

Non-local Means for Scanning Transmission Electron Microscopy Images and Poisson Noise based on Adaptive Periodic Similarity Search and Patch Regularization

N. Mevenkamp¹, A. B. Yankovich², P. M. Voyles² and B. Berkels¹

¹Aachen Institute for Advanced Study in Computational Engineering Science, RWTH Aachen University, Germany

²Department of Materials Science and Engineering, University of Wisconsin–Madison, USA

Abstract

High-Angle Annular Darkfield Scanning Transmission Electron Microscopy (HAADF-STEM) allows to take images at atomic scale with a contrast proportional to the atomic number. STEM acquires an image line-by-line, pixel-by-pixel leading to characteristic distortions. Furthermore, STEM images of beam sensitive materials have to be taken with short exposure times, leading to low contrast images with Poisson noise.

In this paper, we propose an extension of Non-local Means (NLM) tailored to STEM images of crystalline structures. To find similar patches, we introduce an adaptive non-local search strategy that exploits the periodic structure of the crystal images. Furthermore, we extend the patch similarity measure to take into account the horizontal distortions typical for STEM images. Moreover, we discuss the Anscombe transform and the Poisson likelihood ratio to deal with Poisson noise. Finally, the resulting methods are compared to BM3D with Anscombe transform and PURE-LET on simulated and real data.

Categories and Subject Descriptors (according to ACM CCS): I.4.3 [Image Processing and Computer Vision]: Enhancement—Filtering

1. Introduction

Nowadays, electron microscopes are able to acquire images at atomic scale and are used in materials science to determine material properties. An important quality assessment for such images is the so-called *precision*, which essentially describes how precisely the position of atom centers can be identified in an image. The better the precision, the more material properties can be understood [HTU06, JUA*11].

Images in Scanning Transmission Electron Microscopy (STEM) are acquired by moving a focused electron probe over a sample from left to right and top to bottom along a regular grid on the sample surface. In this work, we focus on images where electrons leaving the material on the other side within a certain range of scattering angles are detected. Using this technique, called High-Angle Annular Darkfield STEM (HAADF-STEM) [BBSB*12], the number of detected electrons is proportional to the atomic number of the material at the imaged positions. The line-by-line sequential acquisition, in combination with sample movement due to environmental and instrumental disturbances, leads

to characteristic horizontal distortions of the visible atoms in the acquired images [JN13]. Besides this, low frequency sample drift introduces more global distortions of the otherwise periodic crystalline structures typically observed in experiments. Additionally, the electron counting statistics in the detector follow a Poisson distribution. Thus, the intensity measure in each pixel of the image is affected by Poisson noise. In many applications images with very poor signal-to-noise ratio (SNR) have to be used, e.g. in catalysis: both metallic [OUGB10] and oxide [BBB*14] catalysts are destroyed by the high-energy electron dose required to achieve even moderate SNR in a single image.

The correction of disturbances characteristic to STEM imaging is investigated in both the field of electron microscopy and mathematics. Jones and Nellist [JN13] proposed a method for the correction of scan noise and drift in single images. Kimoto et al. [KAY*10] used a rigid registration scheme that averages a series of STEM images of the same object to achieve a significantly improved precision compared to single shot STEM images. A non-rigid

registration scheme for STEM series has been developed by Berkels et al. [BBB*14] and the best precision achieved so far with STEM [YBD*14] is based on this method. However, the quality of the result of registration approaches depends on the quality of the individual frames. At very low SNR, a large amount of frames is required. Thus, the registration approaches require a very high overall electron dose to be applied to the material.

A major aspect in the improvement of the quality of the individual frames is Poisson noise removal. Since Buades et al. [BCM05] proposed their Non-local Means (NLM) algorithm to remove additive Gaussian white noise from images, numerous variations have been proposed. One of the most successful variations is the BM3D filter developed by Dabov et al. [DFKE07]. Mäkitalo and Foi [MF11] extended this method to the removal of Poisson noise. Key for their success was the derivation of an exact unbiased inverse of the Anscombe variance-stabilizing transformation [Ans48], which approximately transforms the noise distribution from Poisson to Gaussian. There also exist competing approaches that directly deal with the Poisson noise statistics. Deledalle et al. proposed a probabilistic ansatz to formulate patch similarity [DDT09] and developed an NLM variant for the removal of Poisson noise [DTD10]. More recently, Salmon et al. [SHDW13] introduced a patch-based denoising algorithm based on an adaptation of PCA for Poisson noise. Apart from the NLM based denoising algorithms, there has been extensive development of local filters based on shearlets and wavelets. One of the most successful adaptations of this concepts to Poisson noise is PURE-LET [LVBU10].

In this paper, we describe an enhanced version of the NLM algorithm [BCM05] designed to deal with distortions and Poisson noise typical for STEM imaging and tailored to images of crystalline lattices at atomic scale. Note that therefore the algorithm derived in this paper is limited to inputs with periodic structure. The classical NLM is the following weighted non-local average

$$\begin{aligned} D_h[u](x) &:= \frac{\sum_{y \in I} w(x,y)u(y)}{\sum_{y \in I} w(x,y)}, \\ w(x,y) &:= \text{Exp} \left[-\text{dist}(x,y)/h^2 \right], \\ \text{dist}(x,y) &:= \|u(\mathcal{N}_n(x)) - u(\mathcal{N}_n(y))\|_{2,a}^2, \end{aligned} \quad (1)$$

where $\|\cdot\|_{2,a}$ denotes the L^2 -distance with Gaussian kernel and $\mathcal{N}_n(y)$ is the patch of size $n \times n$ with center $y \in I$. In this work we fix $n = 11$. We focus on extensions of the L^2 -distance based patch similarity measure that improve the recognition of self-similarity in the presence of local horizontal distortions of the type described earlier. Additionally, we develop an adaptive similarity search that exploits periodic structure and replaces large local search windows by small non-local ones. We consider a Poisson noise model:

$$\begin{aligned} u : I &:= \{1, \dots, N_x\} \times \{1, \dots, N_y\} \rightarrow \mathbb{N}_0, \\ u(x) &= Y_x \sim \text{Pois}(\lambda(x)), \end{aligned} \quad (2)$$

where $\lambda(x) \in \mathbb{R}_{>0}$ denotes the number of expected counts at x , i.e. the intensity of an underlying ground truth image. We compare two approaches of adapting NLM to this model: Anscombe variance-stabilizing transformation [MF11] and patch similarity based on Poisson likelihood ratios [DTD10].

The layout of the paper is as follows. Section 2 describes the periodic search strategy. Section 3 introduces a patch similarity measure based on patch regularization. In Section 4, we recall the concepts of [MF11] and [DTD10] to extend NLM to Poisson noise. The performance of the proposed methods are illustrated with numerical results in Section 5 and compared to PURE-LET and BM3D with Anscombe transformation. Finally, conclusions are drawn in Section 6.

2. Adaptive periodic search strategy

Performing the full NLM algorithm on the image u requires the calculation of $(N_x \cdot N_y)^2$ weights. In the literature, it is often suggested to reduce this computational effort by restricting the search for similar patches to a local search window for each patch [BCM10]. For images with periodic structures like crystalline lattices such a local strategy disregards most of the periodic self-similarities leaving much of the denoising capabilities of NLM untapped. To close this gap, we propose to use a periodic search grid with small non-local search windows instead. The periodicity analysis necessary to construct the search grid is based on ideas from [Ber13].

We assume that the periodic structures of interest consist of an elementary cell that is placed along two axes with angles α_1, α_2 and spacings or periods $\Delta x_1, \Delta x_2$ to produce the whole pattern. We approximate the parameters as follows.

First, the angles α_1, α_2 are defined from the positions \hat{x}_1, \hat{x}_2 of the two brightest peaks relative to the origin in the modulus of the discrete Fourier transform of the image that are (numerically) non-collinear. Then, for each of these directions, we extract a corresponding 1D intensity profile from u , where the brightest peak in u used as the common origin. Fitting a sine function or a sum of sine functions to these profiles yields the corresponding spacings $\Delta x_1, \Delta x_2$. The spacings could also be estimated from the modulus, but we found the modulus estimates to be less reliable.

With these parameters, we iteratively define the following approximately periodic search grid for a reference pixel $x \in I$

$$\begin{aligned} \pi_{1/2}^0(x) &:= x, \\ \pi_{1/2}^{\pm k}(x) &:= \arg \max_{y \in \mathcal{N}_S(\pi_{1/2}^{\pm(k-1)}(x))} w(x,y) \pm \Delta x_{1/2} \begin{pmatrix} \cos \alpha_{1/2} \\ \sin \alpha_{1/2} \end{pmatrix}, \end{aligned}$$

and extend it to a set of non-local $S \times S$ search windows

$$\pi_n(x) := \left(\bigcup_{k_1, k_2} \mathcal{N}_S \left(\pi_2^{k_2} \left(\pi_1^{k_1}(x) \right) \right) \right) \cap I.$$

The argmax locally resets the pattern structure to the pixel

in the non-local search window $\mathcal{N}_S(y)$ with the highest patch similarity to the reference pixel x . This allows to account for imperfections in the periodic structure of the observed crystal due to signal distortions or material defects, as well as errors in the estimation of the parameters $\alpha_i, \Delta x_i$. However, in case the grid parameters are too far off, the adaptive adjustment with the arg max is unable to correct the misprediction. This leads to a non-matching periodic search pattern resulting in a poor reconstruction, which, in the extreme case of a very high filter parameter h , may produce periodic features not present in the input image. Nevertheless, the proposed algorithm is very robust to noise and gives good approximations of the search grid for all STEM images of crystalline structures we have tested, cf. Figure 1 for an example.




Figure 1: Detected directions of a GaN lattice (left, blue), adaptive periodic search grid (right, yellow) with 5×5 windows and locally most similar pixels (red) for an atom-centered reference pixel (green).

3. A regularized patch similarity measure

In the following, we adapt the NLM patch similarity to local distortions aligned with the scan lines in STEM images. These distortions affect the similarity of patches depicting originally identical atoms. To resolve this issue the distance function $\text{dist}(x, y)$ in (1) is replaced by a regularized version

$$\text{dist}_{\text{reg}}(x, y) := \|u(\mathcal{N}_{n,\text{reg}}(x)) - u(\mathcal{N}_{n,\text{reg}}(y))\|_{2,a}^2,$$

where $\mathcal{N}_{n,\text{reg}}(y)$ denotes an optimally (as defined below) shifted $n \times n$ patch with center $y \in I$.

The patch regularization is based on the optimization of real-valued line shifts for each patch to increase the regularity of u in vertical direction. For a patch $\mathcal{N}_n(x)$ and a vector of horizontal shifts $s \in \mathbb{R}^n$ the shifted patch is

$$\mathcal{N}_{n,s}(x) := \left\{ (i + s^j, j) \mid (i, j) \in \mathcal{N}_n(x) \right\}.$$

Note that superscript indices have to be converted to row indices of the corresponding patch. To evaluate u on shifted patches, we use piecewise linear interpolation in its first variable, representing the horizontal axis in the image, i.e.

$$\begin{aligned} \hat{u} : \mathbb{R} \times \{1, \dots, N_y\} &\rightarrow \mathbb{R}_{\geq 0}, \\ \hat{u}(t, j) &:= u(\lfloor t \rfloor, j) + (t - \lfloor t \rfloor) \cdot (u(\lfloor t \rfloor + 1, j) - u(\lfloor t \rfloor, j)). \end{aligned}$$

For simplicity, we assume an infinite support along the horizontal axis (e.g. periodic boundary conditions). On $\mathcal{N}_n(x)$, we define a discrete derivative of \hat{u} shifted by s with respect to its second variable as follows

$$\hat{\partial}_y \hat{u}(t, j; s) := \hat{u}(t + s^j, j) - \hat{u}(t + s^{j+1}, j + 1).$$

Here, $s_{n+1} := 0$. This allows us to formulate the patch regularization problem:

Problem 1 (Vertical patch regularization) Let

$$f_{ij} : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}, f_{ij}(s) := \|\hat{\partial}_y \hat{u}(\cdot, j, s)\|_{L^2([i, i+1])},$$

and define the target function $F_x : \mathbb{R}^n \rightarrow \mathbb{R}_{\geq 0}^n$ by

$$F_x(s) := (f_{ij}(s))_{i,j \in \mathcal{N}_n(x)}.$$

Then, we define the regularized patch as the optimally shifted patch, i.e. $\mathcal{N}_{n,\text{reg}}(x) := \mathcal{N}_{n,s^*(x)}(x)$ with

$$s^*(x) := \arg \min_{s \in \mathbb{R}^n, s_{j^*} = 0} \frac{1}{2} \|F_x(s)\|_2^2.$$

Since the patches $\mathcal{N}_{n,\text{reg}}(x)$ and $\mathcal{N}_{n,\text{reg}}(y)$ are used to compare the similarity of u at x and y the patch central line j^* must be fixed. Note that shifts larger than the patch size are not reasonable, but do not occur in practice due to local convergence of gradient based optimization and are thus not constrained here.

In principle, this classical non-linear regression problem can be approached with standard gradient based iterative solvers, like Levenberg-Marquardt. However, the constraint $s_{j^*} = 0$ leads to undesirable local minima preventing such solvers from converging to a desirable solution s^* : Consider a 3×3 patch within a binary image. Let the columns 2–3 of row 1 and the columns 1–2 of rows 2–3 of the patch be 1 and the rest 0. Let $j^* = 1$. Then, $\nabla \phi(0) = 0$ for $\phi(s) := \frac{1}{2} \|F(s)\|_2^2$. So, $s = 0$ is a local minimum but $s = (0, -1, -1)^T$ is the optimal solution. Fortunately, this problem can be solved by performing a basis change before linearizing F . Let $e = (e_1, \dots, e_n)$ denote the canonical basis of \mathbb{R}^n . Then $\hat{e} = (\hat{e}_1, \dots, \hat{e}_n)$ with

$$\hat{e}_j := \begin{cases} \sum_{l=1}^j e_l & j < j^*, \\ e_j & j = j^*, \\ \sum_{l=j}^n e_l & j > j^*, \end{cases}$$

is also a basis of \mathbb{R}^n . Note that $\hat{s} \in \mathbb{R}^n$ represented by \hat{e} can be transformed to the canonical basis e using:

$$s_j = \begin{cases} \sum_{l=j}^{j^*-1} \hat{s}_l & j < j^*, \\ \hat{s}_j & j = j^*, \\ \sum_{l=j^*+1}^j \hat{s}_l & j > j^*. \end{cases} \quad (3)$$

If we linearize F according to the variables \hat{s} in the basis \hat{e} instead of s in the canonical basis e , local minima as observed above vanish. Thus, we can calculate the optimal shift with Levenberg-Marquardt after the basis transform.

Evaluating the target function and its derivatives The approximation errors of a numerical quadrature of the integrals in $\frac{\partial}{\partial s_k} f_{ij}(s)$ lead to discontinuities that prevent proper convergence of gradient based solvers for Problem 1. Thus, we derive analytical expressions for $f_{ij}(s)$ and $\frac{\partial}{\partial s_k} f_{ij}(s)$.

Let $s \in \mathbb{R}^n$, $s_{j^*} = 0$ and set $t_1 = i$, $t_4 = i + 1$ and

$$t_2 := i + \min\{\lceil s^j \rceil - s^j, \lceil s^{j+1} \rceil - s^{j+1}\},$$

$$t_3 := i + \max\{\lceil s^j \rceil - s^j, \lceil s^{j+1} \rceil - s^{j+1}\}.$$

Then, the discrete derivative $\hat{\partial}_y \hat{u}$ is linear on the intervals $[t_l, t_{l+1}]$ for $l = 1, 2, 3$ with slope and intercept given by

$$a_{lj} := m_{lj} - m_{l,j+1},$$

$$b_{lj}(s) := m_{lj}s^j - m_{l,j+1}s^{j+1} + (n_{lj} - n_{l,j+1} - m_{lj}i_{lj} + m_{l,j+1}i_{l,j+1}),$$

where $i_{lj} := \lfloor \frac{t_{l+1} + t_l}{2} + s^j \rfloor$, $m_{lj} := u(i_{lj}, j) - u(i_{lj} + 1, j)$ and $n_{lj} := u(i_{lj}, j)$.

Using this, one can calculate the following expression:

$$f_{ij}(s)^2 = \sum_{l=1}^3 \begin{cases} \frac{(a_{lj}t_{l+1} + b_{lj}(s))^3 - (a_{lj}t_l + b_{lj}(s))^3}{3a_{lj}} & a_{lj} \neq 0, \\ (t_{l+1} - t_l)b_{lj}(s)^2 & a_{lj} = 0. \end{cases} \quad (4)$$

From this we can compute the partial derivatives of $f_{ij}(s)$ with respect to the variables \hat{s} in the basis \hat{e} :

$$\frac{\partial f_{ij}}{\partial \hat{s}_k}(s) = \frac{\sum_{l=1}^3 (m_{lj}\hat{\delta}_{kj} - m_{l,j+1}\hat{\delta}_{k,j+1})d_{lj}(s)}{f_{ij}(s)}, \quad (5)$$

$$d_{lj}(s) := \frac{a_{lj}}{2}(t_{l+1}^2 - t_l^2) + b_{lj}(s)(t_{l+1} - t_l),$$

where from (3) we have

$$\hat{\delta}_{kj} := \frac{\partial}{\partial \hat{s}_k} s_j = \begin{cases} 1 & \text{if } j \leq k < j^* \text{ or } j^* < k \leq j, \\ 0 & \text{else.} \end{cases}$$

Note that (5) is not well-defined for $s \in \mathbb{R}^n$ with $f_{ij}(s) = 0$. Thus, we regularize the problem by replacing $f_{ij}(s)$ in (4) and (5) by $f_{ij}^\epsilon(s) := \|\hat{\partial}_y \hat{u}(\cdot, j; s)\|_{L^2((i,i+1), \epsilon)}$, where $\|u\|_{L^2(\Omega), \epsilon} := \sqrt{\int_\Omega u(x)^2 + \epsilon^2 dx}$ denotes a regularized L^2 -norm that is differentiable at zero for $u \in L^2(\Omega)$ and $\Omega \subset \mathbb{R}$.

Regularization of the optimal shifts Numerical tests indicate that the patch regularization works well for patches that contain a significant portion of an atom. However, it does not perform well for patches that contain atomic signal only in single lines that were shifted into the patch due to horizontal distortion. To counter this effect, we regularize the optimal shifts obtained for all patches across horizontally neighboring patches. This assumes that the distortions are irregular in vertical but regular in horizontal direction, which is in line with the effects of the rastering used to acquire STEM images. Theoretically, it would be possible to incorporate the regularization into Problem 1 by locally coupling these problems via a regularization term in the shift

parameters. However, this would significantly increase the computational effort of the method. Thus, we simply post-process the obtained optimal shifts via a weighted average across horizontally neighboring patches:

$$\bar{s}(x) := \frac{\sum_{k=-\lfloor n/2 \rfloor}^{\lfloor n/2 \rfloor} w(x, k)s(x_1 + k, x_2)}{\sum_{k=-\lfloor n/2 \rfloor}^{\lfloor n/2 \rfloor} w(x, k)},$$

where

$$w(x, k) := e^{-|k|/n} e^{\bar{u}(x_1 + k, x_2)} \text{ and } \bar{u}(x) := \frac{1}{n^2} \sum_{y \in \mathcal{N}_n(x)} u(y).$$

The weights w encode the observation that shifts are most reliably detected for patches containing most atomic signal.

4. Adaptation of NLM to Poisson noise

The classical NLM algorithm (1) was designed to remove additive Gaussian white noise (AGWN), i.e.

$$u(x) = u^*(x) + X, \quad X \sim \mathcal{N}(0, \sigma^2). \quad (6)$$

Especially in the low signal-to-noise regime, this model differs significantly from the Poisson noise model (2) that applies especially to low dose electron microscopy. In the following, we recall two strategies that can be used to adapt the NLM weights to the Poisson noise model.

Anscombe variance-stabilizing transformation The major difference between the noise models with Poisson (2) and Gaussian (6) distribution is that for Poisson noise the variance is different in each pixel, whereas for Gaussian noise it is constant over the entire image. This can be approximately corrected for by applying the Anscombe transformation [Ans48] $A(k) := 2\sqrt{k + \frac{3}{8}}$ to the image intensities. A is known to transform an image with Poisson noise model (2) into an image with noise that is similar to AGWN with unitary variance. Then, any model suitable for AGWN can be used to denoise the transformed image. Afterwards, the image is transformed back by inverting the Anscombe transformation. It has been pointed out in [MF11] that both the direct algebraic inverse $A_A^{-1}(D) := (\frac{D}{2})^2 - \frac{3}{8}$ and the asymptotically unbiased inverse $A_B^{-1}(D) := (\frac{D}{2})^2 - \frac{1}{8}$ introduce a significant bias at low counts. Thus, using these as inverse transformations often results in poor quality of the denoised images when the input has low SNR.

Recently, an exact unbiased inverse of the Anscombe transformation has been proposed, which is based on mapping the denoised values $D_h[A(u)] \approx E\{A(u)|u^*\}$ to the desired ones [MF11]: $A_C^{-1}: E\{A(u)|u^*\} \mapsto E\{u|u^*\}$. Here, $E\{X|Y\}$ denotes the expected value of X under the assumption Y . It has been shown that this inverse transformation is unbiased under the assumption that $D_h[A(u)] = E\{A(u)|u^*\}$. Otherwise, a bias may still remain within the estimation error. Nevertheless, in numerical experiments the proposed exact unbiased inverse of the Anscombe transformation in

combination with the previously proposed BM3D method for AGWN removal has proven to be competitive with other methods specifically designed for Poisson noise removal.

Maximum likelihood ratio based weights Instead of stabilizing the variance of the noise, it was proposed to base the NLM weights on a stochastic likelihood ratio that directly accounts for the particular noise distribution [DDT09]. For Poisson distributed observations k_1, k_2 the likelihood ratio for the hypothesis that they share an identical mean value, versus the hypothesis that their mean values are independent, is given by [DTD10]:

$$\begin{aligned} f(k_1, k_2) &:= -\log \frac{\max_{\lambda} p(k_1|\lambda)p(k_2|\lambda)}{\max_{\lambda} p(k_1|\lambda) \max_{\lambda} p(k_2|\lambda)} \\ &= k_1 \log k_1 + k_2 \log k_2 - \frac{k_1+k_2}{2} \log \left(\frac{k_1+k_2}{2} \right), \quad (7) \\ p(k|\lambda) &:= \frac{\lambda^k e^{-\lambda}}{k!}. \end{aligned}$$

In order to adapt NLM (1) to Poisson noise, the authors proposed to replace the L^2 -distance by

$$\text{dist}_P(x, y) := \sum_{z \in \mathcal{N}_n(x)} f(u(z), u(z-x+y)).$$

Note: The implementation of (7) has to ensure, that $k \log k$ is evaluated to zero for $k = 0$ (Recall: $k \log k \rightarrow 0$ for $k \rightarrow 0$).

5. Results and Discussion

We have performed numerical experiments with all possible combinations of the previously described modifications to the NLM algorithm, as well as PURE-LET [LVBU10] (using the implementation of S. Palakkal) and a variation of BM3D using the Anscombe transformation [MF11] (using the implementation of the BM3D authors). All algorithms were tested on STEM images from both measurements and simulations, with varying SNR, i.e. with different electron dose per pixel, and two different structures: single (Gallium-Nitrogen) and double (Silicon) atom elementary cell.

A perfect noise removal algorithm should preserve the horizontal distortions caused by STEM imaging; correcting these is a different problem. The original NLM algorithm was not capable of preserving the horizontal distortions. In the bottom row of Figure 2, one can clearly see how in the left image (without patch regularization) the atoms are blurred while in the two other images (with patch regularization (middle) and regularization of shifts (right)) the characteristic scan line distortions are retained. Comparing these results to the ones shown in the middle row, we see that BM3D (left) introduces less blurring of the scan noise than the original NLM, but more than NLM with patch regularization and Anscombe transform (middle and right). Note that the borders of the NLM denoised images are masked because pixels with patches exceeding the image were not denoised. The peak signal-to-noise ratio (PSNR) values shown

in Table 1 indicate that at moderate SNR the patch regularization does improve the performance of NLM, but only if the shifts are regularized. At low SNR however, the patch regularization has less effect or even reduces the quality of the reconstruction. Most likely the optimal shifts were not detected reliably enough to cope with the extreme noise level in this setting.

The periodic search was developed with the intention to speed up the computations. Indeed, in our experiments, the periodic search based approach was up to 15 times faster than the full search NLM. Our proposed algorithm is still much slower (up to 50 times) than the BM3D implementation by the original authors, which runs within seconds on the images used here. This is because BM3D uses a sliding step that skips 8/9 of the pixels; a smaller patch size (8×8 instead of 11×11); and a smaller search region (39×39 instead of ~ 4500 pixels). Besides the speed-up, we observed that our adaptive periodic search strategy also significantly improves the PSNR compared to the full search NLM results, cf. Table 1. This is because the periodic strategy truncates weights by omitting pixels that are not on the periodic search pattern, whereas the full NLM algorithm averages even highly different patches albeit with a small weight.

Table 1 lists the PSNR values for the NLM methods (middle columns) with Poisson likelihood ratio based similarity measure (top cell value) and with Anscombe transform (bottom cell value). They indicate that for inputs with high to moderate PSNR (15.1dB - 22.8dB) the Anscombe transform performs better, while at low PSNR (9.6dB - 12.6dB) the Poisson likelihood ratio is superior. Note that the overall runtime using the Poisson likelihood ratio is up to 6 times slower than using the Anscombe transformed L^2 -distances. Our NLM method with periodic search (with and without patch regularization) consistently outperforms PURE-LET at all PSNRs. At the lowest presented PSNR, our proposed NLM with adaptive periodic search strategy and Poisson likelihood ratio based patch similarity measure outperforms BM3D. At higher PSNR however, BM3D with Anscombe achieves the highest PSNRs of all tested algorithms.

Nevertheless, Figure 3 indicates that the PSNR alone is not a sufficient measure. The figure shows an experimentally acquired STEM image of a Gallium-Nitrogen lattice (top left), denoised versions obtained from the NLM algorithm with patch regularization and regularization of the optimal shifts (bottom row), PURE-LET (top center) and BM3D (top right). While the scan lines are not blurred as much by PURE-LET and BM3D as in the example shown in Figure 2, both filters produce significant artifacts in the void between the atoms. The artifacts introduced by BM3D in this case may not have a noticeable effect on a measure like PSNR, but are very problematic for an analysis of the image from a materials science perspective. In contrast, the denoised images calculated by our method are free of these artifacts and thus better for the intended application. Note that the lattices

shown in Figures 1 and 3 are not perfectly periodic with respect to the detected axes. Instead, each second atom along the almost diagonal axis of the lattices has a slight offset. Both the accuracy of the adaptive search grid shown in Figure 1 and the quality of the denoised images (e) and (f) in Figure 3 indicate that the proposed adaptive periodic search strategy manages to cope with such slight irregularities.

6. Conclusions

We have presented two major modifications to the classical NLM algorithm that exploit characteristic features inherent to STEM images of crystals, namely periodic structure and local horizontal distortions. Adaptive periodic search based on a Fourier analysis improved both the computational cost and the denoising performance of NLM. Furthermore, we have presented results indicating that patch regularization is a promising ansatz to prevent blurring of scan noise. While it already performs very well at moderate SNR, this technique has to be improved to deal with low SNR as well. Our comparison of the two patch similarity measures based on the L^2 -distance of the Anscombe transformed intensities and the Poisson likelihood ratio has revealed that the inaccuracy of the Anscombe transform in the low count regime is high enough such that the otherwise well performing Anscombe transform is inferior to the Poisson likelihood ratio at low signal-to-noise ratio. For the lowest regarded PSNR image treated in this paper (9.6dB) the BM3D method with Anscombe transformation is outperformed in PSNR by our proposed NLM method with Poisson likelihood ratio.

Based on the performance of BM3D on STEM crystal images, as well as the promising results achieved by our proposed adaptive periodic search and patch regularization, we plan to merge our proposed modifications with features included in BM3D such as hard thresholding (which could also further improve the speed-up of the periodic search) and pre-processing, as well as sparsity analysis of the matched patches. Furthermore, due to the improvement caused by the regularization of the shifts, we plan to investigate more global formulations of the regularization problem.

Acknowledgements

The authors would like to thank R. Bergmann, P. Binev and W. Dahmen for fruitful discussions. Simulation and acquisition of electron microscopy images was supported by the US Department of Energy Office of Science (DE-FG02-08ER46547).

References

- [Ans48] ANSCOMBE F. J.: The Transformation of Poisson, Binomial and Negative-Binomial Data. *Biometrika Trust* 35 (1948), 246–254. 2, 4
- [BBB*14] BERKELS B., BINEV P., BLOM D. A., DAHMEN W., SHARPLEY R., VOGT T.: Optimized Imaging Using Non-Rigid Registration. *Ultramicroscopy* 138 (March 2014), 46–56. 1, 2
- [BBSB*12] BINEV P., BLANCO-SILVA F., BLOM D., DAHMEN W., LAMBY P., SHARPLEY R., VOGT T.: High Quality Image Formation by Nonlocal Means Applied to High-Angle Annular Darkfield Scanning Transmission Electron Microscopy (HAADF-STEM). In *Modeling Nanoscale Imaging in Electron Microscopy*, Vogt T., Binev P., Dahmen W., (Eds.). Springer, 2012, pp. 127–145. 1
- [BCM05] BUADES A., COLL B., MOREL J.-M.: A review of image denoising algorithms, with a new one. *Multiscale Modeling & Simulation* 4, 2 (2005), 490–530. 2
- [BCM10] BUADES A., COLL B., MOREL J.-M.: Image denoising methods. A new nonlocal principle. *SIAM review* 52, 1 (2010), 113–147. 2
- [Ber13] BERGMANN R.: The fast Fourier transform and fast wavelet transform for patterns on the torus. *Applied and Computational Harmonic Analysis* 35, 1 (2013), 39–51. 2
- [DDT09] DELEDALLE C. A., DENIS L., TUPIN F.: Iterative weighted maximum likelihood denoising with probabilistic patch-based weights. *IEEE Transactions on Image Processing* 18 (2009), 2661–2672. 2, 5
- [DFKE07] DABOV K., FOI A., KATKOVNIK V., EGIAZARIAN K.: Image denoising by sparse 3D transform-domain collaborative filtering. *Image Processing, IEEE Transactions on* 16, 8 (2007), 2080–2095. 2
- [DTD10] DELEDALLE C.-A., TUPIN F., DENIS L.: Poisson NL means: Unsupervised non local means for Poisson noise. In *17th Intern. Conf. on Image processing* (2010), pp. 801–804. 2, 5
- [HTU06] HOUBEN L., THUST A., URBAN K.: Atomic-precision determination of the reconstruction of a 90 degree tilt boundary in $\text{YBa}_2\text{Cu}_3\text{O}_{7-\delta}$ by aberration corrected HRTEM. *Ultramicroscopy* 106, 3 (March 2006), 200–214. 1
- [JN13] JONES L., NELLIST P. D.: Identifying and correcting scan noise and drift in the scanning transmission electron microscope. *Microscopy and Microanalysis* 19 (8 2013), 1050–1060. 1
- [JUA*11] JIA C.-L., URBAN K., ALEXE M., HESSE D., VREJOIU I.: Direct observation of continuous electric dipole rotation in flux-closure domains in ferroelectric $\text{Pb}(\text{Zr,Ti})\text{O}_3$. *Science* 331, 6023 (March 2011), 1420–1423. 1
- [KAY*10] KIMOTO K., ASAKA T., YU X., NAGAI T., MATSUI Y., ISHIZUKA K.: Local crystal structure analysis with several picometer precision using scanning transmission electron microscopy. *Ultramicroscopy* 110, 7 (June 2010), 778–782. 1
- [LVBU10] LUISIER F., VONESCH C., BLU T., UNSER M.: Fast interscale wavelet denoising of Poisson-corrupted images. *Signal Processing* 90 (2010), 415–427. 2, 5
- [MF11] MÄKITALO M., FOI A.: Optimal inversion of the Anscombe transformation in low-count Poisson image denoising. *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society* 20, 1 (2011), 99–109. 2, 4, 5
- [OUGB10] ORTALAN V., UZAN A., GATES B. C., BROWNING N. D.: Direct imaging of single metal atoms and clusters in the pores of dealuminated HY zeolite. *Nature Nanotechnology* 5 (2010), 506–510. 1
- [SHDW13] SALMON J., HARMANY Z., DELEDALLE C. A., WILLETT R.: Poisson Noise Reduction with Non-local PCA. *Journal of Mathematical Imaging and Vision* 48, 2 (2013), 279–294. 2
- [YBD*14] YANKOVICH A. B., BERKELS B., DAHMEN W., BINEV P., SANCHEZ S. I., BRADLEY S. A., LI A., SZLUFARSKA I., VOYLES P. M.: Picometre-precision analysis of scanning transmission electron microscopy images of platinum nanocatalysts. *Nature Communications* 5 (2014). 2




Figure 2: (a) Simulated Silicon lattice image (5.3 – 66.4 average counts per pixel), (b) Ground truth with Poisson noise (0 – 88 counts per pixel), Denoised with (c) PURE-LET, (d) BM3D with Anscombe transformation, (e) full search NLM with patch regularization, regularization of shifts and Anscombe transformation, (f) periodic search ($S = 5$) NLM with patch regularization, regularization of shifts and Anscombe transformation, (g) periodic search ($S = 5$) NLM with Poisson likelihood ratio (no patch regularization), (h) periodic search ($S = 5$) NLM with patch regularization and Poisson likelihood ratio (no regularization of shifts), (i) periodic search ($S = 5$) NLM with patch regularization, regularization of shifts and Poisson likelihood ratio

Noisy Silicon lattice	Non-local means						PURE-LET	BM3D
	Full search			Periodic search ($S = 5$)				
	NLM	R-NLM	RR-NLM	NLM	R-NLM	RR-NLM		
22.8848	28.6314 28.8981	28.8152 29.1689	29.2262 29.7389	28.866 29.1757	29.0804 29.4107	29.4047 <u>29.9286</u>	21.5062	31.351
22.7499	29.2409 29.4898	29.2420 29.5418	29.6975 30.0692	29.8343 29.9642	29.8840 30.0368	30.3547 <u>30.6024</u>	25.929	31.2158
15.1247	23.2631 23.3659	23.3462 23.2740	23.5858 23.5635	26.3926 26.3595	26.1856 25.9913	<u>26.4841</u> 26.3662	23.3226	27.0712
12.5788	19.1668 18.4766	19.2129 18.4854	19.3624 18.5833	24.5369 23.9472	24.3904 23.7169	<u>24.5579</u> 23.9050	22.0077	25.7354
9.6277	15.9968 15.0249	15.9895 15.0018	16.0559 15.0613	23.8700 22.5414	23.7243 22.3267	23.8404 22.5080	20.3441	23.7002

Table 1: Peak signal-to-noise ratios (PSNR) of the different methods applied to a simulated Si-lattice image with varying total electron dose. NLM is Non-Local Means without patch regularization, R-NLM with patch regularization (no regularization of shifts), and RR-NLM with patch regularization and regularization of shifts. The top number in each cell indicates usage of the Poisson likelihood ratio and the bottom number indicates usage of the Anscombe transformation.




Figure 3: (a) Original Gallium-Nitrogen lattice image (29Mx magnification, 24 – 117 counts per pixel), denoised with (b) PURE-LET (c) BM3D with Anscombe transformation, (d) full search NLM with Poisson likelihood ratio, (e) periodic search ($S = 5$) NLM with Poisson likelihood ratio, (f) periodic search ($S = 5$) NLM with Anscombe transformation; (d) - (f) with patch regularization and regularization of shifts