

Accurate Extraction of Reciprocal Space Information from Transmission Electron Microscopy Images

Edward Rosten and Susan Cox

Los Alamos National Laboratory, Los Alamos, NM, USA
{edrosten, scox}@lanl.gov

Abstract. As the study of complex systems has become dominant in physics the link between computational and physical science has become ever more important. In particular, with the rising popularity of imaging techniques in physics, the development and application of cutting edge computer vision techniques has become vital. Here we present novel image analysis methods which can be used to extract the position of features in diffraction patterns (reciprocal space) with unprecedented accuracy.

The first contribution we have developed is a method for calculating the nonlinear response of photographic film by using the noise in the image enabling the extraction of accurate intensity information. This allows high-resolution (but non-linear) film to be used in place of low-resolution (but linear) CCD cameras. The second contribution is a method for accurately localising very faint features in diffraction patterns by modelling the features and using the expectation maximization algorithm directly on the image to fit them. The accuracy of this technique has been verified by testing it on synthetic data.

These methods have been applied to transmission electron microscopy data, and have already enabled discoveries which would have been impossible using previously available techniques.

1 Introduction: TEM and the Importance of Image Analysis

Over the last twenty years the ability to image materials with electrons, which have a wavelength considerably smaller than light, has revolutionised the physical sciences. However, generating meaningful data from the images obtained often requires automated image analysis. A number of programs have been created which perform basic operations well, but when dealing with diffraction patterns (reciprocal space) and with situations where accuracy and statistical significance are critical these are not adequate. Here we present methods for extracting extremely accurate and statistically significant data from diffraction patterns.

The diffraction patterns shown here were obtained using transmission electron microscopy (TEM), an important technique in both academia and industry. In a transmission electron microscope [1] a beam of electrons accelerated through




Fig. 1. A diffraction pattern taken from $\text{La}_{0.5}\text{Ca}_{0.5}\text{MnO}_3$. The parent lattice reflections are caused by the cubic parent lattice. The superstructure reflections are caused by the additional ordering which occurs at $\sim 220\text{K}$ [2], and occur along the primary axis (labelled a^*). The cutout has been enhanced to make the weak reflections more visible, since they are present as very faint spots. The superlattice reflections centred at $*$ and $+$ are shown in the other highlighted area of the diffraction pattern. The positions are found using the method described in Section 4.

a high voltage passes through a thin ($< 200\text{ nm}$) area of a sample. The beam then passes through a series of magnets which act as lenses. By adjusting the strength of these magnets either an image of the sample (real space information) or a diffraction pattern (reciprocal space information) can be observed, and recorded using photographic film or a CCD (depending on the microscope). Thus both images and diffraction patterns can be obtained from the same area of the sample. The most advanced TEMs can perform imaging almost at the atomic scale.

A diffraction pattern is essentially a power spectral density of an affine projection of the crystal lattice (the array of repeating units, or unit cells, which comprise the crystal), with some additional effects arising from dynamical (inelastic) diffraction, which can be ignored here. The main crystal lattice (the *parent lattice*) gives rise to a regular grid of spots termed *parent lattice reflections*. In the case of a typical pseudo-cubic parent lattice (see Figure 1), the grid is described by the two wavevectors a^* and c^* . A wavevector is simply a vector defined in reciprocal space. Some materials have a repeating superstructure with a period greater than the unit cell size superimposed on the parent lattice. In reciprocal space this gives rise to reflections at wavevectors smaller than the unit cell (superlattice reflections). In many strongly correlated systems a one-dimensional electronic superstructure forms at low temperatures. Pairs of superlattice reflections appear along one axis in the diffraction pattern. The position of the superlattice reflections is given by a wave vector whose magnitude is denoted q (with units m^{-1}), though it is often given in units of a^* . The relative intensity of the superlattice and parent lattice reflections varies widely depending on the type of sample, and in some experiments the superlattice reflections can be extremely faint as demonstrated by the diffraction pattern for an $\text{La}_{0.5}\text{Ca}_{0.5}\text{MnO}_3$ thin film shown in Figure 1.

Although a great deal of effort and expertise has gone into designing packages for the analysis of certain kinds of TEM images (e.g. SEMPER [3], Digital Micrograph), analysis of diffraction patterns has typically been performed by hand, which is time consuming and inaccurate. In Section 2 we present a technique for finding the nonlinear response of photographic film, which allows analysis techniques which require a linear response to be used. In Section 4, we present a system which not only allows individual measurements of superlattice reflections to be made with high, and quantifiable, accuracy, but can measure the position of every reflection in a diffraction pattern. This allows results to be measured to a high level of statistical significance, and makes it feasible to obtain data from patterns which would previously have been abandoned as unusable. These techniques have been applied to previously unsolvable problems and the results are presented in Section 5.

2 Nonlinear Film Response Correction

CCD image sensors have a linear response with electron intensity which allows the intensity of the image to be accurately measured. However, the use of photographic film has two advantages relative to using a CCD: firstly the resolution of scanned film is high (typically $\sim 4000 \times 4000$ pixels, compared to 512×512 for common TEM CCD sensors); secondly, the physical robustness of the film allows much longer exposure times (excessive illumination of a CCD with electrons causes damage). This heavily saturating the brighter reflections, allowing the faint ones to be visible. However, in order to interpret intensity information from photographic film correctly it is necessary to process the image, so that the image intensity is proportional to electron intensity.

Manufacturers provide calibration curves for the film, but they are only accurate if development conditions are identical. This is unlikely to be the case since the strength of the developing solutions changes over time as they are used or replaced. Instead, we present a method where the film response can be deduced from the image noise, individually for each image.

Noise in the electron intensity is approximately constant over the image, and is caused by random scattering of the diffracted electrons (the diffraction patterns are typically not scanned at a resolution where film grains are visible, and the level of shot noise is negligible, since the exposure time is long). Therefore, at a given film intensity, f_i , the amount of image noise, η , is related to the electron intensity e_i :

$$\eta \approx \frac{d f_i}{d e_i}, \quad (1)$$

since $\frac{d f_i}{d e_i}$ is the film sensitivity and the electron intensity can be arbitrarily set so that the noise variance is unity. We use this relationship to find the response of the film. To do this, we first smooth the image, to remove noise and then find the difference between the smoothed image (\hat{f}_i) and the unsmoothed image (f_i) at every pixel (the difference being due to noise). The measurements are binned by image intensity, and the amount of noise is measured by taking the standard




Fig. 2. Plot of noise intensity against image intensity, taken from a typical diffraction pattern

deviation of the measurements in each bin. Linear smoothing techniques such as convolving with a Gaussian kernel are unsuitable because they cause features to spread out. Instead, we fit polynomials of order r to groups of n consecutive pixels along every scanline to smooth the image. The values of n and r depend loosely on the feature size.




Fig. 3. (a) Film response measured with r in the range $3 \dots 10$ (in steps of 1) and n in the range $3r \dots 50$ (in steps of 10). (b): Noise distributions for the film responses. (c) A diffraction pattern before and after correction with the film response curve, using $r = 5$ and $n = 15$.

Having measured η for all image intensities, we now have an approximate measure of the differential of the film response, which is shown in Figure 2. Note however the increase in noise for very small pixel values. This is an artifact caused by the the image being clipped at an intensity of zero, which is not taken in to account by the polynomial fitting (even though after fitting, the polynomial is clipped). The function mapping the film intensity to the electron intensity, $e_i = \rho(f_i)$ is given by:

$$\rho(f_i) = \int_0^{f_i} \frac{1}{\eta(\tau)} d\tau \tag{2}$$

and

$$\rho(0) = 0. \tag{3}$$




Fig. 4. Plots of scanlines through diffraction patterns of $\text{La}_{1-x}\text{Ca}_x\text{MnO}_3$. The plot shows that the calibrated film response and CCD response are very similar. The differing peak widths are caused by differences in focus.

Using this equation on the data from Figure 2, we obtain the response curve ρ , shown in Figure 3 a. The integral relationship is useful since it results in a curve which is significantly smoother than the measured noise data. It should be noted that the result is not well modelled by a classic gamma correction curve ($\rho(i) = i^\gamma$): the curve for low pixel values is well modelled in this manner, but the response at high pixel values is dominated by the film saturating. The result of applying this to an image is shown in Figure 3 c.

Figure 3 a, b also show the effect of different values of r and n . As can be seen, the computed response is insensitive to these values. Further, it shows that the noise distribution is strongly non-Gaussian. This backs up the calculation showing that the shot noise is small, since at the currents and exposure times used, the distribution of shot noise (which follows a Poisson distribution) would be approximately Gaussian.

2.1 Evaluation

In order to test the film calibration, one would ideally expose the film to a known pattern (such as a ramp), and record the film response to the known electron intensity. Unfortunately, this is not possible with the equipment involved. If a sample viewed with a CCD is compared to a sample viewed on film, then the relative intensities of the spots should be the same (when the film is corrected for its nonlinear response). This is because a CCD responds linearly to electron intensity.

Figure 4 shows linescans through a diffraction pattern taken on a CCD and film. For comparison, the background level has been removed, and the brightness of the images adjusted so that the main peaks are of the same intensity.

The darker spots are significantly stronger on the film than on the CCD. When the film is corrected, the spots are approximately the same height. This gives a good indication that the film calibration produces accurate results.

3 Accurate Measurement of Parent Lattice Reflections

In order to measure the superlattice reflections, the parent lattice has to be found first. This is done by manually identifying two adjacent parent lattice reflections. The positions of these are then refined using the mean-shift algorithm [4]. Since the diffraction pattern is regular, two spots are enough to define the entire grid of reflections. In particular the pair of spots are used to define the primary (a^*) axis, which is the direction along which the superlattice ordering occurs. However, using only two spots to define the grid is not sufficiently accurate without further refinement.

The two parent lattice positions defines the entire grid of reflections. The distortion is sufficiently low that each point lies close enough to the relevant real parent lattice reflection that mean-shift can successfully refine the position.

There currently exist a variety of techniques for finding the parameters of imaging systems, such as [5], which finds the parameters of pinhole cameras, and [6] which find the parameters of a more sophisticated model, which models nonlinearities with radial distortion. However, not only is the imaging system in a TEM is not well modelled by these, but many of the standard calibration procedures require multiple views of a 3D scene. Instead, we have a set of correspondences between points in the image (the lattice reflections) and a known shape (a 2 dimensional square grid, the scale, position and orientation of which can be chosen arbitrarily), so we use a general purpose distortion model. The model we use to cover all the distortions is akin to a nonlinear (higher order) version of the homography (this is similar to the model presented in [7]):

$$\begin{pmatrix} sx \\ sy \\ s \end{pmatrix} = \begin{pmatrix} h_{1,1} & \dots & \dots & h_{1,2n+1} \\ h_{2,1} & \dots & \dots & h_{2,2n+1} \\ 0 & \dots & 0 & h_{3,2n-1} & h_{3,2n} & 1 \end{pmatrix} \begin{pmatrix} X^n \\ Y^n \\ \vdots \\ X \\ Y \\ 1 \end{pmatrix}, \quad (4)$$

where $(x y)$ are the grid coordinates and $(X Y)$ are the image coordinates, normalised so that the range of X and Y is ± 1 . Writing $\mathbf{H} = \begin{pmatrix} h_{1,1} & \dots \\ \vdots & \vdots \\ h_{3,2n-1} & \dots & h_{3,2n} & 1 \end{pmatrix}$, we first find the parameters of \mathbf{H} using a linear solution and then refine this using reweighted least-squares to minimize the image-space error. Reweighting is required because despite the distortion model, some errors are not easily modelled. For instance, a relatively large amount of distortion can occur nearby where the film is clamped in the scanner. Apart from the components of \mathbf{H} required to get the image at the correct orientation, position and scale, the components are typically quite small, and make corrections on the order of 2–3 pixels towards the edge of the image.

4 Accurate Localisation of Superlattice Reflections

In order to find the superlattice reflections, several standard feature detectors were tested, Harris [8], a DoG (difference of Gaussians—the detector used for SIFT [9] features) based detector and FAST [10]. However, due to the noisy background and the faint, sometimes overlapping nature of the features, the performance of these detectors was poor on all but the best images. Instead, a model of the image (as opposed to a general purpose detector) is used to achieve much more accurate detection.

A diffraction pattern can be considered to be an unnormalised probability distribution, with each pixel representing the probability that a given diffracted electron will end up there. In addition to linear diffraction, the electrons also undergo several linear and nonlinear effects. The end result is that the superlattice reflections end up approximately Gaussian. There is also spreading from the main central reflection, which results in the space between the spots not being completely black. This spreading typically takes the form of a very shallow gradient away from the central reflection. On the scale of a pair of superlattice reflections, this can be considered to be flat. Therefore, we can model the patch of image around a pair of superlattice reflections as a mixture model consisting of two isotropic Gaussians (the superlattice spots) and a constant background level, where the size, position and scale of the Gaussians and the scale of the background are the degrees of freedom. We can then fit this model directly to the image patch around the superlattice reflections using the Expectation Maximization (EM) algorithm [11,12]. Since a reasonable initialization for EM is available— a^* is known and the wave vector (q/a^*) is always quite close to $0.5a^*$ —the resulting algorithm is very robust and is capable of finding the positions of very faint spots, in high noise images. This is illustrated in Figure 1.

4.1 Evaluation

In order to test the accuracy of the superlattice localisation, the system was run on simulated TEM images so that a ground truth measurement for the spot position was known. The simulations measure accuracy under the following common noise conditions:

1. Addition of uncorrelated noise to the image.
2. Addition of correlated, unmodelled intensity changes. These occur as a result of spreading of bright parent lattice reflections and are manifested as an intensity ramp aligned with the approximate direction of the superlattice reflections. They are modelled here as a linear ramp, but in practice, the profile is quite variable.

This is tested in the following manner:

1. Generate ideal image (a constant value with two Gaussian spots at known positions).
2. Add a linear ramp (from $-R$ to R in intensity) aligned with the superlattice reflections.

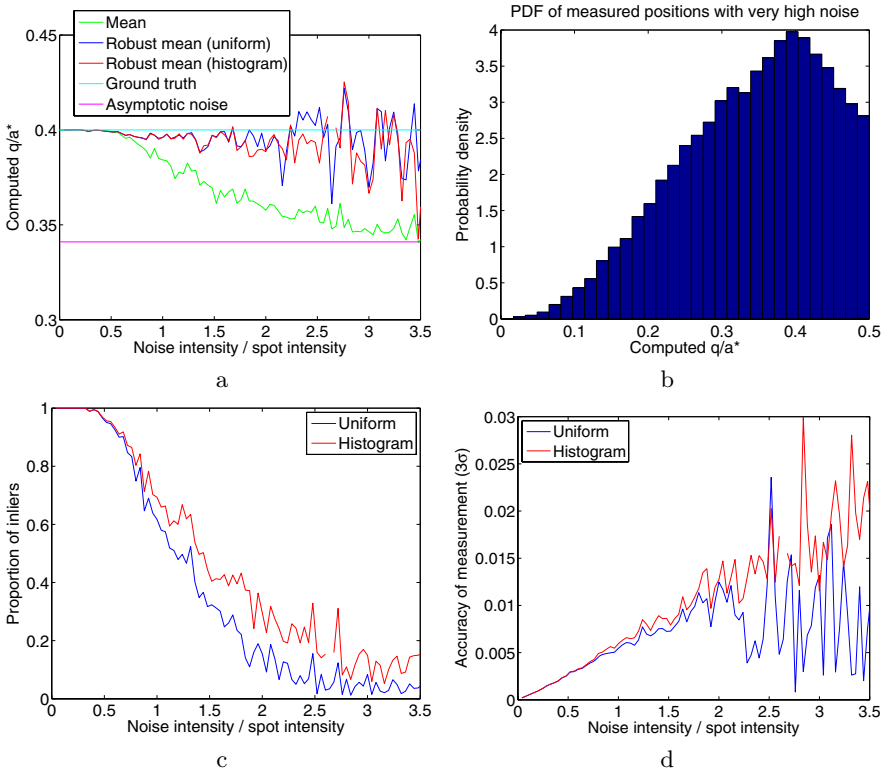


Fig. 5. (a) Graph showing how computed values for q/a^* vary with increasing noise. The background intensity is 0.1, the spot intensity is 0.1 (typical values from images) and the noise intensity (standard deviation) varies from 0 to 0.35. The patch size is 31×31 pixels. (b) Histogram of computed positions. The mean of this distribution is the asymptotic noise mean in (a). (c) Proportion of points which converge on the inlier (as opposed to background) distribution. (d) Accuracy of spot position, when 300 measurements are taken (a typical number).

3. Add uncorrelated noise, of standard deviation σ to each pixel.
4. Clip image intensities to the range 0–1.
5. Attempt to find the spot positions using EM.

When the mixture model fails to converge to a sensible value (for instance when the size of the spots becomes zero, or the computed position or the spots cannot be correct given that $q < 0.5a^*$) the result for that individual computation is rejected.

The results of the computed mean for additive noise only (i.e. $R = 0$) are shown in Figure 5 a. As can be seen, if the mean value of the computed values is taken, then the mean decreases with increasing additive noise. The reason for this is that EM will either converge on the correct spot position (with some small amount of noise), or some a random position with the ‘background’ distribution. When the noise gets large, the background distribution will dominate, so the mean computed

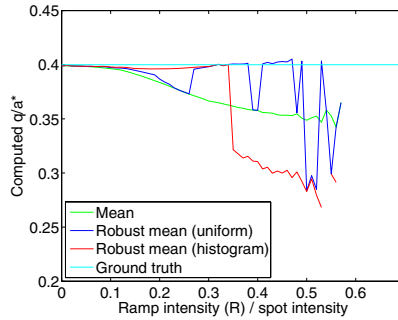


Fig. 6. Plot of the computed spot position against the additive ramp intensity. The parameters are given in Figure 5 and the noise intensity is 0.05. The histogram is taken from the high noise limit with no additive ramp. When the ramp intensity exceeds about 0.6, there are no instances in which the system converges to sensible value.

position will simply be the mean of the background distribution. At intermediate points, the computed mean will be between the two distributions.

The mean can instead be computed robustly by modelling the computed position as a mixture model of a Gaussian (representing correct points) and outliers uniformly distributed between 0 and 0.5. This mixture model can be fitted to the results (using EM), and the mean position of the Gaussian can be taken as the computed spot position. It turns out that the background distribution is not in fact uniform as one might expect, but is instead given by the histogram in Figure 5 b. The reason for this is because there is a background constant component in the image which will slightly bias the mixture model to converge towards the centre. However, as can be seen in Figure 5 a, using this distribution, instead of the uniform distribution in the computation of the mean does not produce significantly better results. However, both robust techniques produce significantly better results than the simple mean computation.

The robust computation of the mean also gives the probability that the measurement is drawn from the Gaussian (foreground) or background distribution. This can be used to compute the proportion of measurements which converge to the correct place. This is shown in Figure 5 c. This can also be used to estimate the accuracy (standard deviation) of the computed mean spot position as shown in Figure 5 d. Note that when the noise gets large, the robust mean using the histogram as the background distribution produces considerably more accurate estimates of the accuracy.

For the correlated noise, the results of the computed mean varying with R are shown in Figure 6. This shows that the system is very sensitive to unmodelled correlated noise. This justifies the decision to treat each pair of superlattice reflections separately, as opposed to computing the average pixel values over all pairs, since a few patches with strongly correlated noise could easily prevent the system from finding the correct spot position.

5 Conclusions

This paper has demonstrated that it is possible to correct the nonlinearity of photographic film, and to extract the positions of superlattice reflections to a very high degree of accuracy, even when the signal to noise ratio is low. These techniques have allowed us to extract information from very faint superlattice reflections. By analysing subtle spatial variations of the superstructure of $\text{La}_{0.5}\text{Ca}_{0.5}\text{MnO}_3$ using the methods presented here, it was possible to show that the periodicity of the superlattice can be altered by altering the strain present in the thin film [13].

References

1. D.B. Williams, B.C.: Introduction to Transmission Electron Microscopy. Plenum Publishing, New York (1996)
2. C.H. Chen, S.-W. Cheong: Commensurate to incommensurate charge ordering and its real-space images in $\text{La}_{0.5}\text{Ca}_{0.5}\text{MnO}_3$. *Phys. Rev. Lett.* **76** (1996) 4042–4045
3. W.O. Saxton: Semper: Distortion compensation, selective averaging, 3-d reconstruction, and transfer function correction in a highly programmable system. *J. Struct. Biol.* **116** (1996) 230–236
4. Comaniciu, D., Meer, P.: Mean shift analysis and applications. In: 7th IEEE International Conference on Computer Vision, Kerkyra, Corfu, Greece, Springer (1999) 1197
5. Faugeras, O.D., Luong, Q.T., Maybank, S.J.: Camera self-calibration: Theory and experiments. In: 2nd European Conference on Computer Vision, Springer (1992) 321–334
6. Stein, G.: Lens distortion calibration using point correspondences. In: 11th IEEE Conference on Computer Vision and Pattern Recognition, Springer (1997) 602–608
7. Claus, D., Fitzgibbon, A.W.: A rational function lens distortion model for general cameras. In: 18th IEEE Conference on Computer Vision and Pattern Recognition, Springer (2005) 213–219
8. Harris, C., Stephens, M.: A combined corner and edge detector. In: Alvey Vision Conference. (1988) 147–151
9. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* **60** (2004) 91–110
10. Rosten, E., Drummond, T.: Machine learning for high speed corner detection. In: 9th European Conference on Computer Vision, Springer (2006)
11. Dempster, A., Laird, N., D.B., R.: Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society B* **39** (1977) 1–38
12. Redner, R.A., Walker, H.F.: Mixture densities, maximum likelihood, and the EM algorithm. *SIAM Review* **26** (1984) 195
13. Cox, S., Rosten, E., Loudon, J.C., Chapman, J.C., Kos, S., Calderon, M.J., Kang, D.J., Littlewood, P.B., Midgley, P.A., Mathur, N.D.: Control of $\text{La}_{0.5}\text{Ca}_{0.5}\text{MnO}_3$ superstructure through epitaxial strain release. *Physics Review B.* (2006)