Adaptive sampling for accelerating neutron diffraction-based strain mapping[*]

To cite this article: S V Venkatakrishnan *et al* 2023 *Mach. Learn.: Sci. Technol.* **4** 025001

View the article online for updates and enhancements.

MACHINE LEARNING
Science and Technology

PAPER

OPEN ACCESS

Adaptive sampling for accelerating neutron diffraction-based strain mapping*

S V Venkatakrishnan[1,*] , Chris M Fancher[2] , Maxim Ziatdinov[3] , Rama Vasudevan[4] , Kyle Saleeby[5] , James Haley[6] , Dunji Yu[7] , Ke An[7] and Alex Plotkowski[8]

1 Multi-modal Sensor Analytics Group, Oak Ridge National Lab (ORNL), Oak Ridge, TN 37831, United States of America
2 Materials for Advanced Manufacturing Group, ORNL, Oak Ridge, TN 37831, United States of America
3 Computational Chemistry and Nano-materials Group, ORNL, Oak Ridge, TN 37831, United States of America
4 Data Nano-analytics Group, ORNL, Oak Ridge, TN 37831, United States of America
5 Georgia Tech Manufacturing Institute, Georgia Tech, Atlanta, GA 30332, United States of America
6 Energy Systems Analytics Group, ORNL, Oak Ridge, TN 37831, United States of America
7 Neutron Scattering Division, ORNL, Oak Ridge, TN, 37831, United States of America
8 Microstructural Evolution Modeling Group, ORNL, Oak Ridge, TN 37831, United States of America
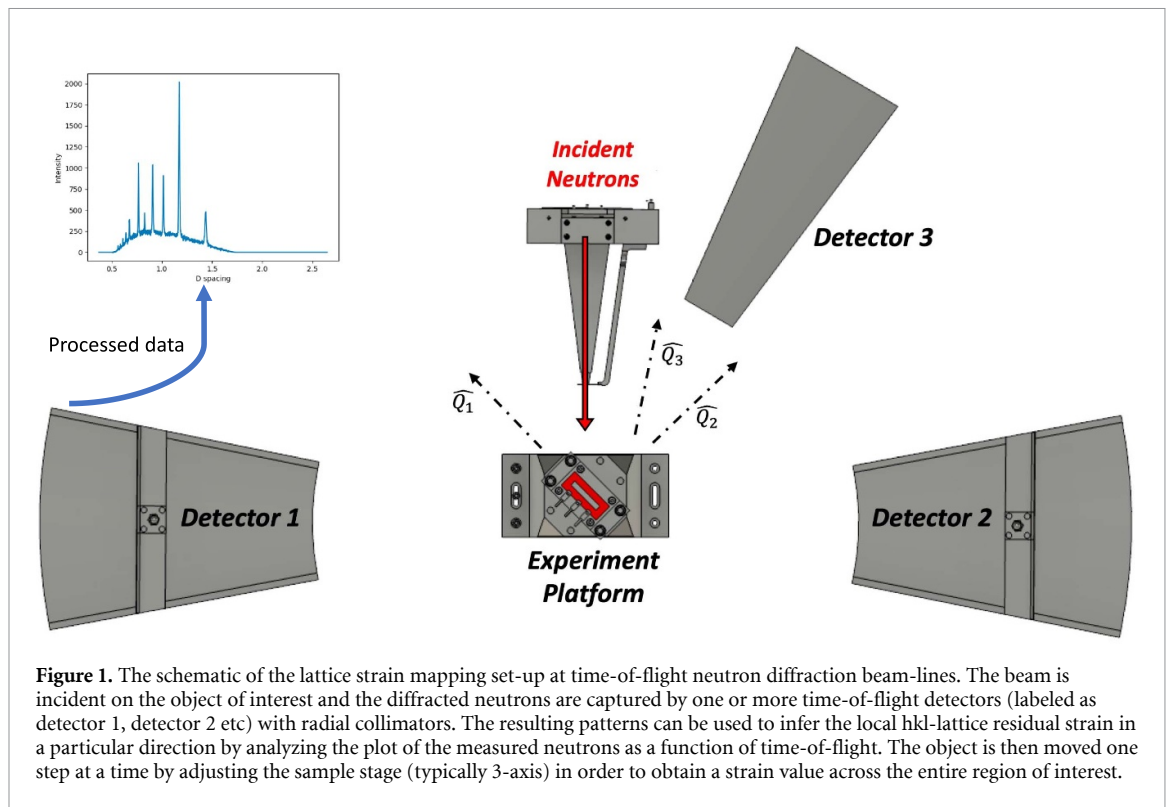* Author to whom any correspondence should be addressed.

E-mail: venkatakrisv@ornl.gov

Abstract

Neutron diffraction is a useful technique for mapping residual strains in dense metal objects. The technique works by placing an object in the path of a neutron beam, measuring the diffracted signals and inferring the local lattice strain values from the measurement. In order to map the strains across the entire object, the object is stepped one position at a time in the path of the neutron beam, typically in raster order, and at each position a strain value is estimated. Typical dwell times at neutron diffraction instruments result in an overall measurement that can take several hours to map an object that is several tens of centimeters in each dimension at a resolution of a few millimeters, during which the end users do not have an estimate of the global strain features and are at risk of incomplete information in case of instruments outages. In this paper, we propose an object adaptive sampling strategy to measure the significant points first. We start with a small initial uniform set of measurement points across the object to be mapped, compute the strain in those positions and use a machine learning technique to predict the next position to measure in the object. Specifically, we use a Bayesian optimization based on a Gaussian process regression method to infer the underlying strain field from a sparse set of measurements and predict the next most informative positions to measure based on estimates of the mean and variance in the strain fields estimated from the previously measured points. We demonstrate our real-time *measure-infer-predict* workflow on additively manufactured steel parts—demonstrating that we can get an accurate strain estimate even with 30%–40% of the typical number of measurements—leading the path to faster strain mapping with useful real-time feedback. We emphasize that the proposed method is general and can be used for fast mapping of other material properties such as phase fractions from time-consuming point-wise neutron measurements.

Figure 1. The schematic of the lattice strain mapping set-up at time-of-flight neutron diffraction beam-lines. The beam is incident on the object of interest and the diffracted neutrons are captured by one or more time-of-flight detectors (labeled as detector 1, detector 2 etc) with radial collimators. The resulting patterns can be used to infer the local hkl-lattice residual strain in a particular direction by analyzing the plot of the measured neutrons as a function of time-of-flight. The object is then moved one step at a time by adjusting the sample stage (typically 3-axis) in order to obtain a strain value across the entire region of interest.

1. Introduction

Neutron diffraction, and in particular engineering diffraction, is a useful tool for measuring bulk residual strains and stresses [1, 2] in metal objects. These data are typically obtained by placing the object of interest in the path of a broad-band neutron beam and using one or more detectors placed at a certain angle with respect to the incident beam (see figure 1) to measure an intensity vs. scattering angle/time-of-flight/d-spacing histogram at a neutron reactor or spallation source [3], repetitively. Engineering diffractometers use neutron optics to capture the signal only from a small probe volume of the part (called gauge volume) which is defined by the incident slits and receiving collimation. Measured diffraction data may be used to extract a variety of information describing the material state, such as phase, lattice parameter, preferred crystallographic orientation, elastic strain, etc. Here, we only utilize the shift in diffraction peak position to compute the residual lattice strain. Lattice strains corresponding to a gauge volume around a location (x, y, z) in the sample are computed by evaluating the shift of a certain peak with respect to a known reference using the following relationship:

$$\varepsilon(x, y, z) = \frac{d_{hkl}(x, y, z) - d_{0, hkl}}{d_{0, hkl}} \qquad (1)$$

where d_{hkl} and $d_{0, hkl}$ are the inter-atomic spacing for the object of interest and stress-free reference corresponding to a particular hkl crystallographic plane respectively. This method can be extended to multiple detector banks to compute strain in multiple strain directions. Neutron diffraction-based strain mapping (NDSM) has been used for several applications such as evaluating strains in additively manufactured (AM) parts [4–8].

NDSM is a scanning probe measurement in which the object is mapped one point at a time in 2D or 3D. A typical workflow involves defining a discrete grid of points and raster scanning by moving the object stage once each point has been measured. Upon measuring all the points, an algorithm is used to fit the peaks and compute the strain value at each location in the object. However, there are several drawbacks for such typical workflows. First, neutron measurements are time consuming (taking of the order of 30 s or up to a few minutes per point when the gauge volume being investigate is of size $5 \times 5 \times 5$ mm^3) depending on other factors such as penetration depth and type of material which influences the signal-to-noise ratio. Thus for a sample of size $5 \times 5 \times 5$ cm^3 corresponding to 1000 points it can take nearly 10 h of scan time in order to map the strain. In the course of this time-consuming measurement, the user of such instruments typically does not receive real-time feedback of the strains measured thus far. Second, even if an estimate of the strain at each position can be computed rapidly, because of the raster scanning strategy the user only gets a progressive

view of the sample one region at a time, taking hours before identifying regions of high strain gradient (if any) and therefore are not able to adapt their measurements, or stop the scan if there are no interesting features in the object. Thirdly, such a raster scan strategy does not adapt to the object under study—thereby leading to potentially uninformative measurements in regions that have largely uniform strain. In summary, while NDSM is a powerful characterization tool, there are several drawbacks with the current workflows.

There has been interest in recent years to be able to reduce acquisition time for slow scanning probe systems (where measurement time is much longer than the time to re-position the probe/sample) while still being able to infer the underlying information of interest. One common strategy has involved making a uniformly random set of measurements and using an advanced image reconstruction algorithm to infer the underlying property [9–11] at all points. However, this approach does not account for the geometry of the quantity of interest (such as residual strain) being measured and can miss important details in the specimen under study. In order to address this drawback, researchers have developed dynamic sampling algorithms based on supervised machine learning (ML) [12] where the methods start with a small uniformly random set of measurements, interpolate these onto a finer grid, and decide on the next set of measurements based on some metric. These algorithms have also been implemented in real-world systems such as scanning electron microscopes and scanning x-ray systems [13, 14] illustrating how the total number of measurements can be reduced to about 30%–40% of the typically made measurements with a minimal loss of information thereby enabling a new era of autonomous instruments [15]. Alternate approaches, such as Bayesian Optimization (BO) based Gaussian process (GP) regression [16] have also been used since these do not require any training data unlike the supervised learning approaches. Recent demonstrations for small-angle x-ray scattering [17], scanning probe microscopy [18] and neutron triple axis spectrometry [19] further highlight the potential of such methods.

In this paper, we present an adaptive sampling approach for neutron diffraction based single-axis strain mapping. Specifically, we use a BO-based GP method to reconstruct the underlying strain fields along with an uncertainty map from a sparse set of measurements, and use the estimated quantities to predict the next best measurement point. This process of *measure, infer and predict* is run in a loop till we obtain an accurate reconstruction of the underlying strain field. We modified the beam-line instrumentation so that the data is measured, streamed, reduced and then fed to the GP model running on a powerful compute node in real time so as to realize a practical system. We demonstrated the utility of our approach by using it to measure the residual lattice strain in an AM low transformation temperature (LTT) steel. Using experimental data collected at the VULCAN [3] beam-line at the Spallation Neutron Source at Oak Ridge National Laboratory we demonstrate that as opposed to conventional raster scan approaches we can reduce number of measurements to about 30%–40% of what is typical for a representative LTT steel sample and yet obtain an accurate representation of the strain map. Thus our approach can be used to dramatically reduce scan time, provide useful real-time feedback and improve the operations of neutron-diffraction instruments when used for strain mapping.
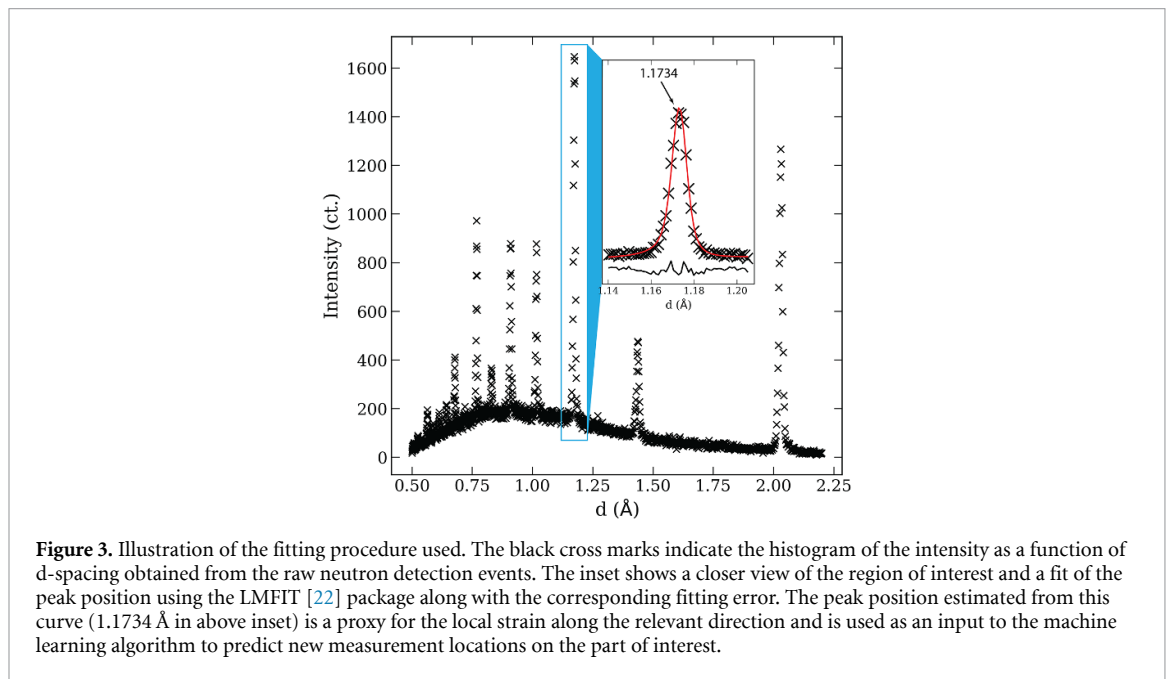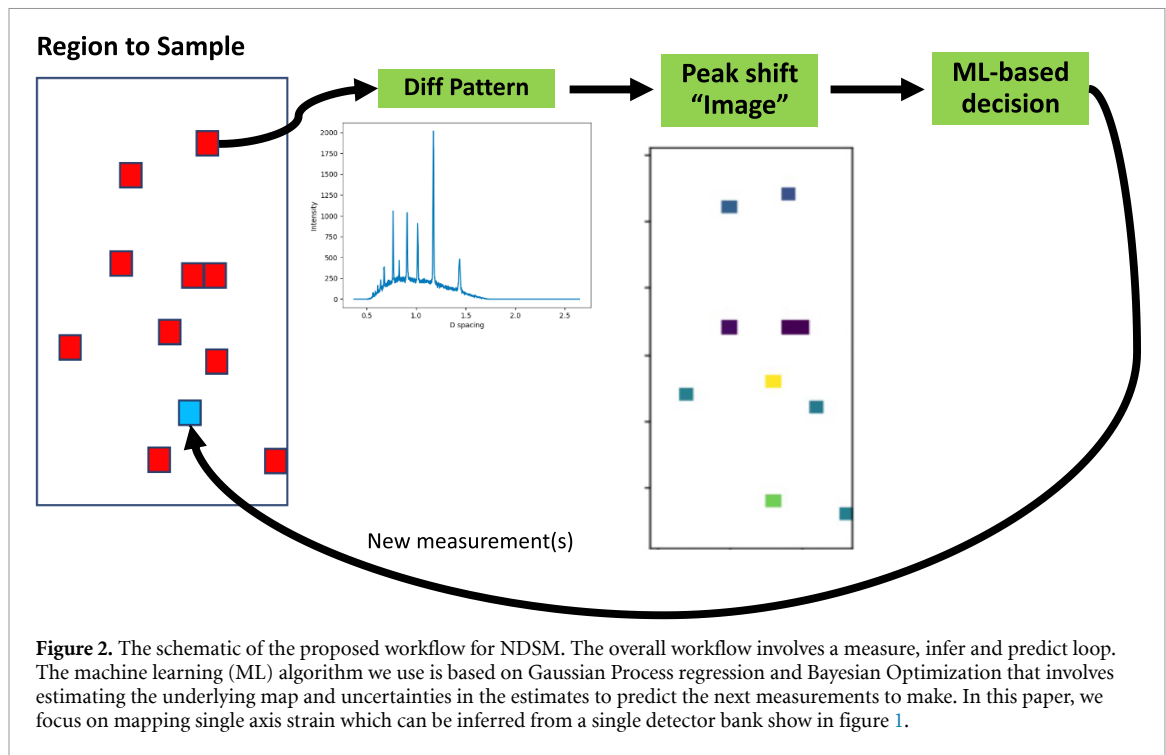
2. Proposed workflow for adaptive sampling

Figure 2 shows an illustration of our proposed approach for fast NDSM. Measured time-focused neutron data (time-focused from neutron events) are continuously updated within the Experimental Physics and Industrial Control System (EPICS) and then after each planned measurement time or total neutron proton charge is met, the histogram is immediately analyzed to extract the position of a particular peak which is the d_{hkl} in equation (1). These are then used as input to an ML model that predicts the locations of next set of measurements. The loop can be terminated based on an automated stopping criteria or a fixed measurement/time budget.

2.1. Real-time reduction, streaming, peak fitting and strain computation
Neutron diffraction data were measured at the VULCAN [3] engineering neutron diffractometer at the Spallation Neutron Source (SNS) at the Oak Ridge National Laboratory. The SNS is a pulsed neutron source that generates neutrons using a spallation process by bombarding a mercury target with high energy proton pulse. Bunches of protons are pulsed at the mercury target at 60 Hz, which produces neutrons with a distribution of energies. Neutron events [20] measured at VULCAN are recorded with specific pixel ids and arrival time. These event data collected in the 2-d detectors are processed to generate representative intensity vs. time-of-flight (or d-spacing) histogram (see figure 3).

The EPICS on VULCAN is used as the overarching system for motor control with the neutron event distributor (nED) data acquisition [21]. These systems work in tandem to facilitate efficient experimental control for neutron diffraction experiments. nED has a modular software design to allow for instrument-specific live-processing of measured neutrons for quick investigation. EPICS V7 implements a

Figure 2. The schematic of the proposed workflow for NDSM. The overall workflow involves a measure, infer and predict loop. The machine learning (ML) algorithm we use is based on Gaussian Process regression and Bayesian Optimization that involves estimating the underlying map and uncertainties in the estimates to predict the next measurements to make. In this paper, we focus on mapping single axis strain which can be inferred from a single detector bank show in figure 1.

Figure 3. Illustration of the fitting procedure used. The black cross marks indicate the histogram of the intensity as a function of d-spacing obtained from the raw neutron detection events. The inset shows a closer view of the region of interest and a fit of the peak position using the LMFIT [22] package along with the corresponding fitting error. The peak position estimated from this curve (1.1734 Å in above inset) is a proxy for the local strain along the relevant direction and is used as an input to the machine learning algorithm to predict new measurement locations on the part of interest.

back-end system facilitating communications between nodes using a publish/subscribe process and allows external clients to subscribe and monitor the live data using the process variable access protocol.

We developed a custom python middle-layer library (MLL) to monitor and analyze the live neutron diffraction data and interact with the VULCAN instrument to start new measurements. The MLL interacted with the VULCAN EPICS system using pyEPICS (version 3.5.1). New measurements are started using a custom input-output control. We note that the live-processed data are linearly binned raw counts that are calibrated with a first order correction for the flight path. While this can be refined based on using calibration data in order to obtain a precise value for the strain, any shifts from a lack of a precise calibration are systematic to all data and are canceled by the formula in equation (1). The live data were analyzed via single peak fitting to extract the peak position (see figure 3 for a representative example). The LMFIT [22] python package was used for peak fitting using an asymmetric Pseudo-Voigt function with a linear background. Once this fitting was complete we obtain a strain value corresponding to a certain object position. Finally, we note a more accurate estimation of the strain values at each position can include a more

detailed use of calibration data and better fitting routines for the peak position and shape; however this is not required for our goal of adaptive sampling and automated experiment steering.

2.2. GP regression and BO

Gaussian process regression (GPR) [23] is a popular ML technique that can be used to infer with a quantified uncertainty the values of unknown points on a grid from a partial set of measurements. The core assumption is that the values of a function, $\mathbf{f}$ (like strain values across a object), over an underlying grid X is a jointly Gaussian random vector with a certain (prior) mean μ and co-variance K,

$$\mathbf{f} = f(X) \sim \mathrm{MVN}(\mu_X, K_{X,X}) \tag{2}$$

where MVN denotes the multivariate Normal/Gaussian and $K_{X,X}$ is a covariance matrix such that $(K_{X,X})_{ij} = k_\theta(x_i, x_j)$. The functional form of the kernel, k_θ, is chosen prior to the experiment and θ represents the parameters of the kernel. The mean μ is commonly set to a constant. From a practical point of view, the definition of GP implies that the underlying structure of the function is locally smooth, which is a reasonable assumption for the strain values at the resolutions typically used in neutron instruments. Because of the assumption of joint Gaussianty, the conditional probability of the function values at new point(s) given known values of the function at other points is also a Gaussian random vector whose conditional mean and co-variance can be computed in closed form,

$$\mathbf{f}_* | X_*, X, \mathbf{y}, \theta \sim \mathrm{MVN}(\mathbb{E}(\mathbf{f}_*), \Sigma(\mathbf{f}_*)) \tag{3}$$

$$\mathbb{E}(\mathbf{f}_*) = \mu_{X_*} + K_{X_*,X} K_{X,X}^{-1}(\mathbf{y} - \mu_X) \tag{4}$$

$$\Sigma(\mathbf{f}_*) = K_{X_*,X_*} - K_{X_*,X} K_{X,X}^{-1} K_{X,X_*} \tag{5}$$

where X and $\mathbf{y}$ are initial(typically sparse) grid points and corresponding (noisy) observations of the function, commonly referred to as training data, and X_* are new grid inputs (for which there no observations yet). Note that for the sake of brevity we assumed that the model noise is absorbed into the kernel computation. For more detailed overview of the GPR, we refer the reader to [23].

In order to use GPR as a technique to predict what points to measure next based on a partial set of measurements, one first computes an estimate of the mean and co-variance function at all points in the (discrete) search space of interest. Using these estimates, an acquisition function [16, 24] A is then computed over all the grid points in order to choose new points on the grid that are likely to provide the most informative measurement. Typically, this is done by

$$X_{\mathrm{new}} \leftarrow \underset{x \in X \setminus X_{\mathrm{previous}}}{\mathrm{argmax}} \; A(x; \mathbb{E}, \Sigma) \tag{6}$$

where X_{prev} is the set of all the previous locations on the grid at which a measurement has been made. This type of approach is called BO [25] because we effectively compute the posterior distribution parameters based on some underlying prior model to infer the next point(s) to measure. A simple acquisition function is one that simply uses the variance estimates ('uncertainty') at each position from equation (5) and chooses to measure the next point at the position corresponding to the highest variance. This is referred to as a pure exploration regime or 'active learning' whose goal is to reconstruct an entire data distribution with a minimal number of measurements. On the other hand, if the goal is also to find sample regions where a specific behavior is maximized (or minimized), more sophisticated acquisition functions such as the expected improvement or knowledge gradient [16] can result in a better exploration-exploitation trade-off for the overall system.

We have used GPR and BO in order to implement our adaptive closed loop system for strain mapping. The software implementation is based on the GPim library—https://github.com/ziatdinovmax/GPim—which provides an easy way to apply GPR and BO in Pyro [26] and GPytorch [27] to scientific imaging and hyperspectral data. In order to make a prediction according to equations (3)–(5), we learn GP model parameters θ from the available measurements at each step by performing a few iterations of gradient ascent with a certain step-size/learning rate on the log marginal likelihood. In this work we have used the Matern kernel [23] for the covariance kernel $k_\theta(x_i, x_j)$ corresponding to equation (2) that yielded satisfactory performance on previously acquired strain maps of similar materials as those used in our experiments. This kernel has the form

$$k(x_i, x_j) = \frac{1}{\Gamma(\nu) 2^{\nu-1}} \left(\frac{\sqrt{2\nu}}{l} d(x_i, x_j) \right)^\nu K_\nu \left(\frac{\sqrt{2\nu}}{l} d(x_i, x_j) \right) \tag{7}$$

where l is a length scale parameter (estimated from the data during the training process in our case), Γ is the gamma function, K is a modified Bessel function, $d(x_i, x_j)$ is the Euclidean distance between the points, and we set $\nu = 0.5$ which is a common choice in the literature especially when we wish to capture sharper variations in the underlying quantity. For the BO, we have chosen the expected improvement [24] function as our acquisition module that computes the values based on the estimated mean and variance at all points in the grid of interest. This expected improvement function—one of the most popular choices for BO—computes the expected improvement in the maximum value of the function when adding each new grid point. Once this value is computed for all the unmeasured points on the grid, we choose the next point based on the one that has the maximum value for the acquisition function. Finally, before applying the GPR and BO, we normalize all the measured strain values approximately to the $[0, 1]$ range by subtracting a reference value, and dividing the result by a fixed number based on past scans of similar materials. We note that alternate schemes can also be used to normalize the data without having a significant impact on algorithm performance. A flow-chart showing the overall approach is shown in algorithm 1.

Algorithm 1. Flowchart showing the adaptive sampling method over a uniform 2D grid.

1: **function** $[\hat{f}] \leftarrow$ ADAPTIVESAMPLING $(p_{init}, p_{final}, N_x, N_y, \theta_{init}, \Phi)$
2:     %**Inputs**: p_{init}—Initial fraction of total grid to measure; p_{final}—Final fraction of total grid at which to stop experiment; N_x, N_y: Total number of points along x and y directions; θ_{init}: Initial parameters for the covariance kernel; Φ: Parameters used for GP learning—step size and maximum number of iterations
3:     %**Outputs**: $\hat{f}$ - $N_x \times N_y$ sized array with estimated strain values
4:     $X_{visit} \leftarrow$ GenerateUniformPoints(N_x, N_y, p_{init}) ▷ Generate initial grid locations
5:     $X_{full} \leftarrow$ GenerateFullGrid(N_x, N_y) ▷ Generate all the grid points
6:     $Y_{visit} \leftarrow$ MEASURE(X_{visit}) ▷ Send information to the instrument, measure, reduce and obtain normalized peak shifts at the set of points as in section 2.1
7:     **while** $|X_{visit}| \leqslant \lceil p_{final} * N_x * N_y \rceil$ **do**
8:         $(\mu, \Sigma, \theta_*) \leftarrow$ EstimateGPparams$(Y_{visit}, X_{visit}, \theta_{init}, \Phi)$
9:         Compute $\mathbb{E}(\mathbf{f}_*)$ and $\Sigma(f_*)$ over X_{full} using equations (4) and (5)
10:         Compute $A(., \mathbb{E}(\mathbf{f}_*), \Sigma(f_*))$ over all points in X_{full} ▷ Using expected improvement acquisition function from [16]
11:         $X_{new} \leftarrow \text{argmax}_{x \in X_{full} \setminus X_{visit}} A(x; \mathbb{E}(\mathbf{f}_*), \Sigma(f_*))$
12:         $Y_{new} \leftarrow$ MEASURE(X_{new})
13:         $Y_{visit} \leftarrow Y_{visit} \cup Y_{new}$
14:         $X_{visit} \leftarrow X_{visit} \cup X_{new}$
15:         $\theta_{init} \leftarrow \theta_*$
16:     **end while**
17:     $\hat{f} \leftarrow \mathbb{E}(\mathbf{f}_*)$
18 **end function**

---

3. Results

We used two AM parts made of LTT steel—one actively cooled, and one passively cooled during printing in order to highlight the strength of our approach. Development of residual stresses during processing has been identified as a key scientific challenge for improving AM materials. The parts were built in two sections of 36.75 mm each, for a total height of 73.5 mm. Each section was allowed to fully cool to room temperature before additional material was deposited. This approach to processing creates a distinct pattern in the thermal cycles and phase transformations of the material. The volume expansion from austenite to martensite upon cooling causes a discernible feature in the residual lattice strain field below the interface between the two sections. This approach was taken here deliberately to generate a distribution of lattice strains to test the validity of the proposed sampling approach. Furthermore, the parts were thin in the 3rd dimension (see figure 4) and so we effectively only require 2D strain mapping. The chemistry of the LTT steel was selected specifically because it forms martensite even for slow cooling rates, and because the martensitic transformation occurs at depressed temperatures. The martensitic finish temperature was slightly below room temperature, resulting in a small fraction of retained austenite (less than 3 volume percent). Here, we assume the martensite strain state is representative of the material. Correspondingly, we used the (211) (hkl) body centered cubic reflection of the martensitic LTT sample as the reference peak ($d_{0,hkl}$ in equation (1)) at which we determine the shift of the sample. For one of the parts we measured the single-axis strain in a conventional raster scan mode and retro-actively apply our method to evaluate the performance of the proposed algorithm since we have the final ground-truth from the full sampling. For the other sample, we only deploy our adaptive sampling method collecting data up to nearly 76% of the full sampling grid. The incident slits were set as 5 mm in horizontal and 5 mm in vertical directions, and a set of 5 mm radial

Figure 4. Illustration of an object made of LTT steel used in our experiments with an overlaid grid. The right panels show the dimensions of the object—which is 10 mm in thickness and 73.5(height) $\times$ 65 mm in the face exposed to the neutron beam. The white boxes illustrate the footprint of the region being sampled (gauge volume) by the neutron beam. The overall object is mapped by stepping the stage one point at a time through the incident neutron beam. The successive steps can correspond to overlapped regions for a finer sampling grid. Depending on the size of the object and the beam-footprint, the total time to map the strains can be of the order of several hours. In our case, the gauge volume is about $5 \times 5 \times 5$ mm^3 and the object is stepped by 2.5 mm each time. Furthermore, we have used a thin sample (10 mm in depth dimension) in order to demonstrate the method on a 2D case.

collimators in front of the 90-degree banks were used to confine the horizontal field of view in the diffraction angle. As a result the defined beam gauge volume is $5 \times 5 \times 5$ mm^3 inside the 10 mm thick wall. The step size of the scan was set to 2.5 mm. The acquisition time per point was approximately 30 s per position to have enough statistics for a lattice strain determination with less than 100 micron strain error. Based on empirical tests of previously acquired strain maps, we choose a Matern kernel for the co-variance function. The learning rate for the GP parameters was set to 0.05 and the total number of training iterations for the GP was set to 200. The lower and upper bound value of the kernel length scales (l in equation (7)) were set to 1.0 and 4.0, respectively, and the actual value was jointly estimated during the GP training. (Pseudo) Strain values to be used were normalized so that

$$f(x,y) = \frac{d_{hkl}(x,y) - n_0}{n_1 - n_0} \tag{8}$$

where $f(x,y)$ is the value used by the GP algorithm, $d_{hkl}(x,y)$ is the peak shift position value estimated by the fitting algorithm discussed in section 2.1, n_0 and n_1 were the minimum and maximum peak-shift values obtained from a previous scan of a sample made of a same LTT steel corresponding to $n_0 = 1.168$ Å and $n_1 = 1.172$ Å used to normalize the data for use by the GP method. We emphasize that we have dropped the z dimension for this proof-of-concept study since the sample is thin along the direction of the beam (see figure 4).

First, we present results by retroactively sampling a full 2D strain map in order to compare the conventional raster scanning approach that is widely used at instruments to the proposed method. The grid was chosen to be of size 29×25 for a total of 725 points corresponding to a total measurement time (ignoring time to re-position the object) of approximately 6.04 h. Figure 5 shows the the initial sampling points (5% of total grid), estimated strain map and ground truth strain from the fully sampled grid. As expected the raster scan only can infer a small portion of the total object, while the GPR based method provides a low-resolution estimate of the full structure. Figure 6 shows the sampling grid, inferred reconstruction and ground truth from 40% of the grid points using the two schemes. We observe that the proposed method produces a strain estimate that is qualitatively very similarly to the ground truth data—highlighting the strength of our method. We also note that the sampled points automatically identify the points of interest in the material and hone in on such regions early in the scan as opposed to the conventional workflows where the method is agnostic to the strain distribution in the object. For neutron experiments this implies that even within a measurement time of nearly 2.4 h we are able to get a very good estimate of the strain map in the sample i.e. a dramatic ability to reduce scan time from the typical 6 h and get useful feedback. In order to quantify the performance of our approach, we also plot the normalized mean absolute deviation (NMAD) and the structural similarity (SSIM) [28] between the GP reconstruction after each new measurement to the fully sampled strain map in figure 7. The NMAD is computed as

$$\text{NMAD} = \frac{1}{C*M} \sum_{i=1}^{M} |\hat{f}(x_i, y_i) - f_{\text{GT}}(x_i, y_i)| \tag{9}$$
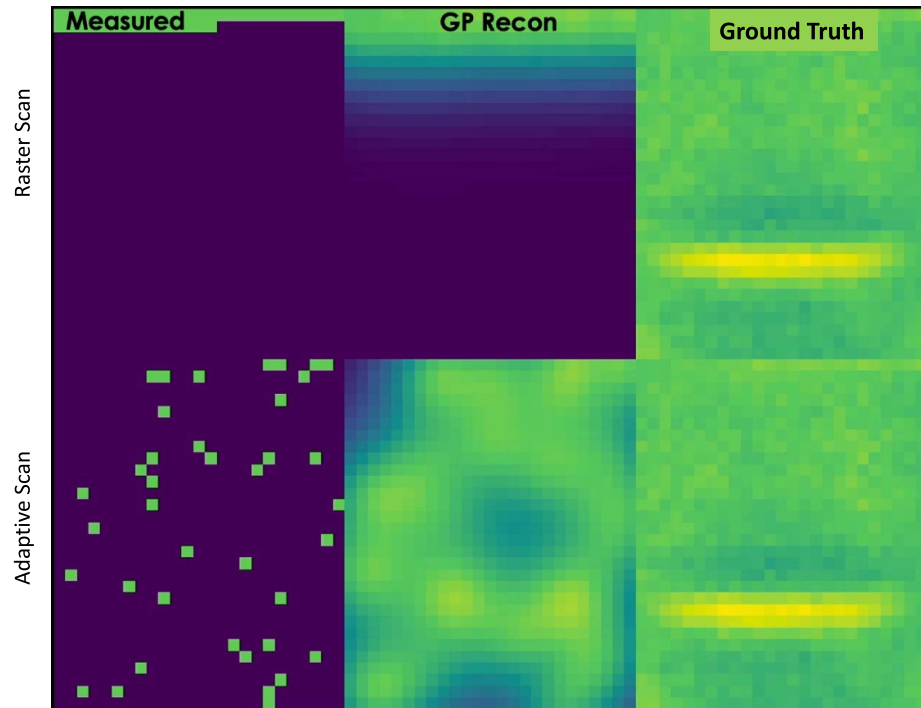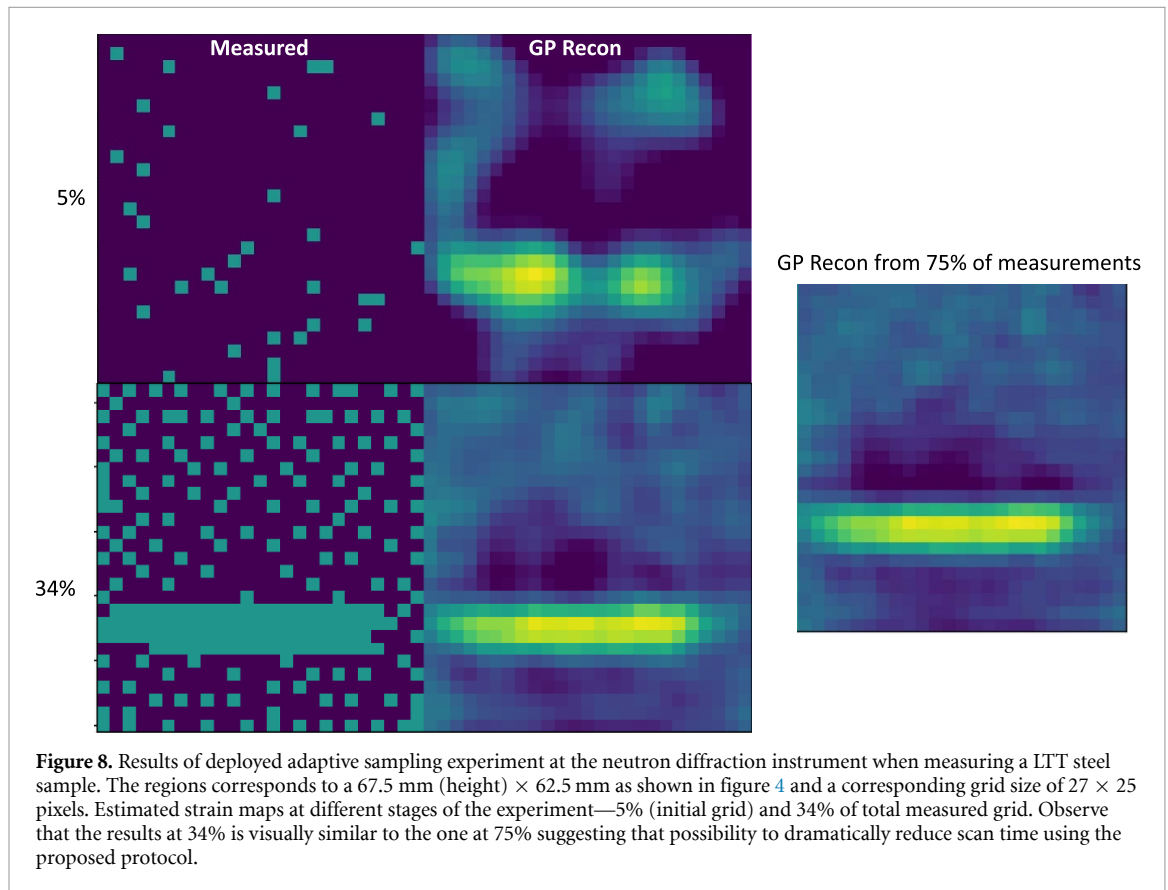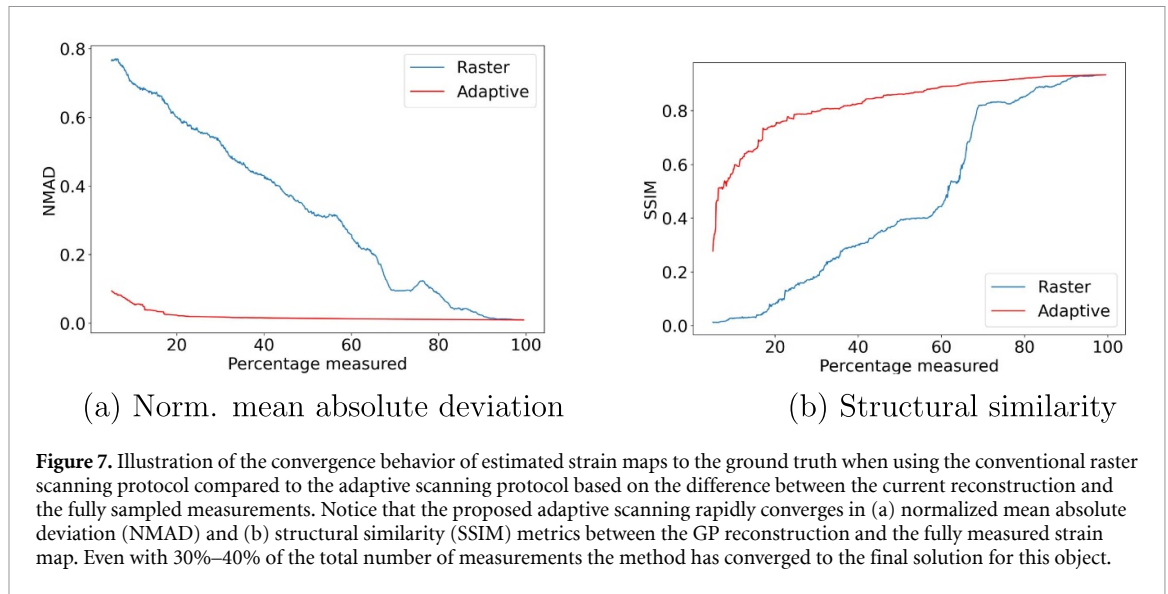
Figure 5. Results of retro-actively applying proposed GP-based method for strain mapping on a fully measured map (shown as ground truth) in order to evaluate the proposed adaptive scanning method compared to the traditional raster scan approach. The regions corresponds to a 72.5 mm (height) × 62.5 mm as shown in figure 4 and a corresponding grid size of 29 × 25 pixels. The first column in a binary image showing measured points in green and un-measured points in blue. The second column is the estimated strain map using GPR from the measured points. The third column is the fully sampled strain map i.e ground truth (GT). The images show the reconstructed strain maps when using 5% of the total grid points with the raster and proposed scanning method.
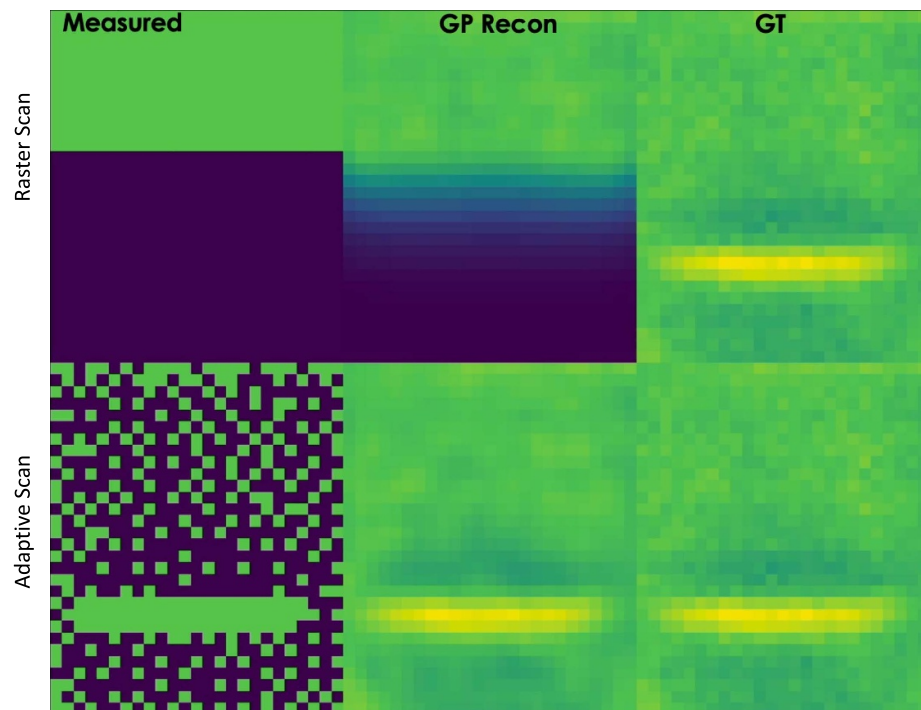
Figure 6. Estimated strain maps when using 40% of the total grid points using the conventional raster scanning approach and the proposed adaptive sampling strategy from the retro-actively sampled data. The regions corresponds to a 72.5 mm (height) × 62.5 mm as shown in figure 4 and a corresponding grid size of 29 × 25 pixels. Notice that the proposed GP method adapts the sampling points to the underlying strain map and hence achieves faster convergence to the ground truth.

(a) Norm. mean absolute deviation   (b) Structural similarity

Figure 7. Illustration of the convergence behavior of estimated strain maps to the ground truth when using the conventional raster scanning protocol compared to the adaptive scanning protocol based on the difference between the current reconstruction and the fully sampled measurements. Notice that the proposed adaptive scanning rapidly converges in (a) normalized mean absolute deviation (NMAD) and (b) structural similarity (SSIM) metrics between the GP reconstruction and the fully measured strain map. Even with 30%–40% of the total number of measurements the method has converged to the final solution for this object.

Figure 8. Results of deployed adaptive sampling experiment at the neutron diffraction instrument when measuring a LTT steel sample. The regions corresponds to a 67.5 mm (height) × 62.5 mm as shown in figure 4 and a corresponding grid size of 27 × 25 pixels. Estimated strain maps at different stages of the experiment—5% (initial grid) and 34% of total measured grid. Observe that the results at 34% is visually similar to the one at 75% suggesting that possibility to dramatically reduce scan time using the proposed protocol.

where $\hat{f}$ is the reconstructed strain map, f_{GT} is the ground truth value from a full sampling, M is the total number of grid points and C is chosen as the maximum value of the ground-truth strain map. We observe that the proposed method rapidly converges to the true solution in terms of NMAD and SSIM even with a fraction of the measurements (30%–40%) further demonstrating the utility of the proposed method.

Next, having studied the potential of the GP+BO method, we deploy the algorithm on a second LTT steel sample. The sampling grid size was of size 27 × 25 and the measurement time per point was approximately 30 s. Each grid point corresponded to a gauge volume of 5 × 5 mm in 2D with an overlap between each point being 2.5 mm. The diffraction data was measured, reduced, streamed to a powerful computing resource and rapidly used to compute the strain map at the measured points which was then fed to an automated loop control software. Due to experimental and time constraints we terminated the measurements when approximately 75% of the total number of grid points were measured by our adaptive sampling system. Figure 8 shows the results (GP reconstruction) from the adaptive sampling approach at 5%, 34% and 75% of

the total number of measurements. Similarly to the previous sample, we notice that the proposed algorithm locates and samples the regions of interest and produces a high quality inference of the underlying strain field even with a fraction of the total number of measurements—paving the way for a fundamentally different approach to making strain measurements at neutron diffraction instruments.

4. Conclusion

In this paper, we presented a method for accelerating neutron diffraction based strain mapping by proposing a fully closed-loop instrument that makes measurements, processes them on-the-fly, and decides the next informative measurement by using a BO method based on GPR. Our fast strain mapping system was enabled by innovations in hardware that enables accurately and rapidly re-positioning an object, fast neutron diffraction detectors, software systems that can read and reduce the data from the detectors, algorithms that can rapidly fit peaks and ML methods that can be designed to infer the next informative measurements to be made. We implemented our system in an experimental set up and demonstrated that the strain mapping task can be performed in nearly 2.4 h instead of the typical 6 h for a specific study on an AM steel part. Our work builds on a larger trend of moving towards autonomous scanning probe imaging systems [18, 25] in cases where the time to measure is significantly higher than the sum of the time to re-position the object and computation of the next informative measurement to make. In such systems, there can be an overall saving in experimental measurement time while providing a progressively improved image resolution by employing a powerful algorithm that can infer the underlying structure from a sparse set of measurement points. In the future, we plan to extend the algorithm to process strain data from multiple detectors using multi-output GP [29] and demonstrate its impact on 3D grids where the time-savings can be even more dramatic due to the increased dimensionality of the problem. This should in principle enable full 3D reconstruction of the strain tensor, with the data savings coming from correlations of the strain tensor components in addition to the spatial correlations.

Data availability statement

At the time of writing this manuscript, we are still using the data towards a scientific publication (not related to the method or algorithms). Hence we are unable to share the data immediately. The data that support the findings of this study are available upon reasonable request from the authors.

Acknowledgment

ORCID iDs

S V Venkatakrishnan ● https://orcid.org/0000-0002-7637-4987
Chris M Fancher ● https://orcid.org/0000-0002-3952-5168
Maxim Ziatdinov ● https://orcid.org/0000-0003-2570-4592
Rama Vasudevan ● https://orcid.org/0000-0003-4692-8579
Kyle Saleeby ● https://orcid.org/0000-0002-7523-8174
James Haley ● https://orcid.org/0000-0002-9023-3048
Dunji Yu ● https://orcid.org/0000-0001-8946-7851
Ke An ● https://orcid.org/0000-0002-6093-429X

References

[1] Wang X-L 2006 *JOM* **58** 52–57
[2] Wissink M L, Chen Y, Frost M J, Curran S J, Rios O, Sims Z C, Weiss D, Stromme E T and An K 2020 *Proc. Natl Acad. Sci.* **117** 33061–71
[3] An K, Chen Y and Stoica A D 2019 *MRS Bull.* **44** 878–85
[4] Watkins T R, Bilheux H Z, An K, Payzant E A, Dehoff R R, Duty C E, Peter W H, Blue C A and Brice C A 2013 *Adv. Mater. Process.* **171** 23–27
[5] An K, Yuan L, Dial L, Spinelli I, Stoica A D and Gao Y 2017 *Mater. Des.* **135** 122–32
[6] Ikeda T, Hirose S, Uozumi H, An K, Chen Y, Seid A, Okayama T and Katsurai T 2020 *SAE International Technical Paper* 2020–01–1071

[7] Tremsin A, Gao Y, Makinde A, Bilheux H, Bilheux J, An K, Shinohara T and Oikawa K 2021 *Addit. Manuf.* **46** 102130

[8] Nycz A *et al* 2021 *Mater. Des.* **205** 109751

[9] Stevens A, Yang H, Carin L, Arslan I and Browning N D 2014 *Microscopy* **63** 41–51

[10] Hwang S, Han C W, Venkatakrishnan S V, Bouman C A and Ortalan V 2017 *Meas. Sci. Technol.* **28** 045402

[11] Kelley K P, Ziatdinov M, Collins L, Susner M A, Vasudevan R K, Balke N, Kalinin S V and Jesse S 2020 *Small* **16** 2002878

[12] Godaliyadda G M D P, Ye D H, Uchic M D, Groeber M A, Buzzard G T and Bouman C A 2017 *IEEE Trans. Comput. Imaging* **4** 1–16

[13] Zhang S, Song Z, Godaliyadda G M D P, Ye D H, Chowdhury A U, Sengupta A, Buzzard G T, Bouman C A and Simpson G J 2018 *Anal. Chem.* **90** 4461–9

[14] Scarborough N M *et al* 2017 *J. Synchrotron Radiat.* **24** 188–95

[15] Pinkard H and Waller L 2022 *Nat. Methods* **19** 1–2

[16] Frazier P I 2018 arXiv:1807.02811

[17] Noack M M, Yager K G, Fukuto M, Doerk G S, Li R and Sethian J A 2019 *Sci. Rep.* **9** 1–19

[18] Ziatdinov M, Kim D, Neumayer S, Vasudevan R K, Collins L, Jesse S, Ahmadi M and Kalinin S V 2020 *npj Comput. Mater.* **6** 1–7

[19] Teixeira Parente M, Schneidewind A, Brandl G, Franz C, Noack M, Boehm M and Ganeva M 2022 *Front. Mater.* **8** 633

[20] Granroth G *et al* 2018 *J. Appl. Crystallogr.* **51** 616–29

[21] Kozubal A, Dalesio L, Hill J and Kerstiens D 1989 Experimental physics and industrial control system *ICALEPCS89 Proc.* vol 288

[22] Matthew N, Till S *et al* 2018 Non-linear least-squares minimization and curve-fitting for Python (available at: https://lmfit.github.io/lmfit-py/)

[23] Williams C K and Rasmussen C E 2006 *Gaussian Processes for Machine Learning* vol 2 (Cambridge, MA: MIT Press)

[24] Shahriari B, Swersky K, Wang Z, Adams R P and de Freitas N 2016 *Proc. IEEE* **104** 148–75

[25] Ziatdinov M, Liu Y, Kelley K, Vasudevan R and Kalinin S V 2022 *ACS Nano* **16** 13492–512

[26] Bingham E, Chen J P, Jankowiak M, Obermeyer F, Pradhan N, Karaletsos T, Singh R, Szerlip P A, Horsfall P and Goodman N D 2019 *J. Mach. Learn. Res.* **20** 28:1–6 (available at: http://jmlr.org/papers/v20/18-403.html)

[27] Gardner J, Pleiss G, Weinberger K Q, Bindel D and Wilson A G 2018 *Advances in Neural Information Processing Systems* vol 31

[28] Wang Z, Bovik A C, Sheikh H R and Simoncelli E P 2004 *IEEE Trans. Image Process.* **13** 600–12

[29] Alvarez M A, Rosasco L and Lawrence N D *et al* 2012 *Found. Trends Mach. Learn.* **4** 195–266