



Universiteit
Leiden
The Netherlands

Sparsity-based algorithms for inverse problems

Ganguly, P.S.

Citation

Ganguly, P. S. (2022, December 8). *Sparsity-based algorithms for inverse problems*. Retrieved from <https://hdl.handle.net/1887/3494260>

Version: Publisher's Version

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/3494260>

Note: To cite this publication please use the final published version (if applicable).

Chapter 3

Sparse grid-free reconstruction of nanocrystal defects

3.1 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 [61]–[63].

Since the first demonstration of atomic resolution tomography of nanocrystals in 2010 by discrete tomography [64], 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 [16], [17]. The theoretical properties of the discrete reconstruction problem have been studied extensively with results on algorithm complexity, uniqueness, and stability [65]–[67]. 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

This chapter is based on:

Atomic Super-resolution Tomography. *P. S. Ganguly, F. Lucka, H. J. Hupkes, K. J. Batenburg.*
International Workshop on Combinatorial Image Analysis. Springer, Cham, pp. 45-61, 2020.

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 [18]. 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 chapter we propose a model for the atomic resolution tomography problem that combines these two characteristics. Inspired by the algorithm proposed in [12], 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.

3.2 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 regularize the image reconstruction.

An atomic configuration is characterized 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

$$\mathcal{R}[\mu](r, \theta) := \int_{l(r, \theta)} d\mu, \quad (3.1)$$

$$l(r, \theta) = \{(x_1, x_2) \in \mathbb{R}^2 \mid x_1 \cos \theta + x_2 \sin \theta = r\}, \quad (3.2)$$

where we parametrized 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 optimization over the

space of measures:

$$\underset{\mu \in \mathcal{M}(X)}{\text{minimize}} \quad \|\mathcal{R}\mu - y\|_2^2, \quad (3.3)$$

which is an infinite dimensional non-negative linear least-squares problem. In the following, we will introduce a series of discretizations of this optimization problem. Numerical schemes to solve them will be discussed in Section 3.3.

In situations where we only have access to data from a few projection angles, introducing a suitable discretization of (3.3) 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 [68]. The discretized measure μ can then be written as

$$\mu_{\text{grid}} = \sum_{i=1}^n w_i (G * \delta_{\mathbf{x}_i}), \quad (3.4)$$

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 (3.4) in (3.3) 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 optimization problem is given by

$$\underset{w \in \mathbb{R}_+^n}{\text{minimize}} \quad \|\Psi w - y\|_2^2, \quad (3.5)$$

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

The choice of the computational grid in (3.4) 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 (3.5) 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 3.4. 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 (3.4), this corresponds




Figure 3.1: Atomic configurations with (a) an interstitial point defect, (b) a vacancy and (c) an edge dislocation.

to a $w \in \mathbb{R}_+^n$ with few non-zeros entries. To obtain such a sparse solution we can add a constraint on ℓ^0 norm of the weights to the optimization problem:

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

However, this problem is NP-hard [69]. A approximate solution can be found by replacing the ℓ^0 norm with the ℓ^1 norm and adding it to the objective function:

$$\underset{w \in \mathbb{R}_+^n}{\text{minimize}} \quad \|\Psi w - y\|_2^2 + \lambda \|w\|_1, \quad (3.6)$$

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 [69].

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 discretizing the range of the reconstructed image. The fully discrete optimization problem then becomes:

$$\underset{w \in \{0,1\}^n}{\text{minimize}} \quad \|\Psi w - y\|_2^2. \quad (3.7)$$

With image range discretization, 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 optimization problems (3.5), (3.6) and (3.7) allow for the recovery of atomic configurations without solving (3.3), all of them rely crucially on discretization 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. 3.1 shows examples of common lattice defects. In order to resolve these defects correctly, the image domain must be discretized to higher resolutions, i.e. the grid of possible atom positions




Figure 3.2: The normalized Lennard-Jones pair potential as a function of normalized interatomic separation.

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 (3.4) 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 discretization introduced above by requiring $w_i \in \{0, 1\}$. Now, (3.7) becomes

$$\underset{\mathbf{x} \in X^n, w \in \{0,1\}^n}{\text{minimize}} \quad \left\| \sum_{i=1}^n w_i \psi(\mathbf{x}_i) - y \right\|_2^2. \quad (3.8)$$

The minimization over \mathbf{x} is a non-linear, non-convex least-squares problem which has been studied extensively in the context of mathematical super-resolution [12], [70], [71]. In these works, efficient algorithms are derived from relating it back to the infinite dimensional linear least-squares problem on the space of measures (3.3). For instance, for applications such as fluorescence microscopy [12] and ultrasound imaging [72], 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 minimization over w adds a combinatorial, discrete flavor to (3.8). 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 regularization of the underlying under-determined inverse problem.

3.2.1 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 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 chapter, 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 [73]

$$V_{\text{LJ}}(r) = \begin{cases} 4\epsilon \left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^6 \right], & r < r_{\text{cut}} \\ 0, & r \geq r_{\text{cut}} \end{cases} \quad (3.9)$$

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. 3.2 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

$$V_{\text{tot}}(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N) = \sum_{i>j} V_{\text{LJ}}(\mathbf{x}_i - \mathbf{x}_j). \quad (3.10)$$

Adding this energy to the objective in (3.8) leads to

$$\underset{\mathbf{x} \in \mathcal{C}, w \in \{0,1\}^n}{\text{minimize}} \quad \left\| \sum_{i=1}^n w_i \psi(\mathbf{x}_i) - y \right\|_2^2 + \alpha V_{\text{tot}}(\mathbf{x}). \quad (3.11)$$

The regularization 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 $|\mathbf{x}_i - \mathbf{x}_j| > r_{\text{min}}, \forall i > j$. 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 $\left(\frac{\sigma}{r}\right)^{12}$ 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.

3.3 Algorithms

In this section we discuss several algorithms to solve the optimization problems introduced in Section 3.2.

3.3.1 Projected gradient descent

The non-negative least-squares problem (3.5) can be solved with a simple iterative first-order optimization 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

$$w^{k+1} = \prod_{+} \left(w^k + t\Psi^T(\Psi w^k - y) \right), \quad (3.12)$$

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

$$\prod_{+}(\cdot) = \max(\cdot, 0). \quad (3.13)$$

In the numerical experiments in Section 3.4, we used the SIRT algorithm [29] as implemented in the tomographic reconstruction library ASTRA [7], which is based on a minor modification of the iteration described above.

3.3.2 Proximal gradient descent

If we add the non-smooth ℓ^1 regularizer and obtain problem (3.6), we need to extend (3.12) to a proximal gradient scheme [74]

$$w^{k+1} = \text{prox}_h \left(w^k + t\Psi^T(\Psi w^k - y) \right), \quad (3.14)$$

where the projection operator (3.13) is replaced by the proximal operator of the convex function

$$h(x) := \begin{cases} \lambda \|x\|_1 & x \geq 0 \\ 0 & \text{elsewhere} \end{cases}, \quad (3.15)$$

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 3.4, we used the fast iterative soft-thresholding algorithm (FISTA) [75] as implemented in the Python library ODL [8], which is based on a slight modification of the iteration described above.

3.3.3 Simulated annealing

For solving the fully discrete problem (3.7), we used a simulated annealing algorithm as shown in Algorithm 2, 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

Algorithm 2 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:   Run acceptance steps 5–13
15:   Increase  $\beta$ 
16: end while
```

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 (3.7)). 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 neighbor locations at random.

Algorithm 3 ADCG with energy

```
1: for  $k = 1 : k_{\max}$  do
2:   Compute next atom in grid  $g$ :
    $\mathbf{x}_{\text{new}} \in \arg \min_{\mathbf{x}_{\text{new}} \in g, (\mathbf{x}^k, \mathbf{x}_{\text{new}}) \in \mathcal{C}} \|\sum_{i=1}^k \psi(\mathbf{x}_i) - y + \psi(\mathbf{x}_{\text{new}})\| + \alpha V_{\text{tot}}(\mathbf{x}^k, \mathbf{x}_{\text{new}})$ 
3:   Update support:  $\mathbf{x}^{k+1} \leftarrow \{\mathbf{x}^k, \mathbf{x}_{\text{new}}\}$ 
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
```

3.3.4 ADCG with energy

Variants of the Frank-Wolfe algorithm (or conditional gradient method) [19], [76] have been proposed for solving problems of the form (3.8) [72], [77] without discrete constraints for w and are commonly known as alternating descent conditional gradient (ADCG) schemes (see [78] for an analysis specific to multidimensional sparse inverse problems). Here, we modify the ADCG scheme to

1. incorporate binary constraints on w
2. handle the singularities of the atomic interaction potentials
3. avoid local minima resulting from poor initializations

The complete algorithm is shown in Algorithm 3. Essentially, the scheme also alternates between adding a new atom to the current configuration and optimizing 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 optimized by a continuous local optimization method. Here, the Nelder-Mead method [79] implemented in SciPy [80] was used.

A continuation strategy is used to avoid problems resulting from poor initializations: Algorithm 3 is run for increasing values of α , starting from $\alpha = 0$. The reconstruction obtained at the end of a run is used as initialization 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 3 as “ADCG” when used for $\alpha = 0$ and as “ADCG with energy” otherwise.

3.4 Numerical experiments

We conducted numerical experiments by creating 2D atomic configurations with defects and using the algorithms discussed in Section 3.3 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.

3.4.1 Ground truth configurations

We generated ground truth configurations using the molecular dynamics software HOOMD-blue [81], [82]. 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 minimizer [83] to give the configurations shown in Fig. 3.1. The following parameter values were used in (3.9) 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

3.4.2 Discretized 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 3.2, 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 (3.1), each atom centre was projected onto a 1D detector, following which it was convolved with a 1D Gaussian of the form $G(z) = e^{-(z-z_0)/\zeta^2}$, where z_0 is the position of the atom centre on the detector and ζ 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 ζ of the Gaussian function was taken to be equal to the discretization of the detector given by the detector pixel size d . Both were taken to be 0.01.

3.4.3 Discretization of the reconstruction volume

For SIRT, FISTA and simulated annealing (described in subsections 3.3.1, 3.3.2 and 3.3.3, respectively), each dimension of the reconstruction area was discretized using the detector pixel size d . Therefore, there were $1/d \times 1/d$ grid points in total.

Gridding is required for our variant of ADCG (subsection 3.3.4, Algorithm 3) at the atom adding step. We found that a coarse discretization, with less than $1/9^{\text{th}}$ the number of grid points, was already sufficient.

3.4.4 Comparison between reconstructions

The reconstructions obtained with the different algorithms are shown in Fig. 3.3. 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.




Figure 3.3: 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.




Figure 3.4: 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.

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 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 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

	Interstitial defect		Vacancy (3 projs.)		Edge dislocation	
	Number of atoms	Mean distance	Number of atoms	Mean distance	Number of atoms	Mean distance
Ground truth	37	0.0000	48	0.0000	39	0.0000
SIRT	36	–	49	–	66	–
FISTA	36	–	49	–	66	–
Simulated an-nealing	37	0.0184	48	0.0164	39	0.0159
ADCG	37	0.0138	48	0.0130	39	0.0049
ADCG with energy	37	0.0018	48	0.0024	39	0.0048

Table 3.1: Number of atoms and mean Euclidean distance from ground truth atoms for reconstructions obtained with different algorithms. Thresholding was used to compute the number of atoms detected in the SIRT and FISTA reconstructions.

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° . These results are shown in Fig. 3.4. Taking projections at different angles (e.g. 0° , 22.5° and 90°) did not improve results. 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. Reconstructions with 3 projections for the interstitial point defect and edge dislocation were not significantly different from those with 2 projections. In Table 3.1, we report the number of atoms detected and (where applicable) the mean Euclidean distance of atoms from the ground truth. Note that for computing the mean distance, we required that the number of atoms detected in the reconstruction matched that in the ground truth. Thresholding with a pre-defined minimum distance between peaks was used to detect atoms in the SIRT and FISTA reconstructions.

3.4.5 Effect of adding energy to the optimization

To resolve atom positions using Algorithm 3, the contribution of the potential energy was increased gradually by increasing α . In Fig. 3.5, we show the effect of adding energy to the optimization 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




Figure 3.5: (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 (indicated by the grey line) is selected. Increasing α beyond this optimal value leads to a large increase in the data discrepancy due to addition of more atoms.

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. 3.5(e) for $\alpha = 100.0$. An optimal value of the regularization 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.

3.5 Discussion

The results of our numerical experiments demonstrate that algorithms like ADCG, which do not rely on domain discretization, are better at resolving the defects in the atomic configurations shown in Fig. 3.1. 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 optimization 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 optimization 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 initialization. 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. 3.3 with two projections. More generally, one would need to include suitable features of global optimization algorithms [84] that do not compromise ADCCG's computational efficiency (note that we could have adapted simulated annealing to solve (3.11) but using a cooling schedule slow enough to prevent getting trapped in local minima quickly becomes practically infeasible). A related problem is to characterize local and global minimizers of (3.11) to understand which configurations can be uniquely recovered by this approach and which cannot. To process experimental data, it may furthermore be important to analyze the impact of the error caused by the approximate nature of the mathematical models used for data acquisition (\mathcal{R}, G) and atomic interaction (V_{LJ}).

3.6 Conclusions

In this chapter we proposed a novel discrete tomography approach in which the locations of atoms are allowed to vary continuously and their interaction potentials are modeled explicitly. We showed in proof-of-concept numerical studies that such an approach can be better at resolving crystalline defects than image domain discretized 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.

