Ultramicroscopy 109 (2009) 1485–1490

Contents lists available at ScienceDirect

Ultramicroscopy

journal homepage: www.elsevier.com/locate/ultramic

A model based atomic resolution tomographic algorithm

W. Van den Broek *, S. Van Aert, D. Van Dyck

EMAT, University of Antwerp, Groenenborgerlaan 171, B-2020 Antwerp, Belgium

ARTICLE INFO

Article history:
Received 11 May 2009
Received in revised form
29 July 2009
Accepted 18 August 2009

PACS:
61.05.J–
41.85.Ew
42.30.Va
42.30.Wb

Keywords:
Atomic resolution tomography
HAADF STEM
Maximum likelihood
Algebraic reconstruction technique

ABSTRACT

Tomography with high angular annular dark field scanning transmission electron microscopy at atomic resolution can be greatly improved if one is able to take advantage of prior knowledge. In this paper we present a reconstruction technique that explicitly takes into account the microscope parameters and the atomic nature of the projected object. This results in a more accurate estimate of the atomic positions and in a good resistance to noise. The reconstruction is a maximum likelihood estimator of the object. Moreover, the limits to the precision have been explored, allowing for a prediction of the amount of expected noise in the reconstruction for a certain experimental setup. We believe that the proposed reconstruction technique can be generalized to other tomographic experiments.

1. Introduction

In tomography a three-dimensional (3D) object is retrieved from a set of two-dimensional (2D) projections. Only those measurements qualify as projections, for which the intensity is proportional to a sum of a property of the sample along the beam direction. This is known as the projection requirement [1]. In bright field transmission electron microscopy (BF TEM) of non-crystalline thin samples, this property is the mass [1]; in energy filtered TEM (EFTEM) it is the concentration of a certain element [2,3]; while in high angle annular dark field scanning TEM (HAADF STEM), it is the atom number raised to a power of approximately 1.7 [4]. The intensity of a BF TEM image of a thin crystalline sample is determined primarily by phase contrast and is thus not considered a projection within the tomographic framework.

The shape of the image of an atom is primarily determined by the electron beam which can be much broader than the atom's potential [5]. In general, the intensity in a pixel of an atomic resolution image will therefore not be proportional to a projection along the beam direction of a physically meaningful quantity (e.g. the atomic potential) of the sample, thus violating the projection requirement. For atomic resolution tomography the projection requirement needs to be restated: The image intensity of any two atoms must equal the sum of the image intensities of the individual atoms. Hartel et al. show in Ref. [4] that this is easily accomplished for atoms separated in a direction normal to the beam by choosing a detector with a high angular width. For atoms only separated in the direction parallel to the beam also a high inner detector radius is necessary if the distance is below approximately 1 nm. Also coherent dynamical effects must be excluded, this means that the sample must be tilted out of zone axis or must be amorphous.

Four difficulties are specific for electron tomography:

- A tilt series typically extends from -70° to $+70^{\circ}$, i.e. there is a missing wedge of 40° (except when specialized sample holders are used [6]). This results in an anisotropic resolution in the reconstruction, with an elongation along the beam direction [7,8].
- The number of projections is often kept low because of the long recording time per image and/or to avoid beam damage. This can result in artifacts and a reduced resolution [9].
- In HAADF STEM the electron count is typically low, leading to noisy projections and a noisy reconstruction.
- We noticed that with iterative reconstruction techniques the resolution improves with iteration number while the noise worsens. To mitigate the noise the reconstruction is terminated at cost of resolution loss.

In Ref. [10], Batenburg et al. introduce the discrete algebraic reconstruction technique (DART) in electron microscopy. The

* Corresponding author. Tel.: +32 3 265 32 46; fax: +32 3 265 33 18.
E-mail address: wouter.vandenbroek@ua.ac.be (W. Van den Broek).

0304-3991/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.ultramic.2009.08.003

authors show that consequences of the problems stated above can be mitigated if prior knowledge is exploited that justifies the use of few discrete grey values. If a reconstruction on the atomic scale would be carried out with a regular technique using the projections themselves, then the reconstructed object would have continuously distributed grey values, and therefore DART would not be applicable. In [11] Jinschek et al. describe a method that works in zone axis orientation and thus is complementary to this work. They use an Argand plot to find the number of atoms in all columns and subsequently use these numbers as input for the DART algorithm, instead of the projections themselves. These achievements suggest that fully exploiting prior knowledge might be very beneficial for atomic resolution tomography.

In Section 2 we present a reconstruction technique for atomic resolution HAADF STEM that explicitly takes into account the spatial distribution of the image of a single atom, and thereby the microscope settings that cause this distribution: the high tension, the spherical aberration, the defocus, the objective aperture and the inner and outer angles of the detector. Also the atomic nature of the object is explicitly taken into account.

In Section 3 the Cramér–Rao lower bound (CRLB) is introduced [12–14]; this is the theoretical lower bound on the variance of unbiased estimates of the unknown parameters, here corresponding to the grey values. This will allow us to make predictions about the noise level in the reconstruction for a given microscope setting.

In Section 4 this technique is tested on a simulation of an amorphous silicon (Si) particle, with a large missing wedge, a low number of projections and a signal-to-noise ratio (SNR) of only 3. It turns out that the algorithm is very resistant to noise. It finds exactly for every voxel whether it contains an atom. The reconstruction shows none of the artifacts typical for a missing wedge or a low number of projections. Since the algorithm converges, no arbitrary stopping criterion is needed. In addition, the probability with which aluminium (Al) and Si can be distinguished is calculated.

In Section 5 the method and the simulation results are discussed and in Section 6 the conclusions are drawn.

2. Algebraic reconstruction technique

2.1. Modelling the projection process

In the algebraic reconstruction technique (ART) [9,15,16] the object f is overlaid with a grid and within each square, f is assumed to be constant. The pixels of f are indexed with one variable $j=1,2,\ldots,K$, like in Fig. 1. The values of the projections p are indexed with one variable i too. If there are N projection angles with n pixels per projection, then i runs from 1 to $M=nN$. According to [9] every projection value p_i can be written as the weighted sum of the pixels of f. The projections are then a set of linear equations:

$$\sum_{j=1}^{K} w_{ij} f_j = p_i, \quad i=1,2,\ldots,M. \tag{1}$$

There are various ways of choosing the weights, and in this article we propose a new definition. For every pixel i in the projections and every voxel j in the object, we define the weight w_{ij} as the value in i of the image of a single atom positioned in j, normalized to a total intensity of 1, i.e. $\sum_i w_{ij}=1$. This is illustrated in Fig. 1. The image of an atom is determined by its atomic number and the microscope parameters; these weights therefore bring prior knowledge in the reconstruction, i.e. make the algorithm model based. The spatial extent of an atom image is now accounted for completely by the weights; the object f therefore only contains

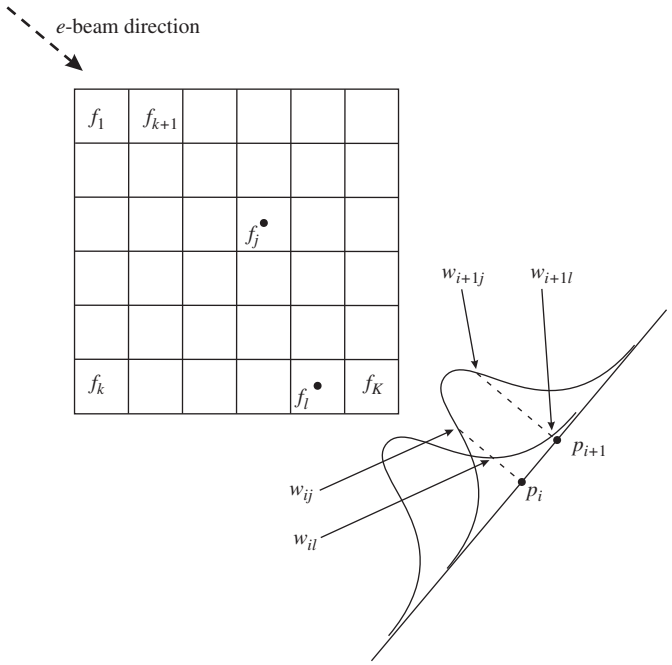

Fig. 1. The atom in position j gives rise to the right intensity profile, the atom in position l to the left profile. The resulting image p_i in i is the sum of the values in i of both profiles, i.e. $p_i=w_{ij}f_j+w_{il}f_l$.

Dirac functions at the atom positions and zeros in between. Envisage an object f like in Fig. 1. It is composed of an atom in voxel j and one in voxel l. Its image intensity is a sum of the image intensities of both individual atoms. The projection value p_i is given by Eq. (1) and reduces to $w_{ij}f_j+w_{il}f_l$ because f is composed of Dirac functions only.

Usually, reconstruction algorithms calculate 2D images out of 1D line projections, and 3D reconstructions by treating the 2D projections as a set of line projections normal to the tilt axis and the object as a stack of 2D slices normal to the tilt axis. The underlying assumption is that neighbouring slices are imaged independently. This assumption is invalid for atomic resolution due to the spatial extent of the image of a single atom in the two dimensions normal to the beam direction. Therefore an algorithm was written that does the 3D reconstruction directly out of the 2D projections.

2.2. Maximum likelihood

Eq. (1) is more compactly written as:

$$W\boldsymbol{f}=\boldsymbol{p}, \tag{2}$$

with $\boldsymbol{f}=(f_1,\ldots,f_K)^T$ the unknown object f written as a vector, $\boldsymbol{p}=(p_1,\ldots,p_M)^T$ the vector of the expectation values of the projections and W the $M\times K$ matrix containing the weights w_{ij}. All pixels in the projections are assumed to be statistically independent and to suffer from the Poisson noise. Therefore the probability $P(\boldsymbol{q};\boldsymbol{f})$ of obtaining a set of measurements q defined by the vector $\boldsymbol{q}=(q_1,\ldots,q_M)^T$ is given as

$$P(\boldsymbol{q};\boldsymbol{f})=\prod_{i=1}^{M}\frac{(W\boldsymbol{f})_i^{q_i}}{q_i!}\exp[-(W\boldsymbol{f})_i]. \tag{3}$$

This function is called the joint probability density function of the observations. It is a function of the observations q.

The maximum likelihood (ML) method for estimating the object vector $\boldsymbol{f}$ is as follows. The available observations $\boldsymbol{q}$ are substituted in the probability density function. Since the

observations are numbers, the resulting expression depends only on the elements of the object vector $\boldsymbol{f}$. The elements of $\boldsymbol{f}$, the hypothetical true values, are now considered to be variables. To express this, they are replaced by $\boldsymbol{t}=(t_1,\ldots,t_K)^T$. The logarithm of the resulting function, $\ln P(\boldsymbol{q};\boldsymbol{t})$, is called the log-likelihood function. The ML estimate $\hat{\boldsymbol{f}}_{\mathrm{ML}}$ of the object vector $\boldsymbol{f}$ is defined as the vector that maximizes the log-likelihood as a function of the object:

$$\hat{\boldsymbol{f}}_{\mathrm{ML}}=\arg\max_{\boldsymbol{t}}\ln P(\boldsymbol{q};\boldsymbol{t}), \tag{4}$$

i.e. it is the object that is most likely to have produced the projections q at hand. It is proven in [16] that $\hat{\boldsymbol{f}}_{\mathrm{ML}}$ satisfies

$$\hat{\boldsymbol{f}}_{\mathrm{ML}}=\hat{\boldsymbol{f}}_{\mathrm{ML}}W^T\frac{\boldsymbol{q}}{W\hat{\boldsymbol{f}}_{\mathrm{ML}}}, \tag{5}$$

where arithmetic operations between vectors are elementwise, and W is normalized such that each column sum is 1. Expectation maximization [16] is the simplest iterative method that solves Eq. (5):

$$\hat{\boldsymbol{f}}^{k+1}=\hat{\boldsymbol{f}}^{k}W^T\frac{\boldsymbol{q}}{W\hat{\boldsymbol{f}}^{k}},\quad k=0,1,\ldots \tag{6}$$

This requires an input $\hat{\boldsymbol{f}}^0$; here it is set to a constant vector with a value equal to the average of f, which is estimated as the sum over all of q's pixels divided by the number of projection angles and the number of voxels in f. Comparing this with the simultaneous iterative reconstruction technique in [9] shows that expectation maximization essentially is a multiplicative version of this method. The joint probability density function is also defined for distributions other than the Poisson distribution; thus, the concept of ML is not restricted to Poisson noise.

We propose a heuristic modification that explicitly takes the atomic nature of the sample into account. One expects f to be composed of individual Dirac functions separated by the vacuum, so every voxel with a value below half of the average of the non-zero voxels is likely to belong to the vacuum. Hence we set its value to zero and keep it zero in all subsequent iterations. The loss in intensity is compensated for by multiplying the values of the other voxels with a constant. This is repeated for every iteration k. It is observed that this greatly accelerates convergence.

3. Cramér–Rao lower bound

The joint probability density function can also be used to define the Fischer information matrix F and the Cramér–Rao lower bound (CRLB), a lower bound on the variance of any unbiased estimator of the object f [17]. The Fisher information matrix with respect to the elements of the object vector $\boldsymbol{f}$ is defined as [13]

$$F_{rs}=-E\left[\frac{\partial^2\ln P(\boldsymbol{q},\boldsymbol{f})}{\partial f_r\partial f_s}\right], \tag{7}$$

where E denotes the expectation value. From Eqs. (1) and (3) it follows that [17]

$$F_{rs}=\sum_{i=1}^{M}\frac{1}{(W\boldsymbol{f})_i}w_{ir}w_{is}. \tag{8}$$

A theoretical lower bound on the variance of the estimated parameters can now be given: the variance $\mathrm{var}(\hat{f}_r)$ of any unbiased estimator $\hat{f}_r$ of f_r is larger than or equal to the corresponding element on the diagonal of the inverse of F,

$$\mathrm{var}(\hat{f}_r)\geqslant[F^{-1}]_{rr}. \tag{9}$$

The right-hand side of this inequality is known as the Cramér–Rao lower bound (CRLB) [12–14]. The square root of the CRLB will be denoted σ_{CR} and constitutes a lower bound on the standard deviation of any unbiased estimator.

The ML estimator has many interesting properties, three of which are of particular importance here:

- *Consistency.* The probability that the ML estimator deviates more than a specified amount from the true value of the object can be made arbitrarily small by increasing the number of observations.
- *Asymptotic normality.* If the number of observations increases, the probability density function of the ML estimator tends to a normal distribution.
- *Asymptotic efficiency.* If the number of observations increases, the variance of an ML estimator tends to the CRLB. In this sense, the ML estimator is most precise.

4. Simulation

We will simulate a cluster of 43 silicon atoms arranged in a random fashion. The particle is spherical with an approximate diameter of 8 Å. The object to reconstruct is a cube of $40\times40\times40$ voxels, with the particle placed in the middle. The width of the pixels and voxels is 0.3 Å. A top view of the object f is depicted in Fig. 2a.

The simulations are done with a spherical aberration C_s of 0.5 mm and a high tension (HT) of 300 kV, with a defocus Δf and a condenser aperture of semi-angle α_0:

$$\Delta f=\sqrt{1.5C_s\lambda}, \tag{10}$$

$$\alpha_0=(6\lambda/C_s)^{1/4}. \tag{11}$$

These are the Scherzer conditions as defined in Ref. [18], with λ the electron wavelength. Filling in the numbers leads to $\Delta f=38.4$ nm and $\alpha_0=12.4$ mrad. The image of an individual Si atom is calculated with STEMsim [19], with a detector with an inner angle Θ_{in} of 60 mrad and an outer angle Θ_{out} of 100 mrad. Its profile is depicted in Fig. 3. In Fig. 2b a projection of f using these weights for a projection angle of 2° is shown, i.e. it shows $W\boldsymbol{f}$. All microscope settings are summarized in Table 1.

4.1. A typical reconstruction

The projection angles range from -70° to $+70^\circ$ in steps of 4°. The projections are calculated as $W\boldsymbol{f}$ and the mean signal of their non-zero pixels was set to 9; this is equivalent to putting the values of the non-zero voxels of the object to 136.2. The application of Poisson noise to the projections leads to an SNR of 3. Fig. 4 shows a noisy projection at 2°. The behaviour and convergence of this new method will be illustrated by using one typical reconstruction.

In order to save disc space, only the results of iteration numbers k equal to round$[2^m]$ for $m=0,0.5,1,1.5,\ldots$ were stored. In Fig. 5 the number of non-zero voxels in a typical reconstruction is plotted against the iteration number, the relevant scale indeed is logarithmic. The number of atoms in f is found before the 181st iteration. All voxels are now correctly classified as either containing an atom or not, but this does not mean that the algorithm has converged yet. There is still a minor difference between the values in $\boldsymbol{f}^{362}$ and $\boldsymbol{f}^{512}$, while $\boldsymbol{f}^{512}$ equals $\boldsymbol{f}^{724}$ exactly, showing that the solution converged between the 362nd and 512th iteration. Eq. (5) is now fulfilled; this means that our modified expectation maximization algorithm does find the ML estimate. In Fig. 6 the top view of the reconstruction is shown.

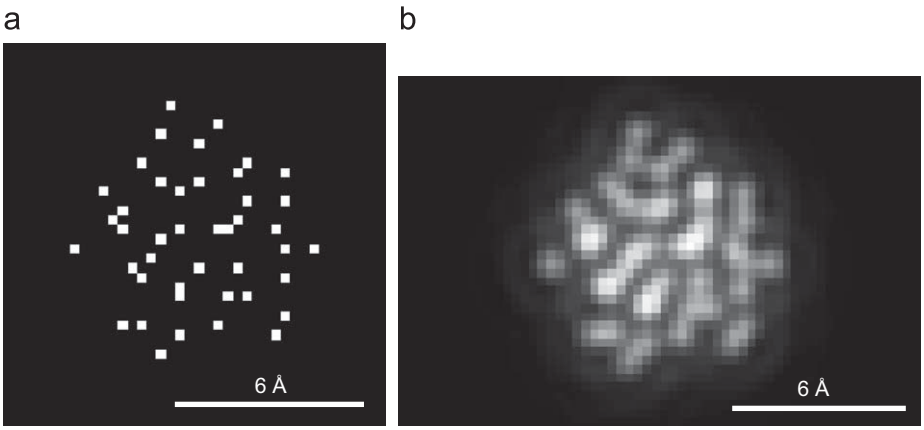

Fig. 2. The amorphous Si particle and one of its projections. (a) Top view of the object f, a Si particle, at the same angle as in Fig. 2b. (b) The noiseless projection P obtained by multiplying $\boldsymbol{f}$ with the weights W for an angle of 2°.

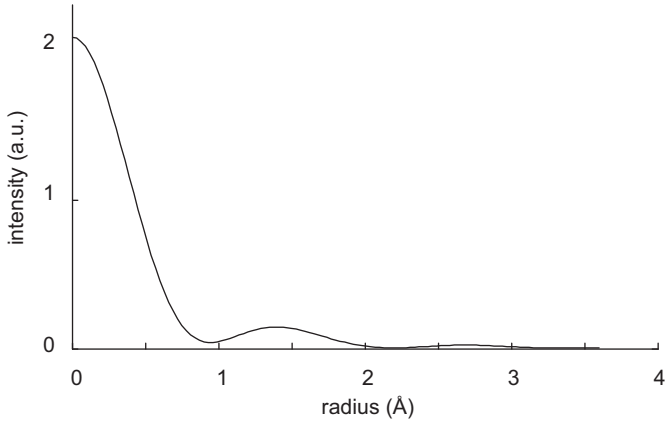

Fig. 3. Profile of the image of a single Si atom, calculated with STEMsim.

Table 1
Microscope parameters used in the simulation.

HT	λ	C_s	Δf	α_0	Θ_{in}	Θ_{out}
300 kV	1.97 pm	0.5 mm	38.4 nm	12.4 mrad	60 mrad	100 mrad

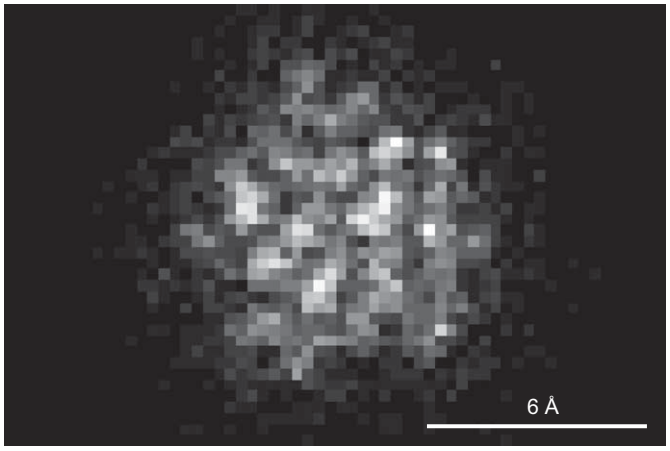

Fig. 4. Projection of the Si particle at 2° with an SNR of 3.

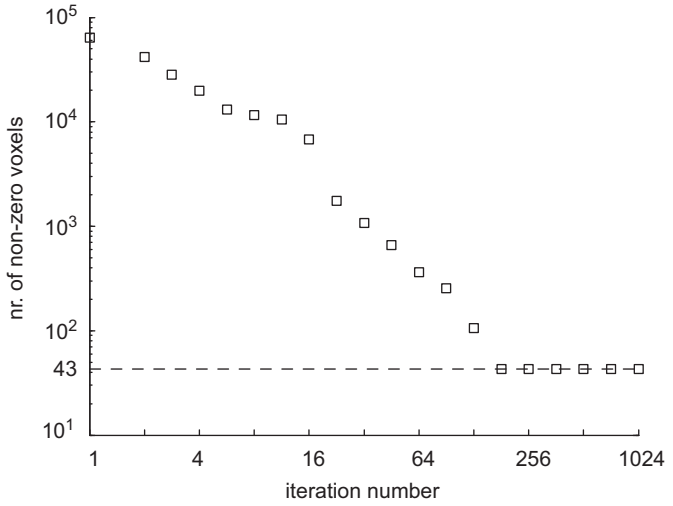

Fig. 5. Number of non-zero voxels vs. the iteration number, on a log–log scale. In iteration 181 the mask size reaches the number of non-zero voxels in f, all atoms are now found. It takes up till the 512th iteration for all the grey values to converge.

Comparison with Fig. 2a makes it clear that for every voxel we found exactly if it contains an atom or not.

4.2. Monte Carlo simulation

In Section 3 we stated that in the asymptotical case the distribution of an ML estimator is given by $N(\mu, \sigma_{CR})$, a normal distribution with a mean equal to the true value μ and a standard deviation σ_{CR}. In order to test if these asymptotic properties are fulfilled, use is made of a Monte Carlo (MC) simulation. We reconstruct the particle 100 times, each time with a different realization of the noise in the projections, with an SNR of 3. The number of iterations in every reconstruction was set to 512 and the convergence was checked by computing the mean absolute error between $\boldsymbol{f}^{512}$ and $\boldsymbol{f}^{513}$ where the mean only runs over the 43 non-zero voxels. The highest mean absolute error among the 100 simulations is 1.1×10^{-14}, so that for all practical purposes the algorithm can be considered to be converged.

Each reconstruction provides an ML estimate of the true value f_c of the voxel containing the atom closest to the particle centre; this estimate is denoted as $\hat{f}_{cML}$. The CRLB is calculated only for the voxels containing an atom, this makes F a 43×43 matrix that is directly invertible. σ_{CR} of f_c is the square root of the

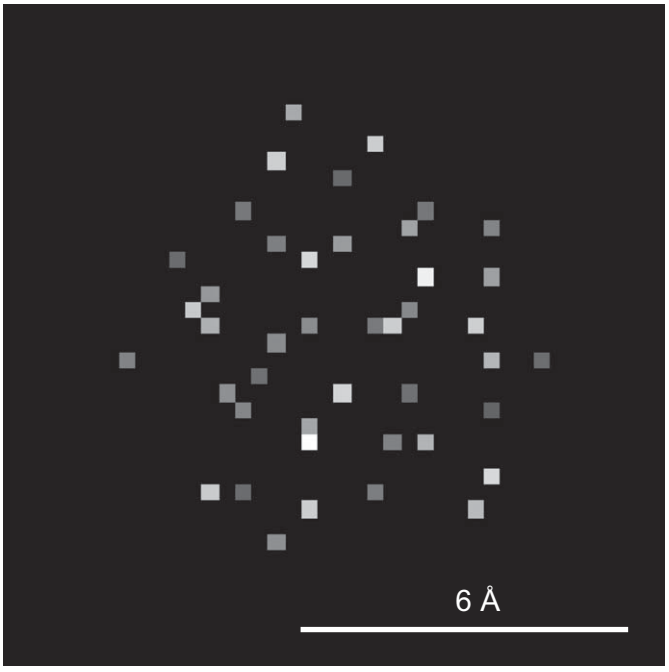

Fig. 6. Top view of the reconstruction, with adjusted contrast to show the variance in the non-zero voxels due to the noise in the projections.

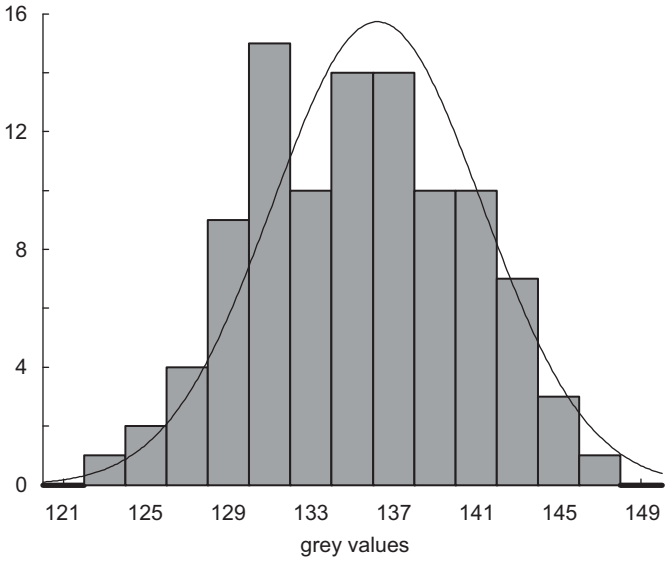

Fig. 7. The histogram of the estimates $\hat{f}_{cML}$ of the value of the voxel containing the atom closest to the particle's centre.

corresponding element on the diagonal of F^{-1}, it equals 5.1. In Fig. 7 the histogram of the estimates $\hat{f}_{cML}$ is depicted and is overlaid with $N(f_c, \sigma_{CR})$.

The Kolmogorov–Smirnov test [20] is a statistical test to determine if a set of observations is drawn from a prespecified distribution. It confirmed at the 5% significance level that the estimates $\hat{f}_{cML}$ are drawn from a normal distribution with mean f_c and standard deviation σ_{CR}, i.e. from $N(f_c, \sigma_{CR})$. Further evidence that the asymptotic case is reached is given by the 95% confidence intervals for the mean and standard deviations of $\hat{f}_{cML}$ which are $[134.3, 136.6]$ and $[4.6, 6.1]$, respectively, and thus encompass $f_c = 136.2$ and $\sigma_{CR} = 5.1$. The SNR of the signal in this voxel can thus be estimated as $f_c/\sigma_{CR} = 27$.

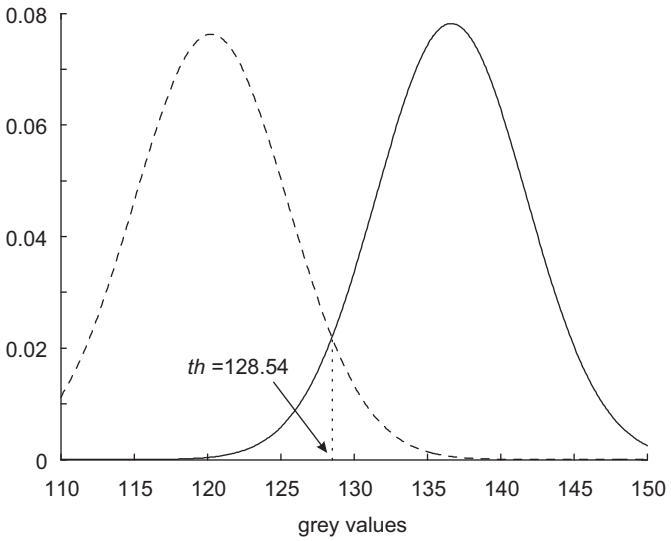

Fig. 8. The probability distributions of the value of Al (broken line) and Si (full line). Their respective means are 120.2 and 136.6, the respective standard deviations are 5.2 and 5.1. The threshold *th* is put to 128.54.

4.3. Detecting aluminium in amorphous silicon

The Si atom closest to the centre of the particle is replaced by an aluminium (Al) atom by giving it a value of 88% of a Si atom. With this percentage we assume an image intensity proportional to $Z^{1.7}$: the atomic number Z_{Si} of Si equals 14 and the atomic number Z_{Al} of Al equals 13, so that indeed $(Z_{Al}/Z_{Si})^{1.7} \simeq 0.88$. It is furthermore assumed that the images of an Al and an Si atom differ only in total intensity, and not in their intensity profiles. This is justified because at the microscope settings used, the atomic potentials are much narrower than the electron probe [5], so that the image is largely determined by the latter. We will therefore use the same weights for both atom types. The object f will be scaled such that the mean image intensity of the non-zero pixels in the projections is 9, yielding an SNR of 3. This scaling differs slightly from that of the pure Si case because of the lesser weight of the Al atom: the expectation value of a voxel with Al is 120.2, that of a voxel with Si is 136.6. The expectation value and the CRLB of the Al atom will be compared to that of its nearest neighbour, which is a Si atom.

In Section 3 it was stated that the ML estimators are asymptotically normal, consistent and efficient, and in Section 4.2 we have proofed that we indeed are in the asymptotical regime. We can therefore safely assume that the ML estimators of the grey values of the Al and the Si atom are normally distributed around their true values with a variance equal to their CRLBs. σ_{CR} of the Al and the Si atom are 5.2 and 5.1, respectively. These probability distributions are drawn in Fig. 8.

If one would choose to make the chance of being misclassified equal for both atom types, then the threshold *th* must be set to 128.54. The tail of the Al distribution to the right of *th* then contains 5.6% of the total area, and the tail of the Si distribution to the left of *th* contains 5.6% of the total area as well.

4.4. Dose fractionation

In this section we investigate how the quality of the reconstruction depends on the number of projections and on the projections' noise level. The dose fractionation theorem states that the variance of a reconstructed grey value depends directly on the

total dose in the projections [21,22]. For example, if the dose per projection is doubled while at the same time the number of projections is halved, then the variance of the reconstruction should not alter.

This is tested by doubling the dose per projection of the Si particle by doubling the values of the non-zero voxels to 272.4. The CRLB is now computed using only every second projection, i.e. the projection angles range from -70° to 66° in steps of 8°. The σ_{CR} of the atom closest to the particle centre is 10.2, which yields an SNR of 27, the same as in Section 4.2, thus validating the dose fractionation theorem.

5. Discussion

We developed a modified expectation maximization algorithm, in which we utilize the fact that in our framework any object consists of Dirac functions peaked at the atom positions and surrounded by the vacuum. For every voxel it finds whether it contains an atom or not. The MC simulation showed that at the 512th iteration the mean absolute error of the grey values in the non-zero voxels was always negligible, so no arbitrary stopping criterion is needed. Since Eq. (5) is fulfilled, the reconstruction is the ML estimate of the object.

A conventional (iterative) reconstruction technique needs to estimate the grey values of all voxels in the object, 64 000 in this example. Our new technique on the other hand uses the same data to estimate only the values of the non-zero voxels, 43 in this case. It is therefore very resistant to noise; this is reflected in the reconstruction's high SNR of 27. Take note that it does not use the number of non-zero voxels or their positions as prior knowledge; in the initial guess of this iterative method all voxels are assigned the same value. In this example we see none of the artifacts typical for a low number of projections and a large missing wedge.

Our reconstruction is an ML estimate of the object. We have shown with an MC simulation that its distribution is normal with a mean equal to the true value and a variance equal to the CRLB. The CRLB can thus be used to predict the noise in the reconstructions. In this paper the CRLB was also used to check if the two elements Si and Al are distinguishable in the reconstruction and to choose a good threshold value. Furthermore, we used the CRLB to validate the dose fractionation theorem, thus showing that the SNR of the result is determined by the total dose rather than by the number of projections or the dose per projection alone.

The CRLB is a function of the weights W that in turn are functions of the microscope settings. Since the CRLB is a measure for the precision with which parameters can be measured, the use of the CRLB to optimize the experimental design would therefore be another application [22–24].

6. Conclusion

To our knowledge, this is the first reconstruction algorithm that takes into account the spatial distribution of the single atom image, and thereby the parameters that determine this distribution: the high tension, the spherical aberration, the defocus, the objective aperture, the detector's inner and outer angle and the element's atomic number.

The algorithm performs well. For every voxel in the object it finds whether it contains an atom or not. It needs no arbitrary stopping criterion, yet is very insensitive to noise. The reconstruction shows none of the artifacts typical for a low number of projections and a large missing wedge. The reconstruction is a maximum likelihood estimator of the object, the variance of which converges to the CRLB. The CRLB therefore allows a prediction of the amount of noise for a certain experimental setup, which can, for example, be used to predict whether two elements are distinguishable or to optimize the microscope settings.

Acknowledgement

This work is carried out as part of the Condor project, a project under the supervision of the Embedded Systems Institute (ESI) and with FEI company as the industrial partner. This project is partially supported by the Dutch Ministry of Economic Affairs under the BSIK program.

S. Van Aert and D. Van Dyck (project nos. G.0220.05 and G.0188.08) gratefully acknowledge financial support from the Fund for Scientific Research—Flanders (FWO).

References

[1] P.W. Hawkes, The electron microscope as a structure projector, in: J. Frank (Ed.), Electron Tomography: Three-Dimensional Imaging with the Transmission Electron Microscope, Plenum Press, New York, London, 1992.
[2] R.F. Egerton, Electron Energy-Loss Spectroscopy in the Electron Microscope, second ed., Plenum Press, New York, 1996.
[3] P.A. Midgley, M. Weyland, Ultramicroscopy 96 (2003) 413.
[4] P. Hartel, H. Rose, C. Dinges, Ultramicroscopy 63 (1996) 93.
[5] D. Van Dyck, S. Van Aert, A.J. den Dekker, A. van den Bos, Ultramicroscopy 98 (2003) 27.
[6] N. Kawase, M. Kato, H. Nishioka, H. Jinnai, Ultramicroscopy 107 (2007) 8.
[7] M. Radermacher, W. Hoppe, Proceedings of the Seventh European Congress Electron Microscopy, Den Haag, 1980, p. 132.
[8] J.M. Carazo, The fidelity of 3D reconstructions from incomplete data and the use of restoration methods, in: J. Frank (Ed.), Electron Tomography: Three-Dimensional Imaging with the Transmission Electron Microscope, Plenum Press, New York, London, 1992.
[9] A.C. Kak, M. Slaney, Principles of Computerized Tomographic Imaging, IEEE Press, New York, 1988.
[10] K.J. Batenburg, S. Bals, J. Sijbers, C. Kübel, P.A. Midgley, J.C. Hernandez, U. Kaiser, E.R. Encina, E.A. Coronado, G. Van Tendeloo, Ultramicroscopy 109 (2009) 730.
[11] J.R. Jinschek, K.J. Batenburg, H.A. Calderon, R. Kilaas, V. Radmilovic, C. Kisielowski, Ultramicroscopy 108 (2008) 589.
[12] A. van den Bos, Parameter estimation, in: P.H. Sydenham (Ed.), Handbook of Measurement Science, Wiley, Chicester, 1982.
[13] B.R. Frieden, Physics from Fisher Information—A Unification, Cambridge University Press, Cambridge, 1998.
[14] A. van den Bos, A.J. den Dekker, Resolution reconsidered—Conventional approaches and an alternative, in: P.W. Hawkes (Ed.), Advances in imaging and Electron Physics, Academic Press, San Diego, 2001.
[15] R. Gordon, R. Bender, G.T. Herman, J. Theor. Biol. 29 (1970) 471.
[16] F. Natterer, F. Wübbeling, Mathematical Methods in Image Reconstruction, SIAM, Philadelphia, 2001.
[17] A.J. den Dekker, S. Van Aert, A. van den Bos, D. Van Dyck, Ultramicroscopy 104 (2005) 83.
[18] E.J. Kirkland, Advanced Computing in Electron Microscopy, Plenum Press, New York, 1998.
[19] A. Rosenauer, M. Schowalter, Springer Proceedings in Physics: Microscopy of Semiconducting Materials MSM Conference 2007, Cambridge.
[20] R. Barlow, Statistics—A Guide to the Use of Statistical Methods in the Physical Sciences, Wiley, Chichester, 1989.
[21] R. Hegerl, W. Hoppe, Z. Naturforsch. 31a (1976) 1717.
[22] S. Van Aert, A.J. den Dekker, D. Van Dyck, A. van den Bos, J. Struct. Biol. 138 (2002) 21.
[23] A.J. den Dekker, S. Van Aert, D. Van Dyck, A. van den Bos, P. Geuens. Ultramicroscopy 89 (2001) 275.
[24] S. Van Aert, D. Van Dyck, A.J. den Dekker, Opt. Express 14 (2006) 3830.