

A deconvolution method for the reconstruction of underlying profiles measured using large sampling volumes

Y.-S. Xiong and P. J. Withers*

Manchester Materials Science Centre, University of Manchester, Grosvenor St., Manchester M1 7HS, UK. Correspondence e-mail: philip.withers@manchester.ac.uk

Received 11 July 2005
Accepted 4 April 2006

A deconvolution method for diffraction measurements based on a statistical learning technique is presented. The radial-basis function network is used to model the underlying function. A full probabilistic description of the measurement is introduced, incorporating a Bayesian algorithm based on an evidence framework. This method allows predictions of both the convolution and the underlying function from noisy measurements. In addition, the method can provide an estimation of the prediction uncertainty, *i.e.* error-bars. In order to assess the capability of the method, the model was tested first on synthetic data of controllable quality and sparsity; it is shown that the method works very well, even for inaccurately measured (noisy) data. Subsequently, the deconvolution method was applied to real data sets typical of neutron and synchrotron residual stress (strain) data, recovering features not immediately evident in the large-gauge-volume measurements themselves. Finally, the extent to which short-period components are lost as a function of the measurement gauge dimensions is discussed. The results seem to indicate that for a triangular sensor-sensitivity function, measurements are best made with a gauge of a width approximately equal to the wavelength of the expected strain variation, but with a significant level of overlap ($\sim 80\%$) between successive points; this is contrary to current practice for neutron strain measurements.

© 2006 International Union of Crystallography
Printed in Great Britain – all rights reserved

1. Introduction

It is not generally possible to make a measurement at a point. Whatever the measurement technique, the quantity of interest must be evaluated over a sampling gauge having finite dimensions. When considering strain measurement by diffraction, it is often the case that a large gauge volume must be used, either because of low signal, or because of poor spatial discrimination. This has the effect of smearing out the measured values relative to the underlying function. Different conflicting pressures govern the choice of gauge volume. A large gauge volume enables one to acquire statistically significant information quickly and accurately, free from the effects of microstructural fluctuation, but tends to smooth out, and in severe cases lose irrevocably, variations in stress or strain. If the gauge is too small, the sampling time for scattering methods may become prohibitively long, while local microstructural features may make the data unrepresentative of the continuum behaviour.

There is currently much interest in the multi-scale behaviour of materials. Some authors suggest that phenomena such as the size effect in fatigue can be related to the fractal nature of the microstructure (*e.g.* Carpinteri *et al.*, 2002). Of course, just like the microstructure itself, the residual stress field varies in a multi-scale manner. In general, the sampling volumes

typical of diffraction are too large to measure type-II (grain scale) and type-III (sub-grain scale) stress fields, except in terms of peak broadening (Fitzpatrick & Lodini, 2003; Hutchings *et al.*, 2005). Nevertheless, a great deal of information about the micromechanics of heterogeneous deformation can be obtained by diffraction, either at the grain scale (Clausen *et al.*, 1998), sub-grain scale (Ungar *et al.*, 1984) or between constituent phases (Oliver *et al.*, 2004; Wang *et al.*, 2005; Van Acker *et al.*, 1996; Noyan & Cohen, 1985). In general, materials engineers neglect the microstructural stress fluctuation and focus on the type-I (macro-stress) variation. This component is regarded as smoothly varying and compares directly with continuum finite-element models.

The various methods for measuring elastic strain by diffraction all present different obstacles to the derivation of the true (underlying) macro-strain or stress profile. Laboratory X-rays have the advantage that they penetrate very short (micrometre) distances through most engineering materials. This means that the depth over which they average is very shallow; however, in order to build up a depth profile, destructive layer-removal techniques must be employed. High-energy synchrotron X-rays and thermal neutrons both have much higher levels of penetration (Withers, 2004), offering the possibility of non-destructive measurements. For thermal neutrons, a scattering angle of 90° is usually adopted,

but because of the low intensities and long associated counting times characteristic of neutron diffraction, high-spatial-resolution measurements are usually made with a gauge size of 1 mm or larger, with relatively poor counting statistics and sparse data. On the other hand, the very high photon intensities encountered at modern synchrotron sources mean that small gauge volumes can be defined with excellent counting statistics and sampling density. In fact, it is often the case in practice that the fineness of the gauge that can be used is limited by poor microstructural sampling of the type-I stress field. If, on the other hand, effective deconvolution procedures can be established, then large gauge volumes can be used with obvious benefits in terms of data collection efficiency and statistical representation. The higher energies (short wavelengths) typically employed result in low scattering angles. As a consequence, the gauge shape is commonly an elongated diamond (5:1 or 10:1 aspect ratio) with the longest dimension often a millimetre in length or so. Further complications arise because attenuation, which can be significant over the gauge dimension, especially for synchrotron X-rays, means that the gauge volume may not be evenly illuminated.

Unless the data are noiseless, one would not expect to be able to reconstruct features in the stress or strain profile having a characteristic 'wavelength' smaller than the gauge dimensions. Indeed, if the macro-stress is of concern, then the smooth underlying variation is of primary interest and an algorithm for the 'best guess' reconstruction of the underlying profile would have real practical utility. This is especially so for near-surface measurements. This is both because near the surface, the gradients of strain and stress can be very steep, but also because the dimension characterizing the effective sampling gauge varies as the instrumental gauge enters the surface. This means that while the strain profile is smeared out to translator positions that have the centre of gravity of the sampling volume located beyond the surface of the object, it may be possible to reconstruct a profile that is much more representative of the underlying near-surface profile.

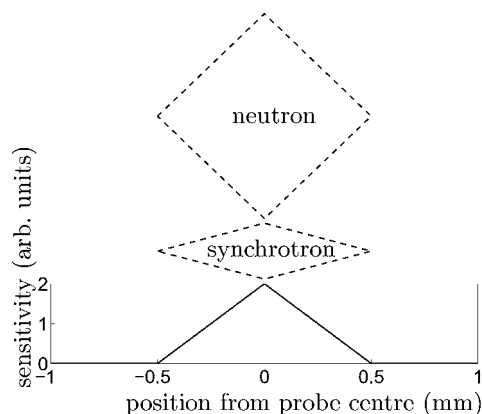


Figure 1
The sensor response function for a typical synchrotron sampling gauge (shown dashed) of length 1 mm neglecting beam attenuation across the gauge volume. A neutron sensor response function would be identical even though the gauge (shown dashed) using a scattering angle of 90° would normally be a square.

Furthermore, in such cases the magnitudes and locations of the near-surface compressive and tensile stresses are often performance critical, and thus of great interest from a structural integrity point of view.

In an ideal (noiseless) case, there is a direct mathematical route from the measured response to the underlying macro-strain or stress profile. However, deconvolution approaches tend to be very sensitive to noise in the measurements, rendering formal deconvolution impractical. Consequently, we have taken a completely different approach. In this paper, we present a method for reconstructing depth profiles based on statistical learning techniques. We apply a probability model to the measured profile, incorporating a Bayesian algorithm based on an evidence framework. The probability model allows predictions for the expectations and error-bars of the underlying strain/stress profile. First, the method is tested on synthetic data in order to assess its capability as a function of measurement sparsity, noise levels and sensor response function (essentially determined by the gauge volume). It is subsequently applied to real data typical of that acquired by neutron and synchrotron strain scanning.

It should be noted that the method is completely general and could be applied to the reconstruction of stress or strain profiles measured by other methods, for example by magnetic methods (Xiong *et al.*, 2006), or indeed more generally to the reconstruction of data of any sort from volume-averaged measurements.

2. Statistical learning method for deconvolution

The framework given below should be applicable to a very wide range of sensors and sensor-sample interactions; however, we shall illustrate its basis and practical utility with reference to the analysis of neutron and synchrotron strain data. Our approach builds upon a rigorous and widely accepted theoretical framework developed by others, which is summarized here only in outline; for more details the reader is directed to the works by MacKay (1992) and Tipping (2001). In order to make the paper accessible, we retain a formal mathematical framework. Those with little interest in the mathematical foundations can omit this section, moving straight to §3.

2.1. Mathematical model of sampling

All practical sensors have a response function $g(x)$ of finite extent. This may vary in both sensitivity and extent as a function of sampling depth. When making measurements by neutron or synchrotron diffraction, the gauge volume over which diffraction data are collected is usually defined by apertures or collimators (Hutchings *et al.*, 2005). This represents the nominal gauge volume and neglects beam divergence and the attenuation of the beam across the sampling gauge. This is quite sufficient in most cases. As a result, when scanning along a line (x), the response function has a triangular form (Fig. 1). If the attenuation length, defined as the distance over which the number of neutrons or photons decreases by

e^{-1} , is comparable to the largest dimension of the gauge, then in reflection geometry at least, a more complex response function would be appropriate that describes the attenuation in illumination over the gauge volume. This level of complexity is neglected here.

In general terms, the sensor-sample signal response can be described mathematically by the convolution

$$y(x) = f(x) * g(x) = \int du g(x - u)f(u), \quad (1)$$

where $f(x)$ and $g(x)$ are signal and sensor response functions, respectively, as they are called in signal processing. In this paper, we will focus on the moving-window-style convolution exemplified by the sampling gauges typical of neutron and synchrotron diffraction shown in Fig. 1. One additional aspect is required. As the sample surface is approached, the effective sampling gauge changes shape, in that no signal is recorded from that portion of the instrumental gauge that lies outside of the sample. As a result, the signal drops in magnitude and the strain is averaged only over the portion of the gauge remaining inside the sample. This can be achieved simply by rewriting the convolution as

$$y(x) = f(x) * g(x) = \frac{\int du g(x - u) \eta(u) f(u)}{\int du g(x - u) \eta(u)}, \quad (2)$$

where $\eta(u)$ is 1 if u is within the sample, and 0 otherwise. The advantage of this modification is that we do not need to adjust the sensitivity of the response function representing the effective volume of the gauge as it passes through the surface of sample.

The aim is to reconstruct the underlying function $f(x)$ based on the measurement $y(x)$ and the given response function, $g(x)$. Reconstruction can be fully achieved by deconvolution of a discrete Fourier transformation (DFT). However, the algorithms developed along this line do not work very well in practice. This is primarily because measurement noise can cause solutions to exhibit unrealistic oscillations. Traditionally, optimal filters have been introduced to solve this problem (Press *et al.*, 1992). Secondly, it should be noted that the noiseless measurement $y(x)$ and the underlying signal $f(x)$ are functions of variables with physical meanings. In our case, the macro-strain (by definition since the local microstructural fluctuations are described by σ_f) and the measured (convoluted) response $y(x)$ have a smooth dependence on position. The deconvolution method based on DFT takes no account of this property.

We model the smooth underlying function as

$$f(x) \rightarrow \sum_{k=1}^K \Phi_k(x)w_k \quad \text{or} \quad \mathbf{f} = \mathbf{\Phi w} \quad (3)$$

where Φ_k is a set of basis functions and w_k are the corresponding weights; K is the number of basis functions. There are various types of basis function available; however, the kernel functions (Vapnik, 1995) are the most successful and are used here. Popular regression models of the kernel function are the radial-basis function networks. Their form can be generally expressed in terms of the Euclidean distance from

their centre x_k , *i.e.* $\Phi(\|x - x_k\|)$. In the statistical learning method, one of the advantages is that the number of weights grows only linearly with the number of training examples. In this work, we use the Gaussian type of radial-basis function, $\Phi(\|x - x_k\|) = \exp[-c(x - x_k)^2]$, where the function is evaluated at position x and the scaling parameter c will be determined by our algorithm. Other radial-basis functions may be used, but the Gaussian type has been proved to be more successful in most cases (Bishop, 1995; Tipping, 2001).

Numerically, the convolution in equation (2) can be written as

$$\begin{aligned} y_m &= [f(x) * g(x)]_m \\ &= \frac{\sum_{n=1}^N g(x_m - u_n) \eta(u_n) f(u_n) \delta u_n}{\sum_{n=1}^N g(x_m - u_n) \eta(u_n) \delta u_n} \\ &= \sum_{n=1}^N G_{mn} f_n \end{aligned} \quad (4)$$

and its matrix form is

$$\mathbf{y} = \mathbf{Gf}, \quad (5)$$

where \mathbf{G} is an $M \times N$ matrix of the given sensor response function, \mathbf{y} is an M -dimensional vector describing the measured convolution, and \mathbf{f} is an N -dimensional vector of the underlying function to be predicted. M is the number of points measured, and N is the number of points used to delineate the underlying function.

Through the convolution relation between the underlying function \mathbf{f} and the measurement \mathbf{y} in equation (5), the measurement can be expressed in terms of the weight parameters

$$\mathbf{y} = \mathbf{G}\mathbf{\Phi w}. \quad (6)$$

The above equation has a form that is very similar to the normal regression model. In principle, we can use the sum-of-the-squares error function to optimize the weights, if given the measured convolution and the forms of the basis function and response function. Once the weight parameters are determined, the underlying function and the measured convolution can be calculated straightforwardly by equations (3) and (6). However, the central goal in modelling is not to obtain a best fit to the training data, but rather to infer the rule (model) generating the data, so that the model can make the best possible predictions for a new input x . The most general and complete description of the data generator (model) is in terms of the probability distribution. On the other hand, the effects of the local fluctuation σ_f and the measurement error σ_y require that the model is probabilistic, *i.e.* a statistical model.

2.2. Statistical description of the convolution process

We consider the measurement of the variation of any underlying property, $f(x)$, using a sensor. In practice, the property may show local fluctuations, perhaps because of local heterogeneities in structure. In our exemplar case, this might be due to the grain-to-grain fluctuations in strain arising from the anisotropy in single-crystal stiffness and plasticity of real crystals. In our treatment we have modelled these (type-II and type-III) fluctuations by a Gaussian distribution having a

standard deviation σ_f . Classically, if a sufficiently large sampling gauge is used, these fluctuations contribute to line broadening (Noyan & Cohen, 1987); however, in practical cases at high spatial resolution, it is not uncommon to measure repeatable point-to-point fluctuations in response arising from grain-sampling effects (Withers, 2003). As discussed in the introduction (§1), while much can be gained from knowledge of the local fluctuations and their standard deviation σ_f from a materials viewpoint, in many cases it is the smoothly varying type-I profile, $f(x)$, that is of engineering interest.

We assume that the probability of the function having a certain form is normally distributed:

$$P(\mathbf{f}|\mathbf{w}) = \frac{1}{[(2\pi\sigma_f^2)^N]^{1/2}} \exp\left(-\frac{1}{2\sigma_f^2} \|\mathbf{f} - \Phi\mathbf{w}\|^2\right), \quad (7)$$

where σ_f represents the point-to-point deviation from the underlying function, in our case having microstructural origins as discussed above. More specifically, this means that the value $f(x)$ detected by the gauge is given by some deterministic function upon which is superimposed a random fluctuation of standard deviation σ_f . This local scatter in the strain can be estimated directly by our algorithm. In other words, the model is able to make some prediction regarding the level of micro-strain in the sample. This is somewhat different from the measure of the micro-strain given by the diffraction line width. It is the variation in the volume-averaged strain that would be measured using a gauge equal in dimension to the spacing between the points used to define the underlying function, namely the N points in equations (4) and (5). As $N \rightarrow \infty$, it describes the variation in the total profile.

In addition, we assume that the M measurements of \mathbf{y} are generated by sampling a smooth function with additive random sensor noise σ_y , so that the probability distribution function of convolution \mathbf{y} for a given function \mathbf{f} has a Gaussian form:

$$P(\mathbf{y}|\mathbf{f}) = \frac{1}{[(2\pi\sigma_y^2)^M]^{1/2}} \exp\left(-\frac{1}{2\sigma_y^2} \|\mathbf{y} - \mathbf{G}\mathbf{f}\|^2\right), \quad (8)$$

where σ_y is the standard deviation of the statistical error associated with measuring \mathbf{y} . In general, the statistical measurement uncertainty is unknown, but can be estimated through our algorithm. In the case of synchrotron and neutron diffraction, the uncertainty in peak position can be estimated in terms of the best fit peak profile to the measured diffraction data, or evaluated in terms of the number of counts in the signal I_n and the standard deviation of the diffraction peak width σ_{diff} as (Withers *et al.*, 2001)

$$\sigma_y = \sigma_{\text{diff}} [1 + 2(2^{1/2})B/H]/(I_n)^{1/2}, \quad (9)$$

where H is the peak height and B is the background.

An important concept in Bayesian inference is that of ‘marginalization’, which involves integrating out unwanted variables (Bishop, 1995). Suppose we are discussing a model with two variables f and y . Then the most complete description is in terms of the joint distribution $p(f, y)$. If we are interested only in the distribution of y , then we should integrate out f .

Thus, the predictive distribution for y is obtained by averaging the conditional distribution $p(y|f)$ with a weighting factor given by the distribution $p(f)$. Thus, the conditional probability distribution of the convolution \mathbf{y} for a given weight vector, \mathbf{w} , is obtained by integrating over \mathbf{f} :

$$\begin{aligned} P(\mathbf{y}|\mathbf{w}) &= \int_{-\infty}^{+\infty} d\mathbf{f} P(\mathbf{y}, \mathbf{f}|\mathbf{w}) \\ &= \int_{-\infty}^{+\infty} d\mathbf{f} P(\mathbf{y}|\mathbf{f}) P(\mathbf{f}|\mathbf{w}) \\ &= \frac{1}{[|\mathbf{C}|(2\pi)^M]^{1/2}} \exp[-(\Delta_{\mathbf{w}}^T \mathbf{C}^{-1} \Delta_{\mathbf{w}}/2)], \end{aligned} \quad (10)$$

with

$$\Delta_{\mathbf{w}} = \mathbf{y} - \mathbf{G}\Phi\mathbf{w}, \quad \mathbf{C} = \sigma_y^2 \mathbf{I} + \sigma_f^2 \mathbf{G}\mathbf{G}^T, \quad (11)$$

where $\Delta_{\mathbf{w}}$ is the difference between the measurements and the model predictions, superscript T is the transpose operator, \mathbf{C} is the covariance matrix of modelled total statistical error, \mathbf{I} is the M -dimensional unit matrix, and \mathbf{G} is defined in equations (4) and (5).

In the statistical description, the output of the convolution is expressed in terms of the most probable value of the distribution in equation (10). Compared with the equation (6), the expectation value of the statistical model is the same as the output of the conventional model. There is a very clear physical meaning of the covariance matrix \mathbf{C} in equation (11), which describes the propagation of the local fluctuations during the convolution process (actually there are some eliminations of the local fluctuations through convolution). The total uncertainty of the measurement comes from two parts: the measurement error σ_y and the convolution of the local fluctuation σ_f .

2.3. Bayesian algorithm

In essence, we are building up the function from a series of Gaussian functions. While it is trivial to fit a data set of M measurements with M Gaussian curves in equation (6) by minimizing the cost function $\chi^2 \equiv \|\mathbf{y} - \mathbf{G}\Phi\mathbf{w}\|^2$, this would lead, in all probability, to over-fitting by which the noise, as well as the signal, are fitted, giving poor predictive capability. In order to avoid over-fitting of the noise in the data and to yield simple (smooth) underlying functions, weight decay (Bishop, 1995) is used. This introduces a tendency for all weights to decay to zero unless there is strong evidence to the contrary.

For the Bayesian algorithm (this has been discussed at length elsewhere), the prior probability distribution of weights, \mathbf{w} , has the Gaussian form (MacKay, 1992; Bishop, 1995)

$$P(\mathbf{w}) = \prod_{k=1}^K P(w_k) \equiv \prod_{k=1}^K \{(\alpha_k/2\pi)^{1/2} \exp[-(\alpha_k w_k^2/2)]\}, \quad (12)$$

The weight decay α_k is connected to weight w_k , and can be written as a diagonal matrix $\mathbf{A} = \text{diag}(\alpha_1, \alpha_2, \dots, \alpha_K)$. One weight decay α_k is given for each weight w_k , as used by Tipping

(2001), with the weight w_k being zero if $\alpha_k \rightarrow \infty$. In practice, most weights are decayed towards zero, and the model solution generally has fewer Gaussian functions ($\ll M$) to capture the trend while discounting the noise as best as it can. Note that this approach is not always successful; for example in Fig. 4(d), the model has overestimated the noise and underestimated the amplitude of the underlying function, in contrast to Fig. 4(f) where the model is informed of the measurement uncertainty as prior knowledge.

Through Bayes' theorem, upon observing data D , the posterior probability of weights \mathbf{w} is then

$$p(\mathbf{w}|D) = \frac{1}{[(2\pi)^K |\mathbf{S}|]^{1/2}} \exp\{-[(\mathbf{w} - \boldsymbol{\mu})^T \mathbf{S}^{-1}(\mathbf{w} - \boldsymbol{\mu})/2]\}, \quad (13)$$

with

$$\boldsymbol{\mu} = \mathbf{S}\boldsymbol{\Phi}^T \mathbf{G}^T \mathbf{C}^{-1} \mathbf{y}, \quad \mathbf{S} = (\boldsymbol{\Phi}^T \mathbf{G}^T \mathbf{C}^{-1} \mathbf{G}\boldsymbol{\Phi} + \mathbf{A})^{-1}. \quad (14)$$

In order to obtain optimal estimations of c , weight decay α_k and noise levels σ_f and σ_y , we apply a Bayesian evidence framework (MacKay, 1992; Tipping, 2001). This framework defines a cost function called the evidence,

$$\int_{-\infty}^{+\infty} d\mathbf{w} P(\mathbf{y}|\mathbf{w}) P(\mathbf{w}), \quad (15)$$

and determines c , weight decay and noise levels iteratively by maximizing the evidence. For convenience, the logarithm of the evidence is maximized generally:

$$\sum_{k=1}^K \ln \alpha_k - \ln |\mathbf{C}| + \ln |\mathbf{S}| - \sum_{k=1}^K \alpha_k \mu_k^2 - \Delta_{\boldsymbol{\mu}}^T \mathbf{C}^{-1} \Delta_{\boldsymbol{\mu}} \quad (16)$$

with $\Delta_{\boldsymbol{\mu}} = \mathbf{y} - \mathbf{G}\boldsymbol{\Phi}\boldsymbol{\mu}$. We set the derivatives of the log-evidence with respect to α_k to zeros and obtain

$$\alpha_k = \gamma_k / \mu_k^2, \quad \gamma_k = 1 - \alpha_k \mathbf{S}_{kk} \quad (k = 1, \dots, K). \quad (17)$$

Since there are no explicit expressions to iterate c , σ_f and σ_y , we will use the gradient of the log-evidence to update these parameters.

In the Bayesian algorithm, the trained model is described in terms of the posterior probability distribution, and the prediction will be made by averaging the probability distribution of the output over all possible models [*i.e.* integrating their probabilities over the posterior probability of the weights in the equation (13)]. When presented with inputs x_m and x_n (the coordinates of the position in our case), the outputs for the measurement y_m and the underlying function f_n are described by probability distributions. The predictions of our model are the most probable values (*i.e.* the means) of these probability distributions:

$$\bar{y}_m = [\mathbf{G}\boldsymbol{\Phi}\boldsymbol{\mu}]_m, \quad \bar{f}_n = [\boldsymbol{\Phi}\boldsymbol{\mu}]_n, \quad (18)$$

with $\boldsymbol{\mu}$ the optimized weight vector given by equation (14). The forms of the solutions in equation (18) are the same as those in equations (3) and (6), while the difference lies in how the optimal weights are calculated. The estimated uncertainties of the predictions of y_m at x_m and f_n at x_n , *i.e.* the respective variances, are given by the variances of the probability distributions:

$$\sigma_{f_n}^2 = \sigma_f^2 + [\boldsymbol{\Phi}\mathbf{S}\boldsymbol{\Phi}^T]_{nn}, \quad \sigma_{y_m}^2 = C_{mm} + [\mathbf{G}\boldsymbol{\Phi}\mathbf{S}\boldsymbol{\Phi}^T \mathbf{G}^T]_{mm}. \quad (19)$$

In order to implement this algorithm, we take the number of the basis functions K to be the same as the number of the convoluted measurements M . The centres of the basis functions x_k are evenly distributed over the sample. We first choose initial values of c and hyperparameters α_k , σ_f and σ_y . Through the covariance matrix of the output \mathbf{C} in equation (11), the weights can be obtained by equation (13). With the optimal weights $\boldsymbol{\mu}$ and the covariance matrix \mathbf{S} , we can improve our estimation of α_k through equation (17), and σ_f , σ_y and c through gradient descent with the log-evidence in equation (16) as the error function. We start the next iteration with the estimated values of hyperparameters until a tolerance criterion is reached. Finally, the expectation value and the uncertainty of the predicted convolution and underlying function can be calculated through equations (18) and (19), respectively.

For any given data set, the number of solutions \mathbf{f} fitting the data \mathbf{y} will be infinite. From the point of view of the Bayesian algorithm, the probability of solutions that have large oscillations is very small, though their number may be large. The most probable solution of the Bayesian algorithm is generally rather smooth and simple, which reduces the risk of giving a seriously 'wrong' answer. However, for a given measured profile (convolution) and sensor response function $g(x)$, obtaining the deconvolution is complex as the underlying function cannot be measured directly. The components with short wavelength of variation have been averaged out by the response function; as a result, there is some information loss through convolution. The lost information cannot be fully retrieved by any method. In such a case, the Bayesian algorithm tries to use a smooth function $f(x)$ to fit the observed convolution $y(x)$, so that the probability of making a large error in predicting $f(x)$ is small. As a result, the prediction of the underlying variation will always tend to be 'conservative' in the sense that the actual underlying variation may be more extreme. Consequently, the prediction is likely to be 'non-conservative' in the engineering sense, in that the underlying function maybe more extreme. However, the predicted variation is likely to be less non-conservative than the response actually measured. As a consequence, the variance σ_f^2 predicted by equation (17) is not likely to be accurate if the underlying function contains significant components having wavelengths smaller than the width of the sensor window. Of course the materials engineer will try to choose the gauge dimensions such that important variations in the underlying profile are not lost in this way. Smaller scale variation will be captured in the diffraction peak width. A great deal of work has been done using Fourier analysis of diffraction peak line broadening to obtain the root mean square (r.m.s.) strain as a function of sampling length as outlined by Noyan & Cohen (1987).

3. Undertaking scanning as efficiently as possible

Clearly, given the high cost and demand for synchrotron and neutron beams, it is important to make strain measurements to

sufficient accuracy to estimate the underlying profile as quickly as possible. For a symmetrical diamond-shaped gauge of the type considered here in Fig. 1 (*i.e.* as would be determined by apertures of equal sizes defining the incoming and outgoing beams), the number of counts per unit time is proportional to the square of the sensor (gauge) width (W^2) for a specific measurement geometry (scattering angle θ). The number of measurements per unit length of scan line is given by $1/\Delta x$, where Δx is the spacing between successive measurements. Using equation (9), the uncertainty in peak position is given by

$$\sigma_y \simeq \sigma_{\text{diff}}/(I_n)^{1/2} \quad (20)$$

if the peak height to the background ratio is high. As a consequence, the time to acquire a measurement to a given accuracy σ_y is proportional to

$$\sigma_{\text{diff}}^2/(W^2\sigma_y^2). \quad (21)$$

As a result, the time to complete a scan is proportional to

$$\sigma_{\text{diff}}^2/(\Delta x W^2\sigma_y^2). \quad (22)$$

Thus for a given diffraction peak width (σ_{diff}) we want to maximize $\Delta x W^2\sigma_y^2$ in order to undertake a scan as quickly as possible consistent with obtaining a good estimator of the underlying profile.

4. Application of the deconvolution algorithm to synthetic data

The chief advantages of synthetic data when assessing the performance of any reconstruction algorithm are that one ‘knows’ the underlying function that one is trying to estimate and one can systematically vary the quality of the data. Here Gaussian ‘noise’ with standard deviation σ_f has been added to the smooth underlying function $f(x)$ characteristic of the type-I macro-stress (strain) profile at a series of N discrete points. As discussed earlier, in practice such point-to-point fluctuations might arise from microstructural heterogeneities.

A sensor response function, $g(x)$, taking the form of a triangle (as seen in Fig. 1), has been used to represent the sampling capability of the sensor. This has then been convoluted with the ‘noisy’ underlying function specified at N ($= 200$) points at each of M measurement points, spaced Δx apart to give the best possible measured response. This represents the measured signal were there no noise associated with the sensor. Finally, the measured training data were obtained by adding conventional Gaussian measurement noise (σ_y) to the response to represent statistical error associated with the sensor. In the case of diffraction, this might be due to the use of poor counting statistics leading to poor determination of the diffraction peak centre and is represented by equation (9). While in practice the level of peak fit uncertainty is often evaluated by the peak fitting routine, except where stated no information about the likely level of measurement uncertainty was input into the model.

4.1. Test example I

As a first test case, we have taken the underlying function (type-I strain) to be

$$f(x) = \sin(\pi x)/(\pi x), \quad (23)$$

with the boundaries of the strained object located at $x = 0$ and $x = 6$. This function has a very simple frequency distribution, with decreasing amplitude of signal with increasing x .

In order to assess the performance of the Bayesian framework compared with a conventional χ^2 minimization, equation (6) was first solved using the cost function $\chi^2 \equiv \|\mathbf{y} - \mathbf{G}\Phi\mathbf{w}\|^2$. The results are shown in Figs. 2(a) and 2(b). We can see that the model predicts very poorly the underlying function, even though the predicted convolution passes through each measurement almost exactly. It is also observed that the predicted convolution (Fig. 2a) exhibits higher oscillation near the boundaries of the sample than within the sample. This is because the effective volume becomes smaller when the centre of the gauge leaves the sample, and the contribution to the cost function is exaggerated by the measurements with most of the gauge outside of the sample. A more reasonable fit could be obtained by reducing the number of basis functions, but that would require prior knowledge of the likely form of the underlying profile. If, on the other hand, a weight decay term $0.01\|\mathbf{w}\|^2$ is added manually to the above cost function, the prediction on the measured convolution is improved somewhat without prior knowledge (see Fig. 2c). However, while the curve through the measurements appears reasonable, the estimate of the underlying function is far from satisfactory (see Fig. 2d). Further increasing the weight decay would compromise the features in the measured response. In the Bayesian algorithm, the weight decays α_i (connected to each weight) are set in an optimal iterative manner. The weights that cause large oscillation are most likely to be removed; consequently the performance of the Bayesian model is far superior (Figs. 2e and 2f), although it does take longer to run (~ 30 s) than the conventional χ^2 minimization. Note that the model is closer to the smooth underlying function than the original noisy (microstructurally fluctuating) profile (open symbols in Fig. 2f).

Fig. 3 shows the comparison between the measurements and the Bayesian model predictions for a fixed sampling interval $\Delta x = 0.2$ and different noise levels and window widths. Unsurprisingly for the noise-free measurements with window width $W = 1$, our method provides a good estimate of the sensor response (continuous curve in Fig. 3a) and reconstructs the unmeasured underlying function almost exactly (Fig. 3b) with very low noise estimates. With a Gaussian measurement uncertainty (noise) level of $\sigma_y = 0.1$ and a Gaussian fluctuation in the underlying function of $\sigma_f = 0.1$, the errors of the model predictions increase. Nevertheless, the method is able to predict accurately the convoluted measured response were there no noise (Fig. 3c) and the deconvolution is able to reproduce details of the underlying function very well, even within the region of smaller amplitude where the scatter is as large as the signal. Note that the data points in Fig. 3(c) are

generated using the same conditions as those in Fig. 2, and the differences are due to the randomness of noise added. The predicted convolution profiles are therefore slightly different, but the predictions of the underlying profile are in close agreement, certainly within the error-bars. Note that here the error-bars are somewhat smaller than in Fig. 2(f), reflecting greater confidence in the prediction.

If we increase the window width to 4, the predicted underlying function has good accuracy towards $x = 0$ due to

the improvement in spatial resolution at a surface where the gauge is only partially filled; however it loses detail in regions of smaller amplitude (as seen in Figs. 3e and 3f). As is seen from the dashed line in Fig. 3(e), the smearing effect of the large window means that the oscillations in the underlying function are lost. As a result, even though the best estimate of the convoluted profile (solid line) is close to the noiseless one (dashed in Fig. 3e), the deconvoluted prediction cannot recover the features in the underlying function Fig. 3(f). This worsening of performance is captured by the large error-bars and the predicted profile is no worse than the microstructurally affected total signal (the open circles in Fig. 3f) as an estimate of the underlying function, though it is more conservative. The fact that the solid lines in Figs. 3(e) and 3(f) are so similar indicates that the model does not seriously amplify noise-related fluctuations, *i.e.* introduce more detail than it should in its predictions of the underlying response. In both Figs. 3(e) and 3(f), the error-bars are reasonable estimates of the point-to-point scatter.

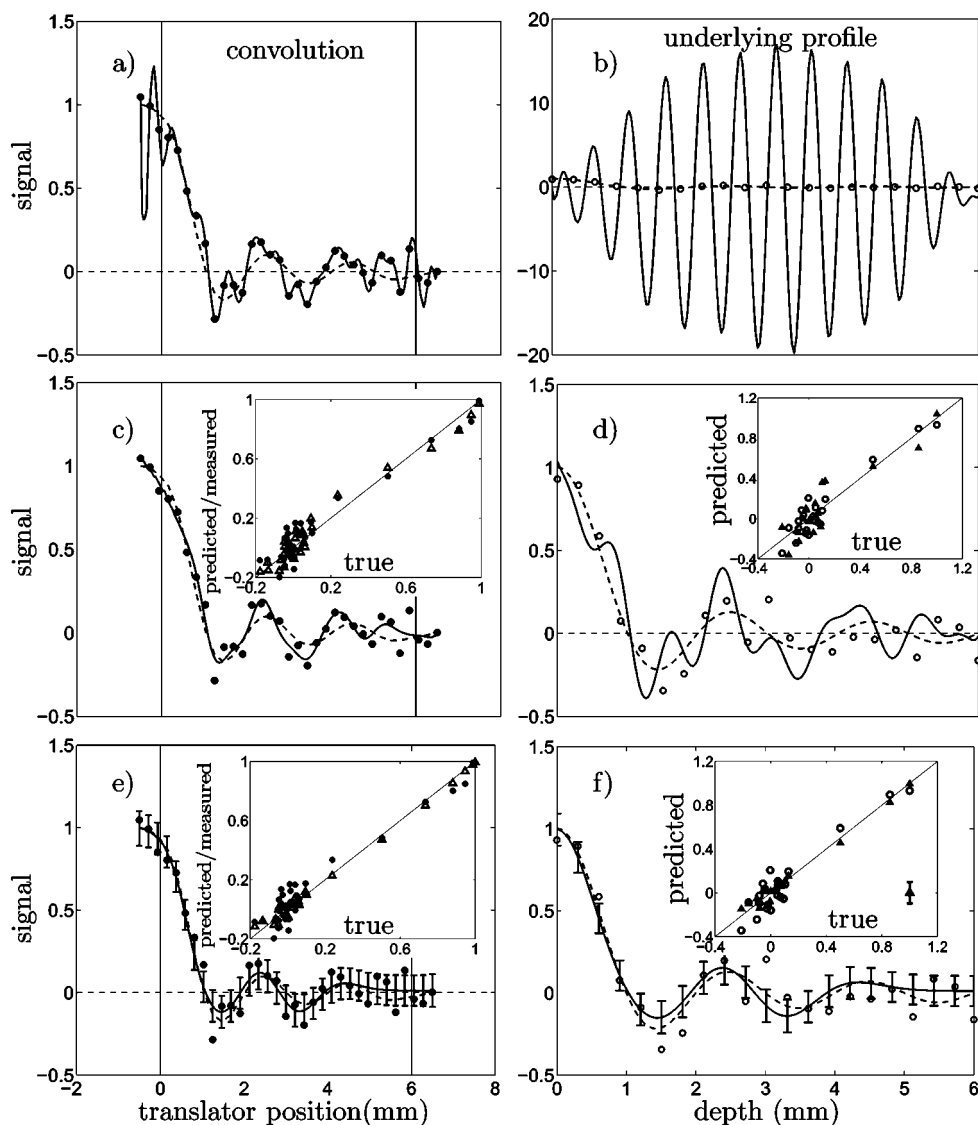


Figure 2

The right-hand panels, (b), (d) and (f), show the underlying smooth function, $\sin(\pi x)/(\pi x)$ (dashed line), along with the actual microstructurally fluctuating ($\sigma_f = 0.1$) values (open symbols). The convolution of the smooth profile with the sensor is shown in the left-hand panels, (a), (c) and (e) (dashed line), along with the noisy ($\sigma_y = 0.1$) measurements (full symbols). In (a) a conventional χ^2 minimization has been applied (solid line) to give an estimate of the underlying profile (solid line) in (b). In (c) a $0.01\|\mathbf{w}\|^2$ weight decay has been added to the χ^2 cost function, while in (e) a Bayesian algorithm has been used to estimate the underlying function (dashed line) in (f). The number of points at which the underlying function f is specified, N ($= 200$), is not important so long as the integrations in equation (2) are accurate enough numerically; for clarity, only a fraction of the 200 points are shown (by open symbols) in the right-hand figures. The insets in (c), (d), (e) and (f) compare the predicted values (triangles) and the data points (circles) against the noiseless true profile. All the predictions (triangles) of the Bayesian algorithm have associated error-bars which have been omitted for clarity, but a typical error-bar is shown at the bottom-right of the inset (f).

For fixed values of noise level, the effect of input noise σ_f is less significant than the sensor noise σ_y . This can be understood as the effect of the response function is essentially that of a moving-window-average, which can eliminate (random) microstructural fluctuations of the signal. Unsurprisingly, the predictions on the observed values (*i.e.* the profile measured by the sensor without sensor noise) are generally better than those on the non-observables (*i.e.* the underlying function), as information in the underlying signal is partly lost through the convolution.

4.2. Test example II

As a second test case, we have taken the underlying function to be

$$f(x) = \cos[\pi x(x + 1)/2], \quad (24)$$

over the range 0 to 4. This function has a richer spectrum structure than $\sin(\pi x)/(\pi x)$ over the defined region, since the half period decreases from essentially 1 to 1/2 to 1/3 to 1/4, *etc.*, between each successive minimum.

Figs. 4(a) and 4(b) show the comparison between the measurements and the model predictions using a sampling interval $\Delta x = 0.05$ and noise levels $\sigma_y = 0.1$, $\sigma_f = 0$ and a window width of $W = 1$. We can see from Fig. 4(b) that the prediction is very good over the longer periods, becoming unacceptable at around $x = 2.5$ ($\lambda \simeq 2/3$). The smearing action of such a wide window ($W = 1$) has filtered out Fourier components with shorter periods. Good performance was recovered near the surface ($x = 4$) despite the very sharp variations there ($\lambda \simeq 1/3$). This is because, as noted earlier, as the gauge leaves the sample, signal is received from only a small part of it, thus decreasing the effective sensor size.

One might expect that a more complete recovery of the underlying function could be achieved by increasing the spatial definition and lowering the uncertainty of the measurements. Figs. 4(c) and 4(d) show the equivalent response using a sensor of higher spatial resolution ($W = 0.2$) along with lower measurement noise levels, $\sigma_y = 0.01$, and a larger

sampling interval, $\Delta x = 0.2$. This corresponds to the case where one slides the sensor across by a distance equal to the width of the sensor in each increment, as is often the case when scanning with neutron diffraction. Accordingly, it would be 5 times quicker to complete the scan because of the reduced measurement frequency, but it would take 25 times longer to acquire each point because the dimensions are 5 times smaller. In fact, because of the much lower measurement uncertainty used here, one would need to acquire 100 times the signal at each point in order to achieve the envisaged 10 times improvement in measurement accuracy [equation (20)]. As a result, the overall acquisition time would be 500 times that for the first setup. Despite this, the response is certainly

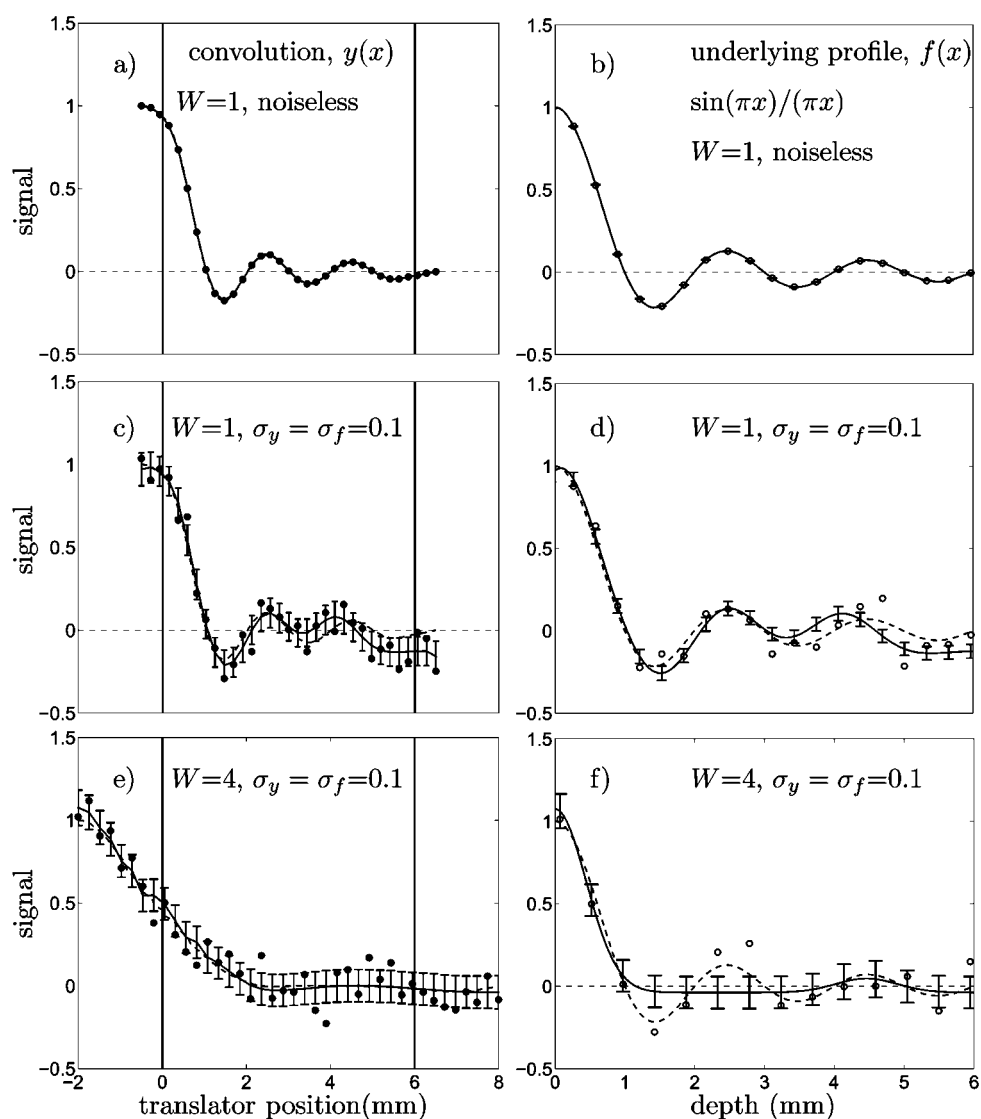


Figure 3 The convolution using window width W of a triangular response function and sampling interval $\Delta x = 0.2$ with sample boundaries at 0 and 6. The left-hand panels (a), (c) and (e) show the measured profile, and the right-hand panels (b), (d) and (f) are for the underlying function, $\sin(\pi x)/(\pi x)$. The open symbols (right-hand side) indicate the microstructural fluctuations in the underlying function; the full symbols (left-hand side) show the measured points. The broken lines are the underlying function (right-hand side) and the noiseless convolution (left-hand side) and the full lines with error-bars are the corresponding predictions of the method. The number of points of the underlying function f is $N (= 200)$; for clarity, only a fraction of the 200 points are shown (by open symbols) in the right-hand figures.

no better than the previous case. This is because without prior knowledge of the improved measurement uncertainty, the sampling rate is too sparse to capture the features of the underlying profile. Consequently, the model has inferred that the point-to-point variation in the measurements is due to noise in the signal rather than any underlying oscillation in the data. As a consequence, the prediction is of a high measurement uncertainty (large error-bars) and a conservative estimate of the strain profile amplitude. In fact, in this case, even near the surface ($x = 4$) the performance is poor due to the sparseness of the data in this region.

Of course, in practice when making strain measurements, one can usually infer the uncertainty in the measurement from

the statistical goodness of fit of the Gaussian or other peak profile to the diffraction peak data. Furthermore, Withers *et al.* (2001) have shown that the uncertainty in the measurement of diffraction angle can be expressed analytically in terms of the diffraction peak width and the number of counts under the peak [equation (9)]. Consequently, provided that the systematic errors are not too large, one can obtain a very good estimate of the measurement uncertainty σ_y . This information can be input into the model directly rather than allowing the model to deduce the noise level on the basis of the evidence. This is especially important where the data are sparse, because it is difficult for the model to differentiate between wavelength components of the underlying function having a spacing less

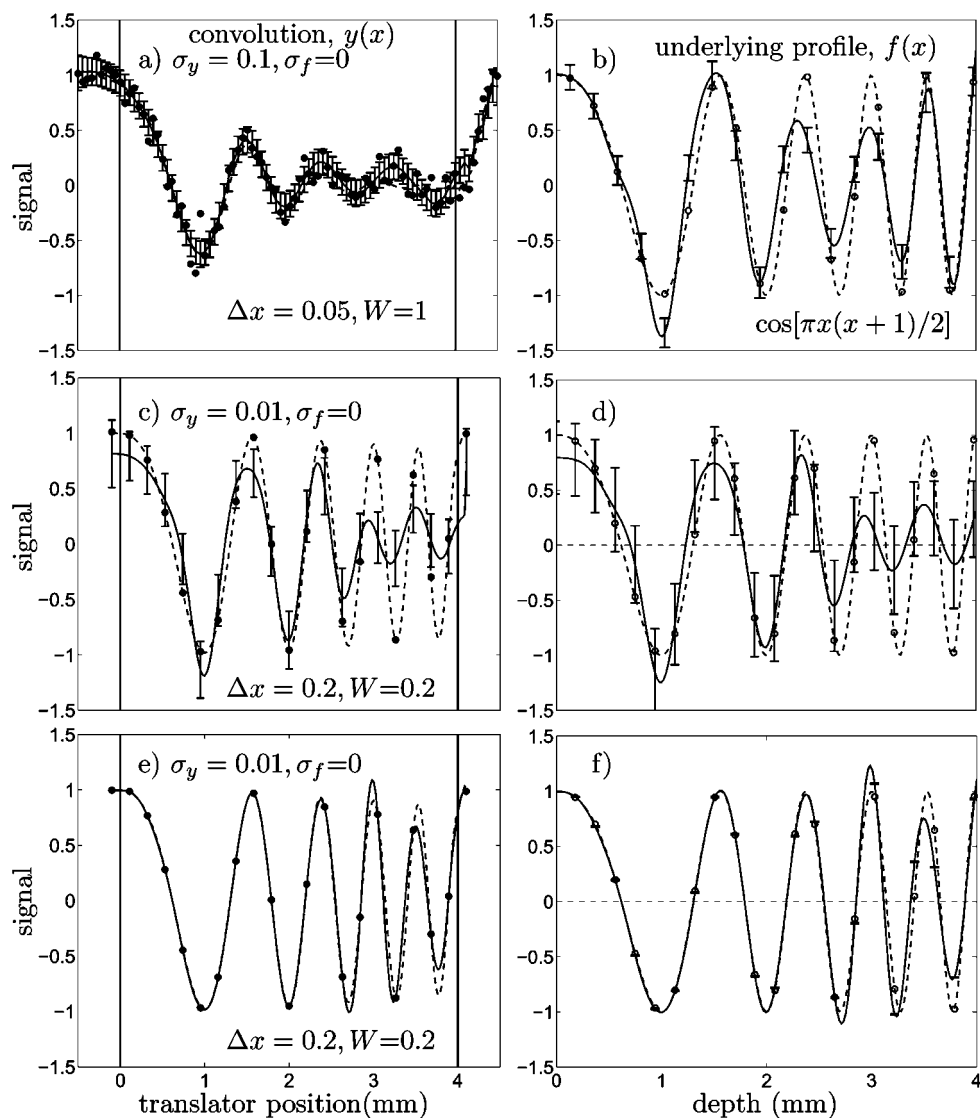


Figure 4 The convolution using window width W of a triangular response function and sampling interval Δx with sample boundaries at 0 and 4. The left-hand panels, (a), (c) and (e), show the measured profile, and the right-hand panels, (b), (d) and (f), are for the underlying function, $\cos[\pi x(x + 1)/2]$. The open symbols (right-hand side) indicate the microstructural fluctuations in the underlying function; the full symbols (left-hand side) show the measured points. The broken lines are the underlying function (right-hand side) and the noiseless convolution (left-hand side), and the full lines with error-bars are the predictions of the method. In (e) and (f), the measurement uncertainty is known and imposed on the model. For clarity, only some of the points defining the underlying function are shown in the right-hand figures.

than or equal to the spacing of the data and noisy fluctuations. Figs. 4(e) and 4(f) show the comparison between the measurements and the model predictions for the same data as in Figs. 4(c) and 4(d). In this case we have input the measurement error ($\sigma_y = 0.01$) directly into the model as prior knowledge instead of inferring it according to equation (17) based on the evidence. The model now performs much better in the light of this new information. In fact, the model is able to predict the underlying function across the whole profile, except in the very-near-surface region where the sparseness of the data prevents accurate deconvolution. In other words, with the prior knowledge that the measurement error is very small, the algorithm simply selects a solution from functions that fit the data points well, allowing shorter wavelength components that would otherwise be viewed as unlikely.

4.3. Assessing the performance over a sine wave

In order to chart the performance of the model more quantitatively, it is useful to consider the underlying test function

$$f(x) = A_f \sin(2\pi x/\lambda), \quad (25)$$

having a constant amplitude A_f and the object boundaries at infinity, since this contains just one wavelength λ . In this case, we can calculate the convolution through

the triangular window (W) analytically:

$$y(x) = A_y(W/\lambda) \sin(2\pi x/\lambda). \quad (26)$$

It can be seen that the convolution has the same form as the underlying function except that the amplitude has the value

$$A_y(W/\lambda) = A_f \left[\frac{\sin(\pi W/2\lambda)}{\pi W/2\lambda} \right]^2, \quad (27)$$

which varies according to the window width, W , normalized by the wavelength of the underlying function. This amplitude is plotted in Fig. 5, where it decreases very quickly as the normalized window width increases. This means that the data measured by a larger window are more vulnerable to the measurement error, since they have a higher noise–signal ratio for a given absolute noise level. Since $A_y(W/\lambda)$ is zero for $W = 2k\lambda$ ($k = 1, 2, \dots$) the information of the underlying function is blocked completely for these window sizes (note the tendency for this in Fig. 7). For a general underlying function having a continuous spectrum, the components with wavelengths $\lambda = W/2k$ will be completely lost. These cannot be retrieved by any method, even for the idealized case of complete sampling, $\Delta x \rightarrow 0$, and no noise, $\sigma_y = 0$. The subsidiary maxima of $A_y(W/\lambda)$ are very small (*i.e.* for $W/\lambda > 2$), so that Fourier components of wavelengths shorter than $W/2$ can be regarded as irrecoverable in most practical applications.

In order to assess the capability of the method and to verify these arguments, we have tested the performance of the method for the underlying function in equation (25) under various conditions. In each trial we have generated data and trained the model 40 times for each set of measurement parameters, namely Gaussian noise level σ_y ($\sigma_f = 0$ in all simulations), sampling interval Δx and window width W . In each case, we have calculated the mean error for the convolution and for the underlying function over the 200 specified locations. In each case, the actual error is the r.m.s difference between the prediction [equation (18)] and the true value of the convolution or the underlying function, while the predicted error is given by equation (19). For each trial, we calculate the average and standard deviation of the resulting errors as the expectation and uncertainty, respectively.

4.3.1. Performance versus measurement noise level. The averaged values of the actual error and the predicted error for the convolution ($\overline{\Delta y}/A_y$) and underlying function ($\overline{\Delta f}/A_f$) are compared in Figs. 6(a) and 6(b) for different measurement noise levels (σ_y/A_y) for fixed window width $W = 0.25\lambda$ and sampling interval $\Delta x = 0.05\lambda$. It is found that the errors vary approximately linearly with the measurement error σ_y/A_y . The 1:1 gradient in Fig. 6(a) suggests that the inferred error is almost exactly the same as the noise actually added (*i.e.* σ_y/A_y), with the consequence that it can be used as an accurate indicator of the measurement error. The actual error of the convolution in Fig. 6(a) is smaller than inferred by the model, as expressed by the error-bars. It is interesting that the measurement error feeds back into our inferred error of the underlying function [*i.e.* we have a 1:1 gradient in Fig. 6(b) for

the inferred errors]. Once again, the actual error tends to be smaller than the inferred error-bars, indicating a good estimation of the underlying function in most cases.

Figs. 6(c) and 6(d) show the actual averaged error in reconstructing the underlying function at different noise levels for fixed window width and sampling interval. A similar linear dependence of the error on measurement noise levels (σ_y/A_y) is observed to that in Figs. 6(a) and 6(b). In the case of $\Delta x = 0.25\lambda$ and $W = 0.25\lambda$, a very poor estimation is achieved and the error-bars are much larger than the scatter in the measurements. This corroborates the observations associated with Figs. 4(c) and 4(d) that for sparse data ($\Delta x \geq W$) there is a tendency for the model to infer a low level of variation in the response by overestimating the noise level. Figs. 6(e) and 6(f) show the actual error of deconvolution for the same conditions as in Figs. 6(c) and 6(d), except in this case we have input the measurement error (σ_y/A_y) directly into the model as prior knowledge, rather than adjusting it according to equation (17) based on the evidence. The expected improvement of performance, as in Fig. 4, occurs only in the case of $\Delta x = 0.25\lambda$ and $W = 0.25\lambda$ in the region of small measurement error (σ_y/A_y). It can be understood that the model cannot track the true value of the convolution accurately in the presence of higher measurement error. This suggests that in cases where the data are sparse (*e.g.* when acquiring neutron diffraction data), it may be necessary to make more accurate measurements and to input the statistical measurement uncertainty into the model as prior information to achieve satisfactory performance. This has an associated penalty in measurement time [equation (22)] and a better strategy may be to use a larger window with overlapping sampling $\Delta x < W$ (see §6). When $\Delta x < W$, the simulations suggest that the model is able to make a good estimate of the actual measurement noise.

4.3.2. Performance versus sampling interval. According to the sampling theorem summarized by Press *et al.* (1992), at least two sample points per period are required to determine a




Figure 5
The normalized amplitude of the convolution of the underlying function $\sin(2\pi x/\lambda)$ with the sensor response as a function of normalized sensor-window width (W/λ).

sine wave. As the period of convolution is the same as the underlying function, it seems to be sufficient to sample at an interval of $\Delta x = \lambda/2$ for the convolution; however, this does not work in practice. Due to the limited resolution of the gauge, measurement error will result in large oscillations to the underlying function during deconvolution. Consequently, the deconvolution method needs greater sampling in order to avoid possible over-fitting of the noisy data, and to reproduce smooth functions for both the convolution and the underlying function.

With regard to the effect of sampling frequency Δx , Figs. 7(a), 7(c) and 7(e) show the actual mean error of the deconvolution for different sampling intervals at various window widths and noise levels. We can see that there is a sharp deterioration in performance for Δx greater than around

0.22λ (*i.e.* when the sampling rate is less than 5 per cycle). It is harder to recognize signal from noisy measurements at low sampling rates, and in such cases the profile amplitude will be underestimated and the noise level overestimated as discussed above. On the other hand, the results also show that increasing the sampling rate much below 0.2λ leads to relatively mild improvements in performance. As a conservative estimation from the simulations, the sampling interval (Δx) should be just less than $\lambda/5$ in order to obtain an acceptable (90%) deconvolution in most cases.

4.3.3. Performance versus window width. With regard to the effect of window width W/λ , Figs. 7(b), 7(d) and 7(f) show the actual mean error of the deconvolution for different window widths at various noise levels and sampling intervals. For small window widths ($W < \lambda$), the actual errors increase

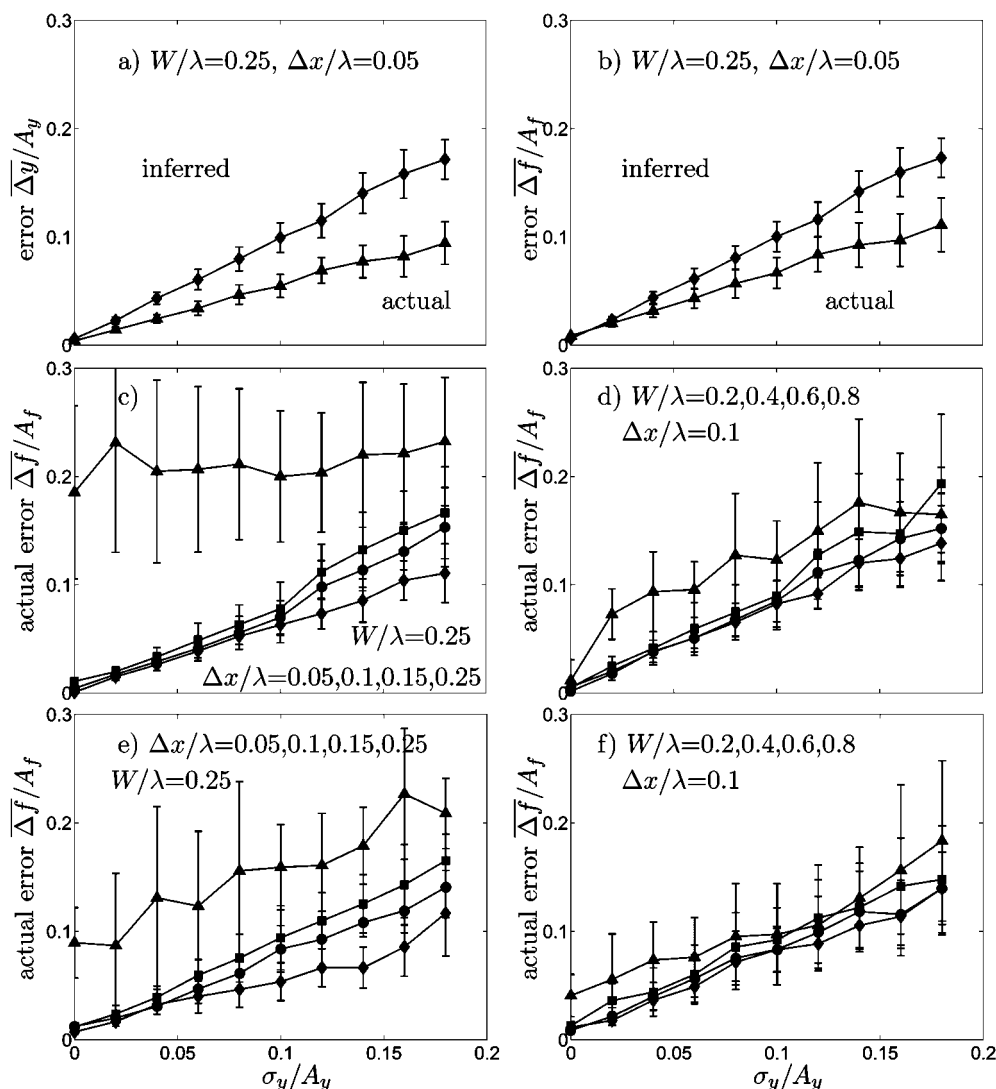


Figure 6 The difference between the predictions and the true error values, where the underlying function is $f(x) = \sin(2\pi x/\lambda)$. Each point is obtained by averaging over 40 simulations (generating synthetic data and training the model 40 times). In (a) and (b), the triangles with error-bars are the actual errors of the convolution (a) and the underlying function (b), and the diamonds with error-bars represent the uncertainty inferred by the model using equation (19). In (c), (e), (d) and (f), diamonds, circles, squares and triangles are used respectively as the varied parameter increases in size. In (c) and (d), the measurement error (σ_y/A_y) is inferred from the data, whereas in (e) and (f) it is input as prior knowledge.

slowly as a function of window width. The actual errors are small, and the deconvolution method can retrieve most of the underlying function. As the window width increases beyond λ , the error in the deconvolution increases very quickly and the solution soon becomes unacceptable. At $W \simeq 2\lambda$, the average error of the deconvolution reaches a maximum value, $1/(2^{1/2})$, and the algorithm gives only the trivial solution (straight line of zero height) in accordance with equation (27) and Fig. 5. If $W > 2\lambda$, the performance improves somewhat, although not to an acceptable level.

With regard to the measurement error σ_y , it is clear that in the regime for which window width and sampling frequency provide acceptable performance, the plateau level is strongly influenced by the measurement error.

If one has no prior expectation of likely dominant wavelengths of the underlying function, one could make an assessment of the critical length scale based on a structural

integrity approach, and use Fig. 7 to assess the largest admissible gauge width. Alternatively one could undertake measurements at W and $2W$, and by comparing the increase in amplitude of the (convoluted) measurements with Fig. 5, it would be possible to read off the dominant stress scale (wavelength).

4.3.4. Optimal scanning time. If one has some expectation of the likely wavelength of the underlying function, using the arguments above one can optimize the scan time by maximizing the quantity $\Delta x W^2 \sigma_y^2$ [equation (22)] consistent with achieving an acceptable average error in the estimation of the underlying function $\overline{\Delta f}$. From our simulations, if an average error of $0.1A_f$ is acceptable, then the fastest scan time per unit length of profile is given by $\Delta x = 0.2\lambda$, $W = 1.08\lambda$, $\sigma_y = 0.1A_y$. This optimization reflects the tendencies identified above, namely to keep Δx below 0.22λ , $W \simeq \lambda$ or less, and $\overline{\Delta y}/A_y \simeq \overline{\Delta f}/A_f$.

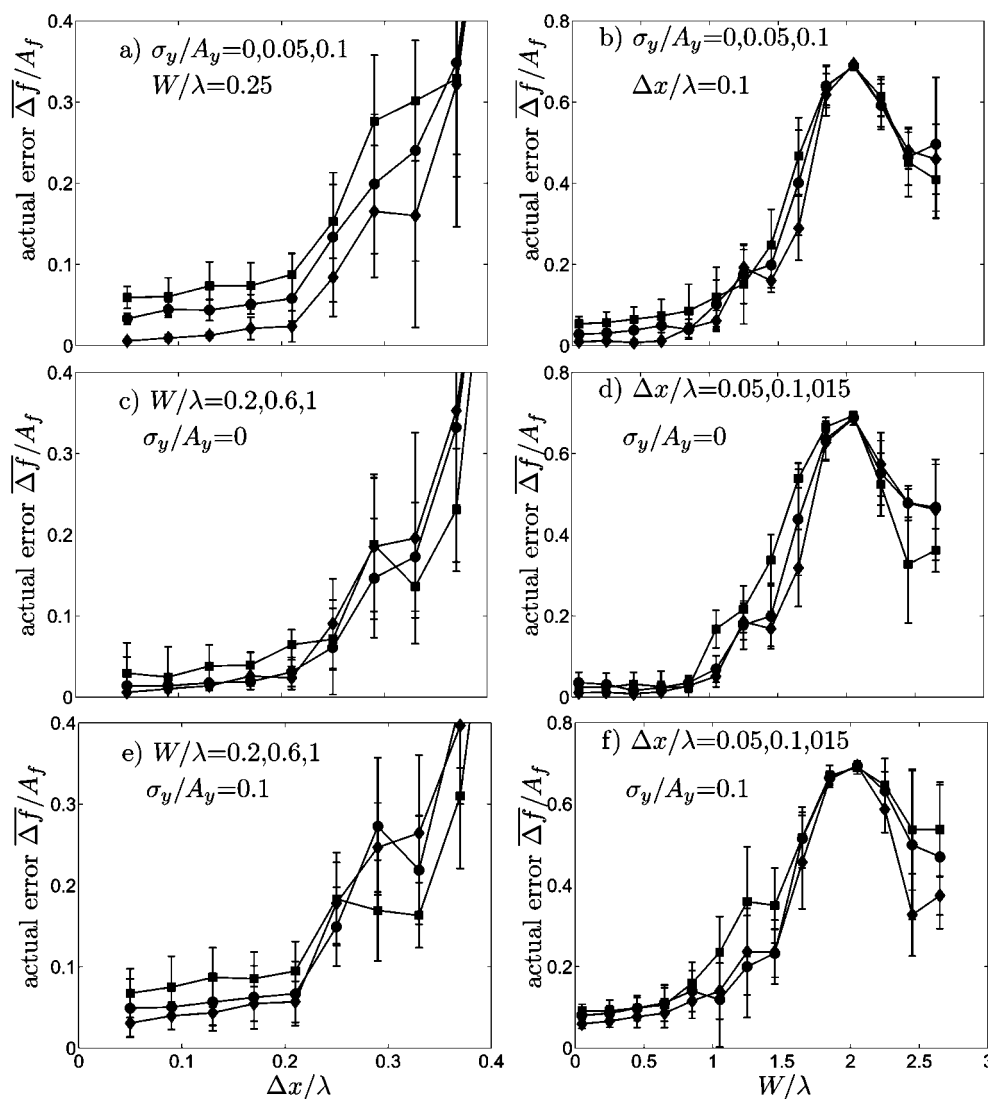


Figure 7

The actual errors, where the underlying function, $f(x) = \sin(2\pi x/\lambda)$, is the same as used in Fig. 6. Similarly, each point is obtained in the same manner as in Fig. 6. Diamonds, circles and squares are used respectively as the varied parameter increases in size.

5. Application to the reconstruction of real diffraction data

While the methods by which strain is measured by neutrons and synchrotron X-rays are essentially the same, the different characteristics of the beams as discussed above (the latter having much larger fluxes and much lower scattering angles) means that the forms of data collected are usually quite different. This is exemplified by the neutron and synchrotron strain data obtained for a 10.6 mm shot-peened 7071 Al block (Webster *et al.*, 1997). The synchrotron data were collected on BM16 at the ESRF, Grenoble, with the scattering vector (also the strain measurement direction) parallel and perpendicular to the surface to measure the in-plane and out-of-plane components of strain, respectively. The neutron diffraction data were collected on D1A at the ILL only for the in-plane direction. In all cases, the 311 reflection was used. In the former case, slits of 100 μm width were used with X-rays of 40 keV at a scattering angle (2θ) of 15.06° , giving a diamond gauge of approximate dimensions 760 and 100 μm across the diagonals. As a result, the spatial resolution for the out-of-plane strain was 7.6 times smaller than for the in-plane strain. An analyser crystal was positioned in the diffracted beam, which had the benefit of minimizing surface effects (Withers, 2003). A slit spacing of 0.5 mm was used at the ILL steady-state neutron source, giving a diamond gauge volume having approximately equal diagonals (710 μm) at a scattering angle of 110° . As a result of the much faster times associated with collection of a diffraction peak using the synchrotron arrangement, the sampling interval is characteristically finer and the measurement uncertainty, σ_y , smaller.

The measured strain profiles for the synchrotron and neutron data are shown in Fig. 8(a). The strain measurement errors representative of the synchrotron are of the same size as the symbols, except at the surface, while the mean error for the neutron measurements is $\sim 2 \times 10^{-4}$. In practice, it is common for shifts in diffraction angle that are not related to elastic strain to occur when the instrumental gauge is partially full. These effects are sometimes called surface strains or spurious strains and are well documented (Hutchings *et al.*, 2005; Edwards, 2003; Webster *et al.*, 1996; Spooner & Wang, 1997). In this work, the strains were measured either in such a way that the spurious strains are small [for example using an analyser crystal on BM16 (Withers, 2003)], or corrected by mathematical analysis (Hutchings *et al.*, 2005). With attenuation lengths of 6.5 and 100 mm for 40 keV X-rays and thermal neutrons, respectively, in both cases the sampling gauge can be considered to be evenly illuminated.

The maximum spatial resolution is achieved when the instrumental gauge just penetrates the surface (when its centre is located a half of the gauge diagonal from the surface). This occurs at $x = -50, -380$ and $-355 \mu\text{m}$ for the out-of-plane synchrotron, in-plane synchrotron and in-plane neutron measurements, respectively. As the sample enters the gauge volume, the effective gauge shape and the position of its effective centre change until the gauge is entirely within the sample, after which the effective and geometrical centres

coincide (neglecting the effect of attenuation). The divisor in equation (2) takes care of the effective increase in the sensor area as the instrumental gauge is progressively filled by the sample. Naturally, the diffracted signal increases over this regime.

In each of the three cases, we take the sensor to be a triangular function having the extents listed above. We can then apply the deconvolution based on our statistical learning method to this data set. Firstly, we can obtain best estimates of the convoluted sensor position-measured strain response free from sensor measurement errors. These are indicated by the lines in Fig. 8(a). The best estimates of the underlying residual strain profile, reconstructed taking account of the convoluting effect of the sensor responses, are plotted in Fig. 8(b).

Since the underlying macro-strain profile is not known, the performance of our deconvolution method cannot be assessed directly. However, as a result of shot peening, the out-of-plane stress is expected to be zero and the in-plane stress isotropic, and therefore we can relate both the in-plane (ε_z) and the out-of-plane (ε_y) strains uniquely to the in-plane stress distribution [$\Sigma_z(x)$]. In simple terms:

$$\Sigma_z(x) = \Sigma_y(x) = E \varepsilon_z(x)/(1 - \nu) = -E \varepsilon_x(x)/2\nu. \quad (28)$$

Given that for aluminium the Young's modulus, E , is around 70 GPa, and Poisson's ratio, ν , is around 0.33, the in- and out-of-plane strain profiles should be essentially equal in magnitude, but opposite in sign. The stress profiles calculated directly from the strain profiles of Fig. 8(b) are shown in Fig. 8(c). It is clear that because of the finer spatial resolution, the out-of-plane synchrotron measurements are most likely to represent the 'true' profile accurately. These show the least difference between the measured and deconvolved profiles. The surface stress is in good agreement with a laboratory X-ray stress measurement made at the surface (-100 MPa). In contrast to the original analysis (Webster *et al.*, 1997), it has been possible to recover the underlying shape of the stress field using the in-plane synchrotron strain measurements. Once deconvolved, the underlying profile shows a high level of agreement and the location and extent of the maximum compressive stress as well as the near-surface stress are in line with the out-of-plane measurements. There is some evidence that the profile is shifted by $\sim 50 \mu\text{m}$ relative to that for the out-of-plane measurements; this degree of mis-registration would not have been exceptional given the sample location procedures available at the time. The neutron measurements show the largest discrepancy (as well as the largest predicted uncertainties), with the sub-surface compressive maxima not recovered upon deconvolution. This is a consequence of the larger measurement uncertainty in the data and the fact that very few points were measured near the surface.

6. Conclusions

In this paper, we have presented a robust deconvolution method based on a Gaussian statistical learning technique. The measurement of convolution is fully described by a probabilistic model that includes a Bayesian algorithm. This

method allows predictions of both the convolved profile and the underlying function, together with their error-bars, while minimizing the effect of measurement scatter or fluctuation. The deconvolution method has been tested on both synthetic and real data sets. The results have demonstrated the conditions under which the model works well in retrieving the underlying function. The conditions under which the data are too noisy or too sparse, or where the sensor has inadequate spatial resolution, have also been determined. The *MATLAB* code for our statistical deconvolution algorithm together with the mathematical formalism is available from our website (<http://www.sdc.manchester.ac.uk/soft/>) for those who wish to apply it to their own data, or for evaluation.

Some conclusions may be drawn as follows.

(i) The actual error (the difference between the predicted and the true value of the underlying function) is proportional to the measurement uncertainty σ_y . The effect of the noise or microstructural fluctuations σ_f in the underlying profile on the reconstruction is very small, and can be neglected.

(ii) The sampling interval does not affect the performance very much, provided that it is less than 20% of the dominant wavelength. However, the sampling rate should be increased accordingly in order to compensate for loss of information due to measurement error and the limited spatial resolution of the sensor window.

(iii) In theory, for the idealized case of Δx and σ_y tending to zero, the window width can be set at any value except a multiple of 2λ in order to reconstruct the function. However, in practice, the window width needs to be of the order of λ or smaller, in order to obtain an acceptable deconvolution (*i.e.* to recover the underlying function within 90%). Wavelengths shorter than $W/2$ can be regarded as irrecoverable practically.

(iv) Our method performs best for overlapping measurements; in cases where they do not overlap it may be necessary to input the measurement uncertainty directly into the model as prior knowledge.

Of course conclusions (ii) and (iii) are made on the basis of simulations carried out for idealized profiles of constant wavelength. For a real function with a continuous spectrum, the situation will be more complex. It should be noted that provided that the acquisition time is increased to maintain a constant level of measurement uncertainty σ_y , the sampling gauge can resolve short-wavelength components of the profile near the surface, even for large sampling gauges.

The implications of these results have been explored for the optimal choice of gauge size and measurement spacing needed to recover the underlying macro-stress or strain profile within minimum time. Perhaps contrary to current practice, the underlying function can be most efficiently recovered using our algorithm for neutron measurements using relatively large gauges (approximately equal to the expected wavelength), but with a high degree of overlap between successive measurements. Unless one is prepared to use long acquisition times to achieve high levels of point-to-point measurement accuracy, this is a more effective strategy than increasing the spatial resolution by decreasing the window size. If it is acceptable to recover the profile to within $0.1A_f$, then a measurement

uncertainty of $0.1A_f$ is probably sufficient for each measurement. Finally, it should be noted that while we have applied the deconvolution algorithm to the recovery of the underlying

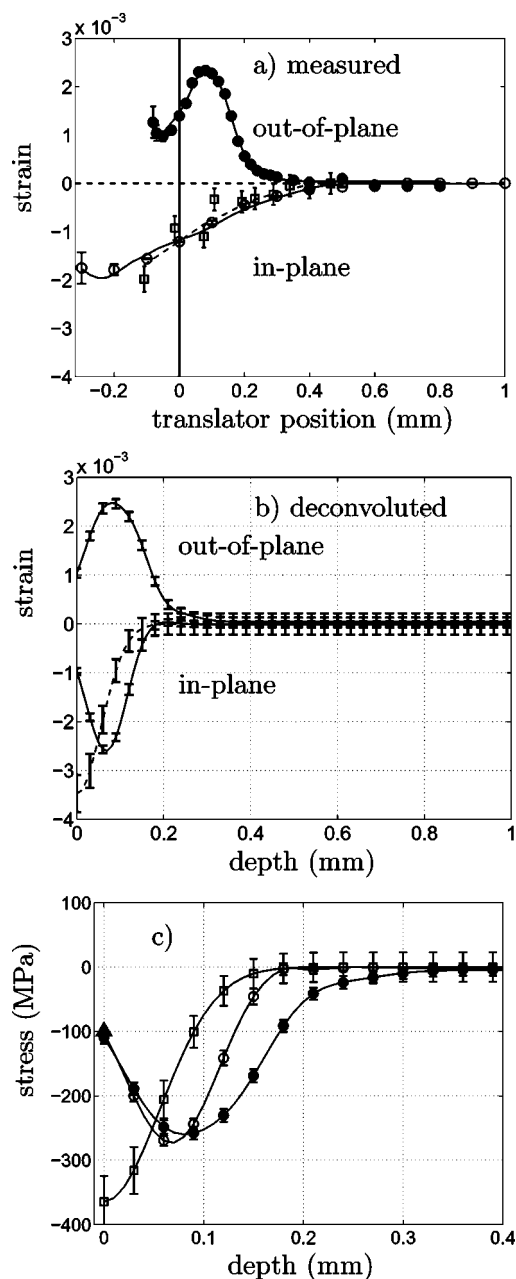


Figure 8

The performance of our algorithm on the experimental data (Webster *et al.*, 1997) collected near a shot-peened Al surface using $W = 0.10, 0.76$ and 0.71 mm and near-surface sampling intervals Δx of 0.02, 0.1 and 0.1 mm for the out-of-plane synchrotron (filled circles), in-plane synchrotron (open circles) and in-plane neutron (open squares) measurements, respectively. The full lines are profiles reconstructed using our method for synchrotron measurements, while the dashed lines are for the in-plane neutron data. (a) The measured strain profile. The vertical line at $x = 0$ denotes the edge of the sample. (b) The reconstructed residual strain profile versus depth below the peened surface. (c) The inferred residual stress profile versus depth below the peened surface; a laboratory X-ray $\sin^2 \psi$ surface-stress measurement (triangle) is also included for comparison.

profile for diffraction strain data, the approach is completely general and is expected to be of wide utility.

PJW is grateful for the support of the Royal Society Wolfson Merit Award. Y-SX is grateful for EPSRC Platform Grant funding. Helpful discussions with M. W. Johnson and the comments of the referees are also acknowledged.

References

- Bishop, C. M. (1995). *Neural Networks for Pattern Recognition*. Oxford University Press.
- Carpinteri, A., Spagnoli, A. & Vantadori, S. (2002). *Fatigue Fract. Eng. Mater. Struct.* **25**, 619–627.
- Clausen, B., Lorentzen, T. & Leffers, T. (1998). *Acta Mater.* **46**, 3087–3098.
- Edwards, L. (2003). *Analysis of Residual Stress by Diffraction using Neutron and Synchrotron Radiation*, edited by M. E. Fitzpatrick & A. Lodini, pp. 233–248. London: Taylor and Francis.
- Fitzpatrick, M. E. & Lodini, A. (2003). Editors. *Analysis of Residual Stress by Diffraction using Neutron and Synchrotron Radiation*. London: Taylor and Francis.
- Hutchings, M. T., Withers, P. J., Holden, T. M. & Lorentzen, T. (2005). *Introduction to the Characterization of Residual Stresses by Neutron Diffraction*. London: CRC Press.
- MacKay, D. J. C. (1992). *Neural Comput.* **4**, 720–736.
- Noyan, I. C. & Cohen, J. B. (1985). *Mater. Sci. Eng.* **75**, 179–193.
- Noyan, I. C. & Cohen, J. B. (1987). *Residual Stress – Measurement by Diffraction and Interpretation*. New York: MRE Springer Verlag.
- Oliver, E. C., Daymond, M. R. & Withers, P. J. (2004). *Acta Mater.* **52**, 1937–1951.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. (1992). *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed. Cambridge University Press.
- Spooner, S. & Wang, X. L. (1997). *J. Appl. Cryst.* **30**, 449–455.
- Tipping, M. E. (2001). *J. Machine Learning Res.* **1**, 211–244.
- Ungar, T., Mughrabi, H., Ronnpagel, D. & Wilkens, M. (1984). *Acta Metall.* **32**, 333–342.
- Van Acker, K., Root, J., Van Houtte, P. & Aernoudt, E. (1996). *Acta Mater.* **44**, 4039–4049.
- Vapnik, V. (1995). *The Nature of Statistical Learning Theory*. New York: Springer.
- Wang, X.-L., Wang, Y. D., Stoica, A. D., Horton, D. J., Tian, H., Liaw, P. K., Choo, H., Richardson, J. W. & Maxey, E. (2005). *Mater. Sci. Eng. A*, **399**, 114–119.
- Webster, P. J., Mills, G., Wang, X.-D. & Kang, W.-P. (1997). *5th International Conference on Residual Stress*, edited by T. Ericsson, M. Odén & A. Andersson, pp. 551–558. Institute of Technology, Linköping University, Sweden.
- Webster, P. J., Mills, G., Wang, X.-D., Kang, W.-P. & Holden, T. M. (1996). *J. Neutron Res.* **3**, 223–240.
- Withers, P. J., Daymond, M. R. & Johnson, M. W. (2001). *J. Appl. Cryst.* **34**, 737–743.
- Withers, P. J. (2003). *Analysis of Residual Stress by Diffraction using Neutron and Synchrotron Radiation*, edited by M. E. Fitzpatrick & A. Lodini, pp. 170–189. London: Taylor and Francis.
- Withers, P. J. (2004). *J. Appl. Cryst.* **37**, 596–606.
- Xiong, Y.-S., Buttler, D. J. & Withers, P. J. (2006). In preparation.
- Xiong, Y.-S. & Withers, P. J. (2005). *J. Mater. Process. Technol.* **170**, 551–562.