ELSEVIER

Ultramicroscopy 62 (1996) 43–52

ultramicroscopy

Wiener-filter enhancement of noisy HREM images

L.D. Marks *

Department of Materials Science and Engineering, Northwestern University, Evanston, IL 60208, USA

Received 27 April 1995; accepted 25 July 1995

Abstract

The use of Wiener filters to restore noisy HREM images is discussed, exploiting the fact that the noise spectrum can be rather simply estimated. Conventional Wiener, Cannon, parametric versions of these and a random-phase form are found to be very effective. Quantitative analysis indicates that the signal-to-noise ratio is improved by a factor of 3–7, with better results for larger pictures. Such filtering techniques should have rather wide applicability in electron microscopy, and could be applied on-line with TV systems.

1. Introduction

In many cases in electron microscopy the experimental images are degraded by noise either from the counting-statistics of the imaging and/or readout noise if an electronic recording mode is used. Even with a strong signal it is useful to reduce or remove the noise; with a weak signal this is essential. Handling noise is a rather old topic in the image processing literature where it has been discussed extensively, e.g. Refs. [1,2], and also but to a much lesser extent in electron microscopy. With an image which is somewhat periodic in character, probably the most common method is to window certain regions in the Fourier plane. At the heart of this approach is the fact that the strong spacings in a somewhat periodic image can be readily detected in a power spectrum. Although this technique can be powerful, it has one major drawback; it is strongly biased by the choice of windows in the Fourier plane. (I will quote here the standard result that applying windows to a pure noise object produces lattice fringes.) Thus Fourier filtering is highly subjective, and must be treated with extreme caution.

Abandoning this type of Fourier filter, is there another approach that can be used to reduce or eliminate noise, exploiting the typical spectral distribution of an HREM image? If one assumes the existence of certain features in the object one can use maximum-entropy methods, but such an assumption may not be valid. Given that the strong spacings are readily apparent in the Fourier plane for a somewhat periodic sample, one also has available a rather good estimate of the noise spectrum in the regions away from these spacings. If this noise is taken as known (or estimable), it is then possible to exploit a number of techniques which I will collectively refer to as Wiener filters (see below for further details). Having made no assumptions about the signal, the results are more objective.

In this note Wiener-filter methods are tested for use in HREM using two different high noise objects under rather extreme conditions of noise. It is demonstrated that very good results can be obtained

* Corresponding author.

SSDI 0304-3991(95)00085-2

either with conventional Wiener, Cannon or parametric Wiener filters, the most robust results coming from a random-phase version of the parametric Wiener filter.

2. Background

It is appropriate to provide a very brief background to the various filters that will be used herein (more information with more rigor can be found in Refs. [1,2]). None of them are necessarily any better than the others from a viewpoint of mathematical rigor, instead of importance is how well they behave with images that are representative of experimental HREM results. The various filters (F, implicitly in the Fourier plane) can be defined as acting on the image I to yield an estimate S of the true signal S_t contaminated by a certain noise η where:

$$FI = S \tag{1}$$

and

$$I = S_t + \eta. \tag{2}$$

2.1. Conventional Wiener filter

The conventional Wiener filter [3] is generated by looking for the filter which minimizes:

$$W = \Sigma(FI - S_t)^2 \tag{3}$$

and it is assumed that the signal and noise are uncorrelated. The appropriate solution is:

$$F = |S_e|^2 / |I|^2, \tag{4}$$

where S_e is an estimate of the signal. A conventional choice of this is:

$$|S_e|^2 = \begin{cases} |I|^2 - |\eta|^2, & |I| > \eta, \\ 0, & |I| < \eta. \end{cases} \tag{5}$$

We note that since $|F| < 1$, there is no justification for reusing the filtered data in an iterative fashion back into Eq. (4).

2.2. Conventional Cannon filter

Although the conventional Wiener filter is often called the "optimum" filter, this is only true if Eq. (3) corresponds to the correct function to minimize, i.e. the errors in amplitude have a normal (Gaussian) distribution. It is also reasonable to consider that the errors in intensity are normally distributed, in which case one minimizes the error in the power spectrum, i.e.

$$W = \Sigma\left(|FI|^2 - |S_t|^2\right)^2. \tag{6}$$

The resulting (Cannon [4,5] or homomorphic) filter is the square root of the Wiener filter, and is similarly less than 1. For completeness, it should be mentioned that the Cannon filter is rather similar to the "background subtraction" method employed by Sattler and O'Keefe [6,7].

2.3. Parametric Wiener filters

An alternative methodology is to avoid specifying in such detail (i.e. Eqs. (3) or (6)) the difference between the image and signal. One method of deriving this approach (see Ref. [2] for more details and rigor) is to look for a minimum (differentiating with respect to S and λ) of the constrained function:

$$W = |QS|^2 + \lambda\Sigma|I - S - \eta|^2, \tag{7}$$

where Q is some linear operator acting upon the signal, and λ is a Lagrangian multiplier [2,8]. If we take Eq. (7) as equivalent to maximizing the signal-to-noise of the result ($Q = |\eta/S_e|$) and no correlation between the signal and noise, one obtains the parametric Wiener filter [2,8] of form:

$$F = |S_e|^2 / \left\{|S_e|^2 + \lambda|\eta|^2\right\}. \tag{8}$$

It should be noted that the Lagrangian multiplier forces the variance of $(I - S)$ and η to be the same, i.e. $\Sigma|I - S|^2 = \Sigma|\eta|^2$, and for the special case of $\lambda = 1$ this reduces to a Wiener filter. Unlike the Wiener or Cannon filters it is possible to iterate on this form using the signal from one cycle as part of a further cycle; however, this does not appear to lead to any improvement when only noise is considered. It is possible to take the root of this filter as a "Parametric Cannon Filter".

Since this filter will play an important role later in this paper, a little further explanation is appropriate. As mentioned earlier, the conventional Wiener filter is "optimum" (in a least-squares sense) with the assumption that the signal and noise are uncorrelated. The parametric Wiener filter solves the same problem in essence as the conventional Wiener filter, but enforces a constraint upon the noise.

It is also appropriate to note the similarity of Eq. (8) and many non-linear minimization schemes such as Quasi-Newton and the Levenberg–Marquardt methods; the parameter λ is adjusted to make the denominator, a pseudo-inverse of the second derivative matrix, sufficiently positive. To continue the analogy, the conventional Wiener method would be a direct Newton method. Although mathematically the Newton method appears to be the best, it is known that Quasi-Newton methods are generally better in practice.

It should also be noted here that Q need not be taken as the noise-to-signal ratio; for instance, it can be taken as unity to give the minimum-norm solution. Andrews and Hunt [2] make the interesting comment, which will have some relevance later, that it could also be matched to the human physiology.

2.4. Random-phase parametric Wiener filter

A rather simple extension of the parametric filter above turns out to be more effective in numerical

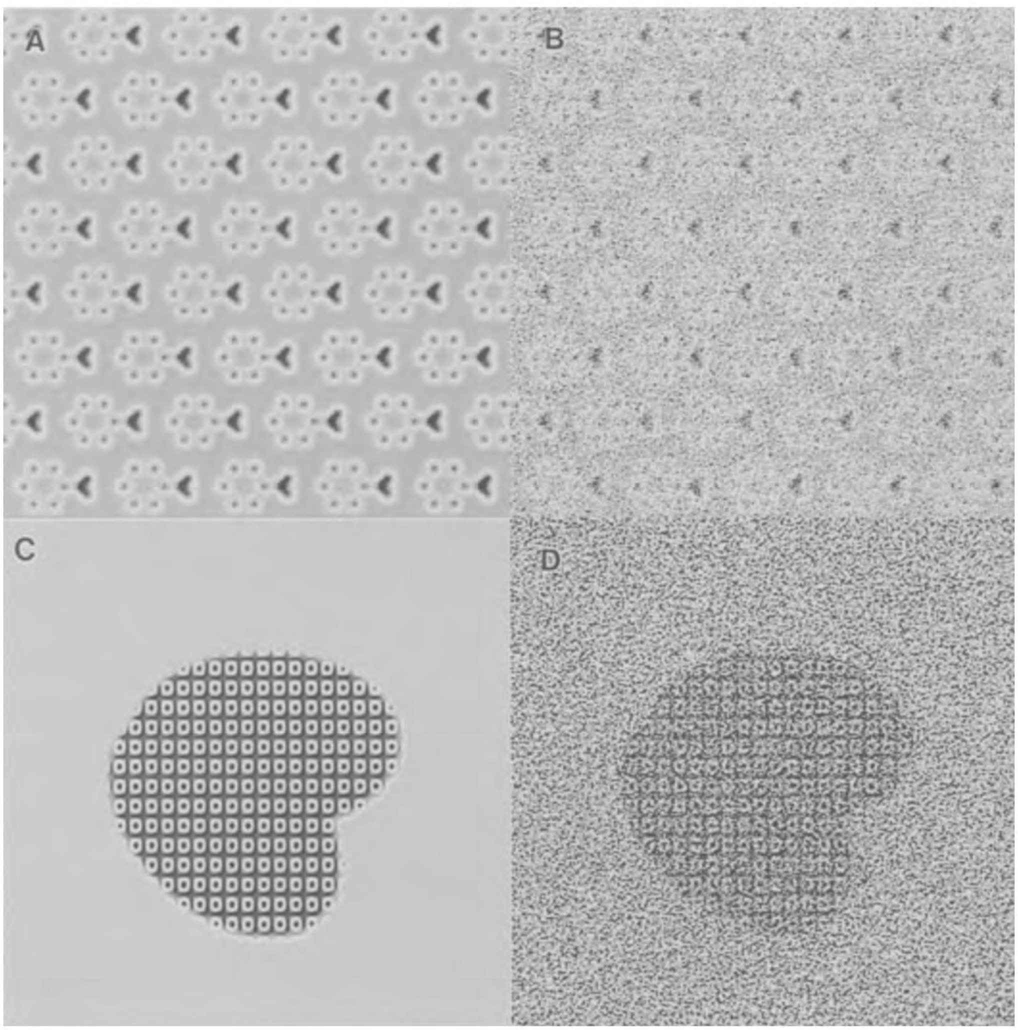

Fig. 1. The two test objects used to test the effects of the filters in (a) and (c) with corresponding images including noise in (b) and (d).

tests. This extension involves making the assumption that the phase of the noise is random, in which case one minimizes:

$$W = |QS|^2 + \lambda_1 \Sigma | I - S - \eta|^2 + \lambda_2 \Sigma(I - S). \quad (9)$$

The second Lagrangian multiplier is approximately equivalent to removing a background component. Its effect is to stabilize the filter (against over- or under-filtering), and adds a small (5%–10%) improvement in performance (see later), the stabilization effect probably being more important.

3. Numerical methodology

Rather than using standard test objects from the image processing literature, it was decided to employ two test objects that more closely resembled typical HREM problems. These are shown in Fig. 1, both the true image and that with a rather high level of noise. All the analysis was performed using Semper 6 software, with two small changes to the code: (a) a standard BSD random number generator replacing the simpler one that came with the code, and (b)

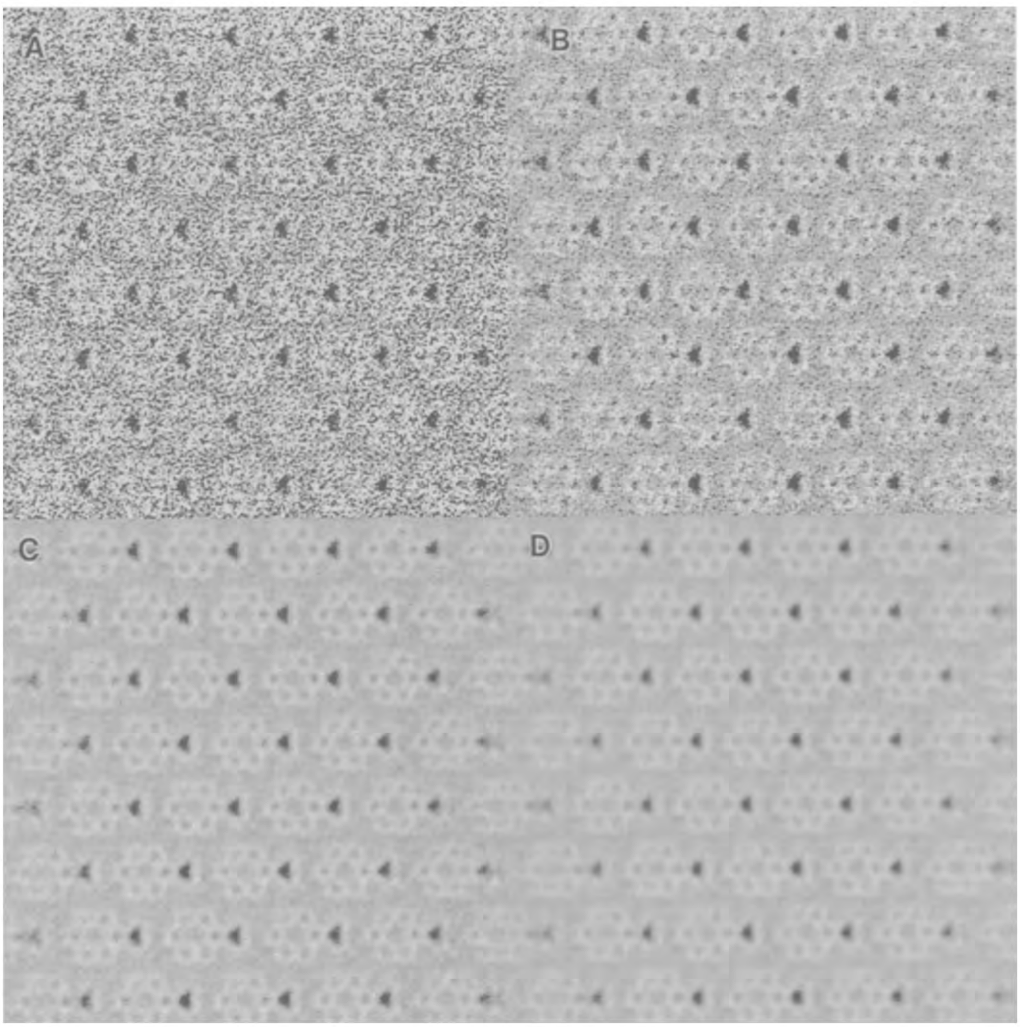

Fig. 2. Comparison of results for different constant multipliers of the known noise level amplitude ($|\eta|$ rather than η^2) using a conventional Wiener filter and test object 1: in (a–d) with multipliers of 1, 1.5, 2 and 2.5, respectively.

obtaining the mean and standard deviation of images was performed using double precision arithmetic.

Before presenting the results, a small clarification is appropriate concerning how the Lagrangian multiplier constraint was satisfied in the parametric Wiener filters. This was done by satisfying either of the equations:

$$\int \eta^2 \, d\theta \, dr = \int |I - S|^2 \, d\theta \, dr \tag{10}$$

or

$$\int r\eta^2 \, d\theta \, dr = \int r|I - S|^2 \, d\theta \, dr. \tag{11}$$

In practice, with a good noise estimate at least at the visual level there was no difference between the two. The parameter λ was determined by a Newton–Ralphson search using numerical derivatives, and converged in general exceedingly quickly (about ten iterations with an initial estimate of unity). For the random phase filter the second Lagrangian multiplier

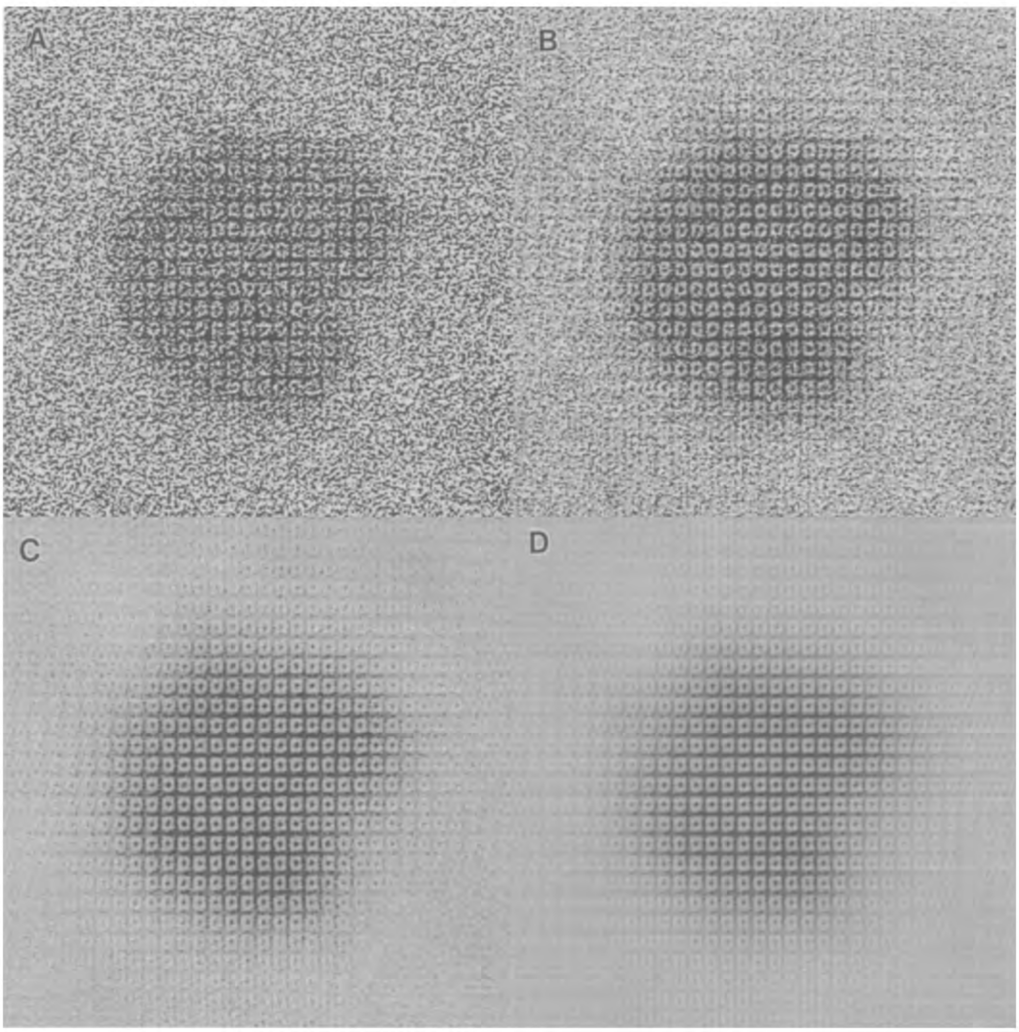

Fig. 3. Comparison of results for different constant multipliers of the known noise level using a conventional Cannon filter and test object 2: in (a–d) with multipliers of 1, 1.5, 2 and 2.5, respectively.

was solved analytically for any given value of the first one.

4. Numerical results

Only the key features of the results will be presented here, for reasons of space. First, both the conventional Wiener and Cannon filters performed rather well (see Figs. 2 and 3 for specific examples of the Wiener filter) although better results were obtained if the noise was overestimated. (Unfortunately the optimum noise overestimation level varied with different test samples.) Large overestimates of the noise tended to produce ringing around the particle-like test object, although the contrast level was somewhat better. (Depending upon the final use of the image, there is some merit to overfiltering the data.)

The parametric Wiener filter consistently performed well, see Fig. 4b, the optimum performance being for the correct value of the noise. Based upon numerical calculations of the signal-to-noise enhancement which are shown in more detail below, it

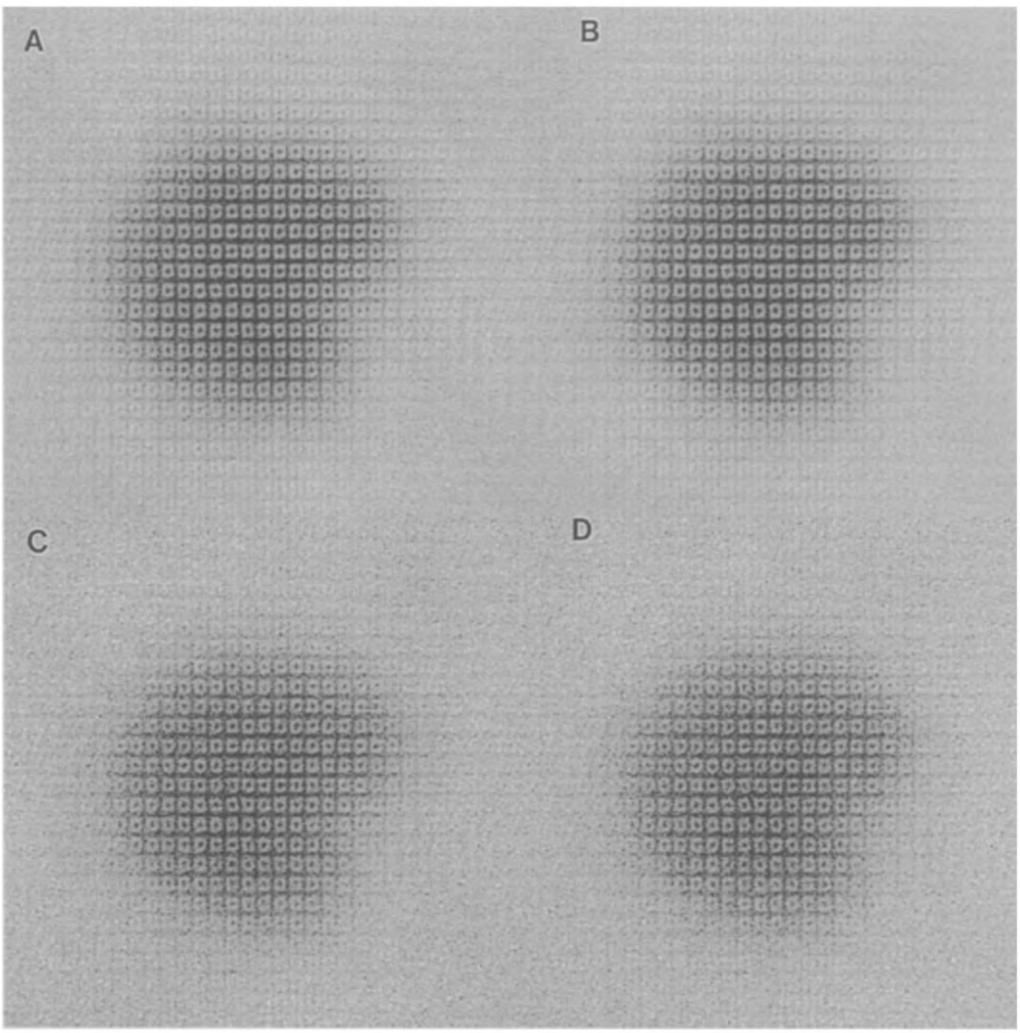

Fig. 4. Results with the known noise level for test object 2: (a) the parametric Wiener filter with the random-phase constraint, (b) the parametric Wiener filter, (c) the parametric Cannon filter with the random-phase constraint and (d) the parametric Cannon filter.

was slightly better than the parametric Cannon filter (Fig. 4d), and almost the same as the random-phase filters (Fig. 4a and Fig. 4c). (It is rather difficult visually to see substantial differences, and the eye does not necessarily agree with the computer as to which filter is best.) An important point to note is that for the parametric filters one can use the known noise spectrum rather than attempting by trial and error to determine an appropriate overfilter level for a conventional Wiener filter. This is an important advantage with an experimental object where the correct result is by definition unknown.

To provide a more quantitative measure, Fig. 5 compares the five different filters in terms of the enhancement of the signal-to-noise (S/N) ratio versus the ratio of the true noise level to that assumed in the filter operation (i.e. multiplying the noise level by a constant). As mentioned above, the conventional Wiener and Cannon filter perform better when the noise level is overestimated, while the parametric filters work with the correct noise estimation. Fig. 6 shows the enhancement of the S/N for the second test object as a function of the input S/N and the picture size. There is a very strong size dependence, with images smaller than 256 × 256 not showing any

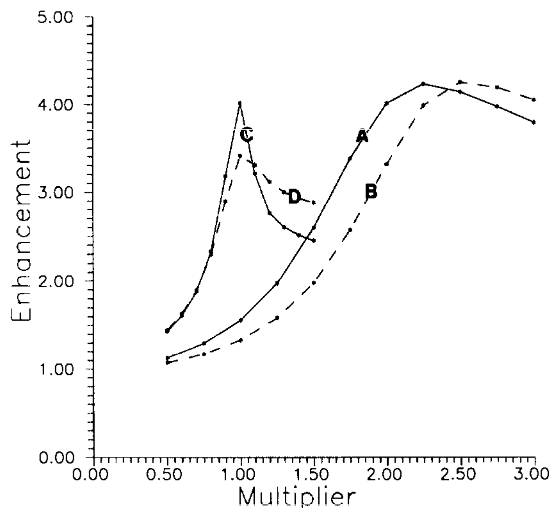

Fig. 5. Quantitative comparisons of the S/N enhancement (y-axis) versus a scaler multiplication of the noise level assumed for the filtering (x-axis) for the different filters: (a) Conventional Wiener, (b) Cannon, (c) parametric Wiener and (d) parametric Cannon. The random phase versions of (c) and (d) lie very slightly above the data shown, indistinguishable on this graph.

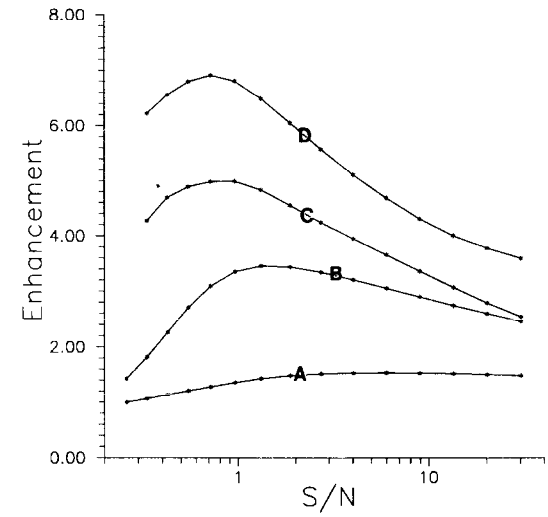

Fig. 6. S/N enhancement (y-axis) versus different initial S/N ratios (x-axis) for different initial picture sizes using the parametric Wiener filter: (a) 128×128, (b) 256×256, (c) 512×512 and (d) 1024×1024.

substantial improvement. This is understandable since with larger sizes in the Fourier plane the signal level increases linearly with the number of pixels, whilst the noise (per pixel) scales as the square-root of the number of pixels. A point of relevance for use of the filter for experimental images is that the S/N of the optimally filtered images with respect to the (noise containing) images was essentially the same as the true S/N. This fact can be used to determine the S/N of experimental data.

An example applying the technique to real experimental data is shown in Fig. 7. This image highlights the fundamental difference between Fourier window filtering and Wiener filtering, and a little further explanation is appropriate. The image is for 0.4 monolayers of gold on Si (111) taken in the off-zone imaging mode [9]. The unit cell of the surface phase is 5 × 2 [10,11], which is best described in terms of a rectangular 10 × 2 cell (in bulk notation $A = (\bar{4}22)/30$, $B = (02\bar{2})/4$), but only diffracted beams (hk) where $h = 2n + m$, $k = 2m$ appear clearly in the power spectra due to a translational symmetry element of 0.5 ± 0.25. A Fourier window filter using only the clearly resolved spots is equivalent to enforcing a centered rectangular 10 × 1 unit cell. Alternatively, if the windows form a 10 × 2 cell then windowing the noise will lead to possibly spurious features. The Wiener filtering approach overcomes these problems since no assumptions are made about the signal, only about the noise. We have extensively used these filters both for high noise levels and low ones with many experimental images (e.g. Refs. [12,13]), and the results in Fig. 7 are quite typical.

5. Discussion

The results obtained appear to be very encouraging, both qualitatively and quantitatively, and unlike other types of Fourier filtering, quite objective. (For reference, a Semper version of these filters is available at http://risc1.numis.nwu.edu/ftp/pub/filters.) It is appropriate here to mention a few additional points, extensions, some possibilities that were explored and found not to be useful and some issues on using such filters in practice.

In terms of extensions, all the methods can be applied for inversion of an image wave given some form of contrast transfer function; in fact the simple Wiener case is well known [14] and its application for linear HREM images has been discussed recently

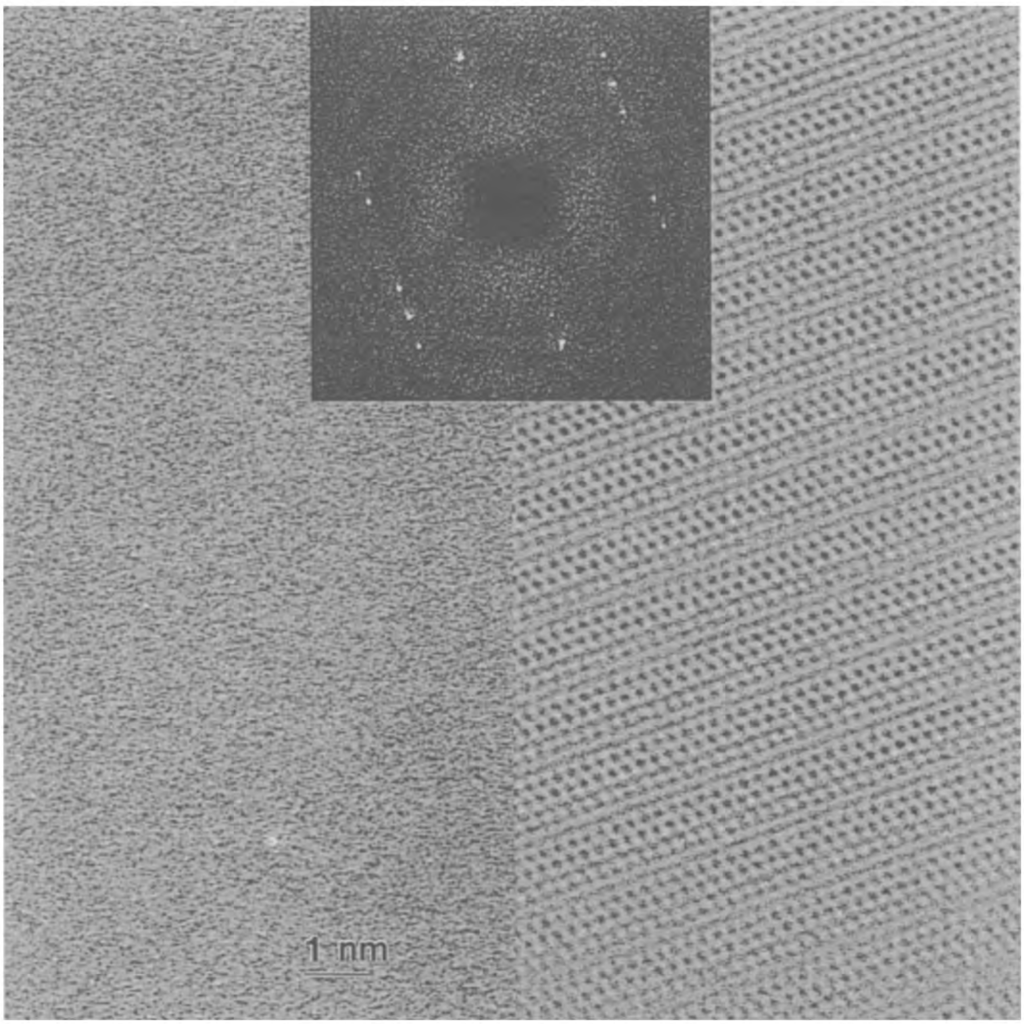

Fig. 7. The central region (256 × 512) or a 1024 × 1024 experimental HREM image of the 5 × 2 Au on Si (111) surface, on the left the original and on the right after filtering, with a power spectrum of the original at the top (high-pass filtered to remove low frequencies). Only noise is clear in the original image, but the power-spectrum shows strong spacings well separated from the noise in the Fourier plane.

[15]. The parametric Wiener methods perform at least as well, since they make the least assumptions about the character of the signal; this will be discussed in more detail elsewhere [16].

One extension that can be used in some cases is to bias the parametric filter to favor certain frequencies, i.e. introduce a Fourier windowing component. For instance, one can minimize:

$$W = |QS|^2 + \lambda_1 \Sigma B |I - S - \eta|^2 + \lambda_2 \Sigma (I - S), \tag{12}$$

where B is some bias function which, for instance, is large in regions where there is signal of interest and small elsewhere.

It may be useful to employ different forms of the operator Q in Eq. (7), for instance unity to give a minimum norm solution [2] or r which is equivalent to minimizing the "energy" of the image [2]. We noticed in many cases that the visually "best" filter was not always the filter with the optimum S/N enhancement in a strict numerical sense. The human visual system tends to act as a low-pass filter system (for further information see texts on psychophysics), and matching Q or the filter level to the human eye as suggested by Andrews and Hunt [2] has advantages.

For real experimental images, a number of techniques were explored to determine the noise, in all cases a radial average of the noise. (One can readily use a noise picture which has anisotropy in the filters.) In principle the best technique for this might be a Wiener filter of the power spectrum; unfortunately in tests this did not always work since the absolute values are rather critical. In practice we are currently performing a small Gaussian convolution (for the noise estimation) prior to taking a radial average that misses the strong signal peaks, and either using this data straight, masking out the low frequency end or one of a linear, exponential, gaussian or manual fits to the noise. It can help to remove edge effect by, for instance, Walsh windows or removal of slowly varying components of the image. From many different experimental cases that have been examined to date the parametric-filters particularly the random-phase version appears to be the best. Their advantage is that they are optimized for the experimental noise level, as against some unknown overfilter level. A disadvantage is their sensitivity to over- or underfiltering. The random-phase version appears to be more robust, i.e. to work a little better in practice than the straight filters. The best method that we are aware of to test the validity of filtered image is a Monte-Carlo simulation involving re-introduction of similar noise and then repeating the filtering; alternatively a model calculated image can be used and processed in this way for comparison with the experimental results.

It should also be noted that there is a tendency for these filters to produce artifacts for small spatial frequencies, i.e. large spacings in the image. In part this is because this region of Fourier space is poorly sampled. As a consequence it is advisable (as in many filtering operations) to remove slowly varying components and ensure that there are not large discontinuities at the edges (because of periodic continuation effects).

A few comments are appropriate about the effect of these filters on spreading information in the image plane. The location of features in an image correlates to extended information in the Fourier plane, some of which may be rather weak. There will always be some loss of this weaker signal, primarily that which lies near the noise level in the Fourier plane, and therefore some corruption of the image. In the sense that the filters remove the weak long-range tails (in reciprocal space) of strong peaks they will tend to delocalize the information. (It is not possible to make a stronger or more quantitative statement than this since delocalization depends upon the signal-to-noise level and the shape transform of the object.) If the signal is very weak or too high a noise level is estimated then the final results will be effectively lattice averaged. It should be mentioned that they will not in many cases remove sharp, strong features such as scratches.

A final point that deserves mention is that the simpler Wiener filters, and even the parametric one are rather simple to implement and no more complicated than some of the existing real-time power spectrum display routines for HREM images. It should not be at all difficult to implement these at the microscope at rates of maybe a few Hz, although it may be a few years before the technology is available for real-time Wiener filtering at TV rates.

6. Conclusions

Wiener filters, particularly the parametric Wiener filters, appear to be very powerful for HREM images since it is rather simple to obtain a good estimate of the noise spectrum. They avoid the subjectivity of Fourier window filters and appear to be both qualitatively and quantitatively rather reliable with an enhancement of the signal-to-noise ratio by a factor of 3–7.

Acknowledgements

I would like to acknowledge the support of the National Science Foundation on grant number DMR-9214505 in supporting this research. I would also like to thank my students for testing these filters on their experimental images.

References

[1] R.H.T. Bates and M.J. McDonnell, Image Restoration and Reconstruction (Clarendon Press, Oxford, 1989).
[2] H.C. Andrews and B.R. Hunt, Digital Image Restoration (Prentice-Hall, Englewood Cliffs, NJ, 1977).
[3] N. Wiener, Extrapolation, Interpolation, and Smoothing of Stationary Time Series (MIT Press, Cambridge, MA, 1942).
[4] T.M. Cannon, Digital Image Deblurring by Nonlinear Homomorphic Filtering, PhD Thesis, Computer Science Department, University of Utah, Salt Lake City, UT, 1977.
[5] T.R. Stockham, R. Ingebretsen and T.M. Cannon, IEEE Proc. 62 (1975) 679.
[6] M.L. Sattler and M.A. O'Keefe, in: Proc. 45th Annual EMSA Meeting, Ed. G.W. Bailey (San Francisco Press, San Francisco, 1987).
[7] M.A. O'Keefe, unpublished, private communication, 1994
[8] B.R. Hunt, Proc. IEEE Trans. Comput. C-22, 9 (1973) 805.
[9] P. Xu, D. Dunn, J.P. Zhang and L.D. Marks, Surf. Sci. 285 (1993) L479.
[10] H. Lipson and K.E. Singer, J. Phys. C 7 (1974) 12.
[11] A.A. Baski, J. Nogami and C.F. Quate, Phys. Rev. B 41 (1990) 10247.
[12] N. Doraiswamy, G. Jayaram and L. D. Marks, Phys. Rev. B 51 (1995) 10167.
[13] B. G. Storey, M. A. Kirk and L. D. Marks, Physica C 246 (1995) 4.
[14] P. Schiske, in: Image Processing and Computer-Aided Design in Electron Optics, Ed. P.W. Hawkes (Academic Press, London, 1972) p. 82.
[15] W.O. Saxton, Ultramicroscopy 55 (1994) 171.
[16] L.D.Marks, in preparation.