator is given by

$$\prod_+ (\cdot) = \max(\cdot, 0). \quad (13)$$

In the numerical experiments in Section 4, we used the SIRT algorithm [22] as implemented in the tomographic reconstruction library ASTRA [1], which is based on a minor modification of the iteration described above.

Proximal gradient descent If we add the non-smooth ℓ^1 regulariser and obtain problem (6), we need to extend (12) to a proximal gradient scheme [26]

$$w^{k+1} = \text{prox}_h \left(w^k + t\Psi^T(\Psi w^k - y) \right), \quad (14)$$

where the projection operator (13) is replaced by the proximal operator of the convex function

$$h(x) := \begin{cases} \lambda \|x\|_1 & x \geq 0 \\ 0 & \text{elsewhere} \end{cases}, \quad (15)$$

which is given by the non-negative soft-thresholding operator

$$\text{prox}_h(x) = \begin{cases} x - \lambda, & x \geq \lambda \\ 0, & \text{elsewhere} \end{cases}.$$

In the numerical experiments in Section 4, we used the fast iterative soft-thresholding algorithm (FISTA) [9] as implemented in the Python library ODL [2], which is based on a slight modification of the iteration described above.

Algorithm 1 Discrete simulated annealing

```

1: while  $\beta < \beta_{max}$  do
2:   Select new atom location:  $\tilde{w}^k \in \arg \min_{k \in \mathcal{C}} \Psi w^k - y$ 
3:   Add new atom to current configuration:  $\tilde{w}^{k+1} \leftarrow \{w^k, \tilde{w}^k\}$ 
4:   Accept new configuration with a certain probability:
5:   if  $\beta \|\Psi \tilde{w}^{k+1} - y\|_2^2 < \beta \|\Psi w^k - y\|_2^2$  then
6:      $w^{k+1} \leftarrow \tilde{w}^{k+1}$ 
7:   else
8:     Generate random number:  $t \in \text{rand}[0, 1)$ 
9:     if  $t < e^{-\beta \|\Psi \tilde{w}^{k+1} - y\|_2^2} / e^{-\beta \|\Psi w^k - y\|_2^2}$  then
10:       $w^{k+1} \leftarrow \tilde{w}^{k+1}$ 
11:    end if
12:  end if
13:  Move atom:  $w^{k+1} \leftarrow \text{random move}(w^{k+1})$ 
14:  Repeat acceptance steps 5–13
15:  Increase  $\beta$ 
16: end while

```

Simulated annealing For solving the fully discrete problem (7), we used a simulated annealing algorithm as shown in Algorithm 1, which consists of two subsequent accept-reject steps carried out with respect to the same inverse temperature parameter β . In the first one, the algorithm tries to add a new atom to the existing configuration. In the second one, the atom locations are perturbed locally. As β is increased towards β_{max} , fewer new configurations are accepted and the algorithm converges to a minimum.

In the atom adding step at each iteration k , the algorithm tries to add an atom at one of the grid location i where the residual $\Psi w^k - y$ is minimal (this corresponds to flipping w_i^k from 0 to 1 in (7)). The allowed grid locations belong to a constraint set \mathcal{C} , such that no two atoms are closer than a pre-specified minimum distance r_{\min} . To perturb the atom positions locally, the algorithm selects an atom at random and moves it to one of its 4 nearest neighbour locations at random.

ADCG with energy Variants of the Frank-Wolfe algorithm (or conditional gradient method) [17, 21] have been proposed for solving problems of the form (8) [3, 15] without discrete constraints for w and are commonly known as alternating descent conditional gradient (ADCG) schemes (see [27] for an analysis specific to multidimensional sparse inverse problems). Here, we modify the ADCG scheme to

1. incorporate binary constraints on w

Algorithm 2 ADCG with energy

```

1: for  $k = 1 : k_{max}$  do
2:   Compute next atom in grid  $g$ :
      $\mathbf{x}_{new} \in \arg \min_{\mathbf{x}_{new} \in g, (\mathbf{x}^k, \mathbf{x}_{new}) \in \mathcal{C}} \|\sum_{i=1}^k \psi(\mathbf{x}_i) - y + \psi(\mathbf{x}_{new})\| + \alpha V_{tot}(\mathbf{x}^k, \mathbf{x}_{new})$ 
3:   Update support:  $\mathbf{x}^{k+1} \leftarrow \{\mathbf{x}^k, \mathbf{x}_{new}\}$ 
4:   Locally move atoms:
      $\mathbf{x}^{k+1} \leftarrow \min_{\mathbf{x} \in X} \|\Psi\mu(\mathbf{x}^{k+1}) - y\|_2^2 + \alpha V(\mathbf{x}^{k+1})$ 
5:   Break if objective function is increasing:
6:   if  $\|\Psi\mu(\mathbf{x}^{k+1}) - y\|_2^2 + \alpha V(\mathbf{x}^{k+1}) > \|\Psi\mu(\mathbf{x}^k) - y\|_2^2 + \alpha V(\mathbf{x}^k)$  then break
7:   end if
8: end for

```

2. handle the singularities of the atomic interaction potentials
3. avoid local minima resulting from poor initialisations

The complete algorithm is shown in Algorithm 2. Essentially, the scheme also alternates between adding a new atom to the current configuration and optimising the positions of the atoms.

In the first step, the image domain is coarsely gridded and the objective function after adding an atom at each location is computed. Locations closer to existing atoms than r_{\min} are excluded. In the second step, the atom coordinates are optimised by a continuous local optimisation method. Here, the Nelder-Mead method [24] implemented in SciPy [31] was used.

A continuation strategy is used to avoid problems resulting from poor initialisations: Algorithm 2 is run for increasing values of α , starting from $\alpha = 0$. The reconstruction obtained at the end of a run is used as initialisation for the next. In the following section, we demonstrate the effect of increasing α on the reconstructions obtained and discuss how an optimal α was selected. In the following section, we refer to Algorithm 2 as ‘‘ADCG’’ when used for $\alpha = 0$ and as ‘‘ADCG with energy’’ otherwise.

4 Numerical experiments

We conducted numerical experiments by creating 2D atomic configurations with defects and using the algorithms discussed in Section 3 to resolve atom positions. In this section we describe how the ground truth configurations were generated and projected, and compare the reconstruction results of different algorithms.

4.1 Ground truth configurations

We generated ground truth configurations using the molecular dynamics software HOOMD-blue [5, 19]. We created perfect square lattices and then induced defects by adding or removing atoms. The resulting configuration was then relaxed to an energy minimum using the FIRE energy minimiser [10] to give the configurations shown in Fig. 1. The following parameter values were used in (9) for specifying the Lennard-Jones pair potential between atoms.

Defect type	ϵ	σ	r_{cut}
Interstitial defect	0.4	0.15	0.4
Vacancy	0.4	0.14	0.4
Edge dislocation	0.4	0.13	0.17

4.2 Discretised projection data

We generated two 1D projections for each ground truth atomic configuration at projection angles $\theta = 0^\circ, 90^\circ$. As discussed in Section 2, the projection of a single atom centre is given by a Gaussian convolution followed by the Radon transform. The Radon transform of a Gaussian is also a Gaussian. Therefore, we interchanged the two operations in the forward transform to speed up the computations. The sum over individual projections of atom centres was used as the total (noise-free) projection. Using the Radon transform in (1), each atom centre was projected onto a 1D detector, following which it was convolved with a 1D Gaussian of the form $G(z) = e^{-(z-z_0)/\varsigma^2}$, where z_0 is the position of the atom centre on the detector and ς controls the width of the Gaussian. Finally, the continuous projection was sampled at a fixed number of points to give rise to a *discrete* projection. For our experiments, the ς of the Gaussian function was taken to be equal to the discretisation of the detector given by the detector pixel size d . Both were taken to be 0.01.

4.3 Discretisation of reconstruction volume

For SIRT, FISTA and simulated annealing, each dimension of the reconstruction area was discretised using the detector pixel size d . Therefore, there were $1/d \times 1/d$ grid points in total.

Gridding is required for Algorithm 2 at the atom adding step. We found that a coarse discretisation, with less than $1/9^{\text{th}}$ the number of grid points, was already sufficient.

4.4 Comparison between reconstructions

The reconstructions obtained with the different algorithms are shown in Fig. 3. For each reconstruction, data from only two projections were used. Note that two projections is far from sufficient for determining the correct atomic configuration and several different configurations have the same data discrepancy.

In the SIRT reconstructions, atom positions were blurred out and none of the defects were resolved. In all cases, the number of intensity peaks was also different from the true number of atoms. Although FISTA reconstructions, which include sparsity constraints on the weights, were less blurry, atoms still occupied more than one pixel. For both these algorithms, additional heuristic post-processing is required to output atom locations. In the edge dislocation case, both algorithms

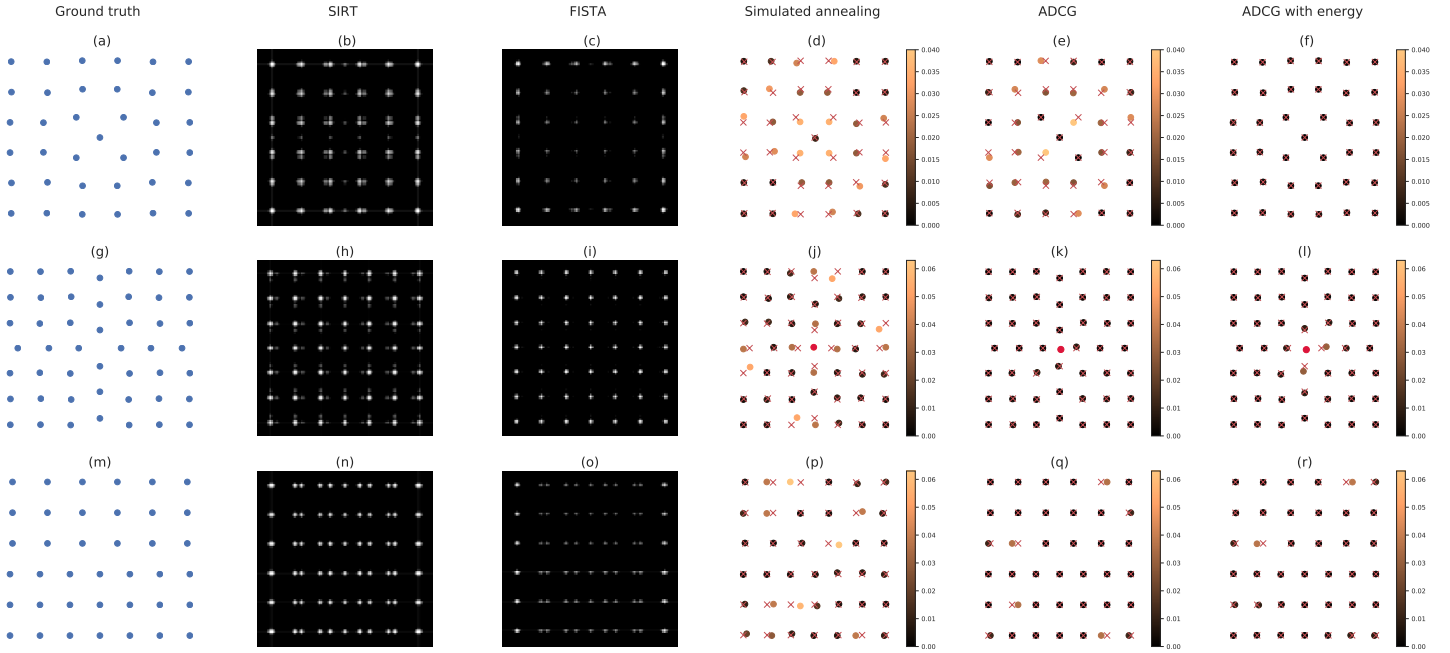


Fig. 3. Reconstructions of atomic configurations with (a)–(f) an interstitial point defect, (g)–(l) a vacancy and (m)–(r) an edge dislocation from two projections. For the simulated annealing, ADCG and ADCG with energy reconstructions, atoms are coloured according to their Euclidean distance from the ground truth. The ground truth positions are marked with red crosses. In (j)–(l) an extra atom (shown in red) was present in the reconstructions but not in the ground truth.

gave rise to a configuration with many more atoms than were present in the ground truth.

The discrete simulated annealing algorithm performed better for all configurations. For the interstitial point defect and edge dislocation, the number of atoms in the reconstruction matched that in the ground truth. The positions of most atoms, however, were not resolved correctly. Moreover, the resolution, like in previous algorithms, was limited to the resolution on the detector. We ran the simulated annealing algorithm for comparable times as the ADCG algorithms and picked the solution with the least data discrepancy.

Already the ADCG algorithm for $\alpha = 0$ performed far better than all the previous algorithms. For the configurations with an interstitial point defect and edge dislocation, all but a few atom locations were identified correctly. For the




Fig. 4. Reconstructions of a vacancy defect from three projections. For the simulated annealing, ADCG and ADCG with energy reconstructions, atoms are coloured according to their Euclidean distance from the ground truth. Ground truth positions are marked with red crosses.

configuration with a vacancy, all atoms were correctly placed. However, an additional atom at the centre of the configuration was placed incorrectly.

Adding the potential energy (ADCG with energy) helps to resolve atom positions that were not identified with $\alpha = 0$. For the interstitial point defect and edge dislocation, these reconstructions were the closest to the ground truth. Adding the energy to the configuration with a vacancy moved the atoms near the defect further apart but was not able to correct for the extra atom placed. For this case, we performed an additional experiment with three projections (at 0° , 45° and 90°), the results of which are shown in Fig. 4. The defect was still not resolved in the SIRT and FISTA reconstructions. However, the number of atoms in the simulated annealing, ADCG and ADCG with energy reconstructions was correct. Once again, the reconstruction obtained with our algorithm was closer to the ground truth than all other reconstructions, with all but one atom placed correctly.

4.5 Effect of adding energy to optimisation

To resolve atom positions using Algorithm 2, the contribution of the potential energy was increased gradually by increasing α . In Fig. 5, we show the effect of adding energy to the optimisation problem. For $\alpha = 0$, an initial guess for the true configuration was obtained. This configuration, though data optimal, was not the ground truth. A quantitative measure of this mismatch is the Euclidean distance between the reconstructed atom locations and those in the ground truth. As α was increased, the reconstructions evolved from being data-optimal to being energy-optimal. At a certain value of α , the Euclidean distance between reconstructed and ground truth atom locations decreased to zero. Increasing α beyond this point led to a large increase in the data discrepancy term due to the addition of more atoms. For very high values of α , the configurations obtained were essentially global minima of the potential energy, such as the honeycomb configuration in Fig. 5(e) for $\alpha = 100.0$. An optimal value of the regularisation

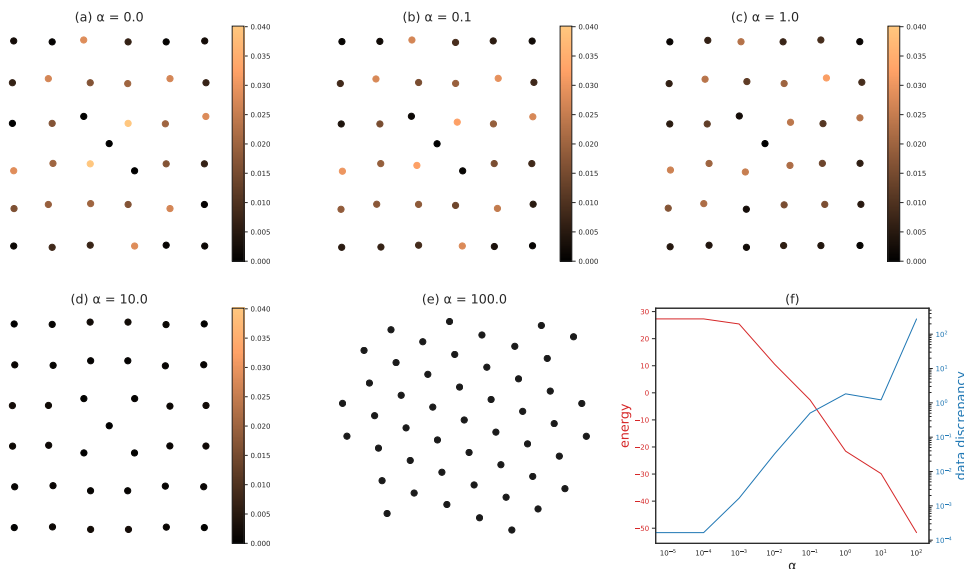


Fig. 5. (a)-(d): Increasing the weighting of the energy term from $\alpha = 0.0$ to $\alpha = 10.0$ helps to resolve the correct atomic configuration. The reconstructed atoms are coloured according to their Euclidean distance from the atoms in the ground truth. (e) At high values of α , the reconstructions have a high data discrepancy and correspond to one of the global minima of the potential energy. (f) From the plots of potential energy and data discrepancy, an optimal value of α can be selected.

parameter was selected by increasing α to the point at which more atoms were added to the configuration and a jump in the data discrepancy was observed.

5 Discussion

The results of our numerical experiments demonstrate that algorithms like ADCG, which do not rely on domain discretisation, are better at resolving the defects in the atomic configurations shown in Fig. 1. Moreover, the output from ADCG is a *list of coordinates* and not an image like that of SIRT or FISTA, which requires further post-processing steps to derive the atom locations. Direct access to coordinates can be particularly useful because further analysis, such as strain calculations, often require atom positions as input.

Adding the potential energy of the atomic configuration to the optimisation problem resulted in reconstructions that were closer to the ground truth. One challenge of the proposed approach (with or without adding the potential energy)

is that the resulting optimisation problem is a non-convex function of the atom locations. The numerical methods we presented are not intentionally designed to escape local minima and are therefore sensitive to their initialisation. To improve this, one important extension would be to also remove atoms from the current configuration, which might make it possible to resolve the vacancy defect in Fig. 3 with two projections. More generally, one would need to include suitable features of global optimisation algorithms [25].

6 Conclusions

In this paper we proposed a novel discrete tomography approach in which the locations of atoms are allowed to vary continuously and their interaction potentials are modelled explicitly. We showed in proof-of-concept numerical studies that such an approach can be better at resolving crystalline defects than image domain discretised or fully discrete algorithms. Furthermore, in situations where atom locations are desired, this approach provides access to the quantity of interest without any additional post-processing. For future work, we will extend our numerical studies on this atomic super-resolution approach to larger-sized scenarios in 3D, featuring realistic measurement noise, acquisition geometries, more suitable and accurate physical interaction potentials and different atom types. This will require additional computational effort to scale up our algorithm and will then allow us to work on real electron tomography data of nanocrystals.

References

1. van Aarle, W., Palenstijn, W.J., De Beenhouwer, J., Altantzis, T., Bals, S., Batenburg, K.J., Sijbers, J.: The astra toolbox: A platform for advanced algorithm development in electron tomography. *Ultramicroscopy* **157**, 35–47 (2015)
2. Adler, J., Kohr, H., Oktem, O.: Operator discretization library (odl). Software available from <https://github.com/odlgroup/odl> (2017)
3. Alberti, G.S., Ammari, H., Romero, F., Wintz, T.: Dynamic spike superresolution and applications to ultrafast ultrasound imaging. *SIAM Journal on Imaging Sciences* **12**(3), 1501–1527 (2019)
4. Alpers, A., Gritzmann, P.: On stability, error correction, and noise compensation in discrete tomography. *SIAM Journal on Discrete Mathematics* **20**(1), 227–239 (2006)
5. Anderson, J.A., Lorenz, C.D., Travesset, A.: General purpose molecular dynamics simulations fully implemented on graphics processing units. *Journal of computational physics* **227**(10), 5342–5359 (2008)
6. Baake, M., Huck, C., Gritzmann, P., Langfeld, B., Lord, K.: Discrete tomography of planar model sets. *Acta Crystallographica A* **62**(6), 419–433 (2006)
7. Batenburg, K.J.: A network flow algorithm for reconstructing binary images from discrete x-rays. *Journal of Mathematical Imaging and Vision* **27**(2), 175–191 (2006)
8. Batenburg, K.J., Sijbers, J.: Generic iterative subset algorithms for discrete tomography. *Discrete Applied Mathematics* **157**(3), 438–451 (2009)
9. Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences* **2**(1), 183–202 (2009)

10. Bitzek, E., Koskinen, P., Gähler, F., Moseler, M., Gumbsch, P.: Structural relaxation made simple. *Physical review letters* **97**(17), 170201 (2006)
11. Boyd, N., Schiebinger, G., Recht, B.: The alternating descent conditional gradient method for sparse inverse problems. *SIAM Journal on Optimization* **27**(2), 616–639 (2017)
12. Bredies, K., Pikkarainen, H.K.: Inverse problems in spaces of measures. *ESAIM: Control, Optimisation and Calculus of Variations* **19**(1), 190–218 (2013)
13. Candès, E.J., Fernandez-Granda, C.: Towards a mathematical theory of super-resolution. *Communications on pure and applied Mathematics* **67**(6), 906–956 (2014)
14. Chen, C.C., Zhu, C., White, E.R., Chiu, C.Y., Scott, M., Regan, B., Marks, L.D., Huang, Y., Miao, J.: Three-dimensional imaging of dislocations in a nanoparticle at atomic resolution. *Nature* **496**(7443), 74–77 (2013)
15. Denoyelle, Q., Duval, V., Peyré, G., Soubies, E.: The sliding frank-wolfe algorithm and its application to super-resolution microscopy. *Inverse Problems* (2019)
16. Foucart, S., Rauhut, H.: A mathematical introduction to compressive sensing. *Applied and Numerical Harmonic Analysis*, Birkhäuser/Springer, New York (2013)
17. Frank, M., Wolfe, P.: An algorithm for quadratic programming. *Naval research logistics quarterly* **3**(1-2), 95–110 (1956)
18. Gardner, R.J., Gritzmann, P.: Discrete tomography: Determination of finite sets by x-rays. *Transactions of the American Mathematical Society* **349**(6), 2271–2295 (1997)
19. Glaser, J., Nguyen, T.D., Anderson, J.A., Lui, P., Spiga, F., Millan, J.A., Morse, D.C., Glotzer, S.C.: Strong scaling of general-purpose molecular dynamics simulations on gpus. *Computer Physics Communications* **192**, 97–107 (2015)
20. Goris, B., De Beenhouwer, J., De Backer, A., Zanaga, D., Batenburg, K.J., Sanchez-Iglesias, A., Liz-Marzan, L.M., Van Aert, S., Bals, S., Sijbers, J., et al.: Measuring lattice strain in three dimensions through electron microscopy. *Nano letters* **15**(10), 6996–7001 (2015)
21. Jaggi, M.: Revisiting frank-wolfe: Projection-free sparse convex optimization. In: *ICML* (1). pp. 427–435 (2013)
22. Kak, A.C., Slaney, M., Wang, G.: Principles of computerized tomographic imaging. *Medical Physics* **29**(1), 107–107 (2002)
23. Montroll, E.W.: Theory of the vibration of simple cubic lattices with nearest neighbor interactions. In: *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*. vol. 3, pp. 209–246. Univ of California Press (1956)
24. Nelder, J.A., Mead, R.: A simplex method for function minimization. *The computer journal* **7**(4), 308–313 (1965)
25. Pardalos, P.M., Zhigljavsky, A., Žilinskas, J.: *Advances in stochastic and deterministic global optimization*. Springer (2016)
26. Parikh, N., Boyd, S., et al.: Proximal algorithms. *Foundations and Trends® in Optimization* **1**(3), 127–239 (2014)
27. Poon, C., Peyré, G.: Multidimensional sparse super-resolution. *SIAM Journal on Mathematical Analysis* **51**(1), 1–44 (2019)
28. Rez, P., Treacy, M.M.: Three-dimensional imaging of dislocations. *Nature* **503**(7476), E1–E1 (2013)
29. Van Aert, S., Batenburg, K.J., Rossell, M.D., Erni, R., Van Tendeloo, G.: Three-dimensional atomic imaging of crystalline nanoparticles. *Nature* **470**(7334), 374 (2011)
30. Van Dyck, D., Jinschek, J.R., Chen, F.R.: ‘big bang’ tomography as a new route to atomic-resolution electron tomography. *Nature* **486**(7402), 243–246 (2012)

31. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al.: Scipy 1.0—fundamental algorithms for scientific computing in python. arXiv preprint arXiv:1907.10121 (2019)