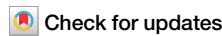


<https://doi.org/10.1038/s41524-025-01771-7>

Demonstration of an AI-driven workflow for dynamic x-ray spectroscopy

**Ming Du** ¹✉, **Mark Wolfman** ^{1,2}, **Chengjun Sun** ^{1,2}, **Shelly D. Kelly** ¹✉ & **Mathew J. Cherukara** ¹✉

X-ray absorption near edge structure (XANES) spectroscopy is a powerful technique for characterizing the chemical state and symmetry of individual elements within materials, but requires collecting data at many energy points which can be time-consuming. While adaptive sampling methods exist for efficiently collecting spectroscopic data, they often lack domain-specific knowledge about the structure of XANES spectra. Here we demonstrate a knowledge-injected Bayesian optimization approach for adaptive XANES data collection that incorporates understanding of spectral features like absorption edges and pre-edge peaks. We show this method accurately reconstructs the absorption edge of XANES spectra using only 15–20% of the measurement points typically needed for conventional sampling, while maintaining the ability to determine the x-ray energy of the sharp peak after the absorption edge with errors less than 0.03 eV, the absorption edge with errors less than 0.1 eV; and overall root-mean-square errors less than 0.005 compared to traditionally sampled spectra. Our experiments on battery materials and catalysts demonstrate the method's effectiveness for both static and dynamic XANES measurements, improving data collection efficiency and enabling better time resolution for tracking chemical changes. This approach advances the degree of automation in XANES experiments, reducing the common errors of under- or over-sampling points near the absorption edge and enabling dynamic experiments that require high temporal resolution or limited measurement time.

X-ray absorption spectroscopy measures the probability of materials to absorb x-rays impinging on them with varying energies. These x-ray energy-space measurements are correlated with the atomic or molecular properties of materials, making spectroscopy an important tool for material characterization. In particular, X-ray absorption near-edge structure (XANES) spectroscopy is an indispensable technique in revealing chemical properties such as the oxidation and bonding states of specific atoms within the material, and has been used in studying actinide bonding¹, spatial phase mapping^{2,3}, and investigating the charging/discharging cycles of batteries⁴, to name a few. When x-rays impinge on the atoms of a material, they are absorbed significantly more when the x-ray energies surpass the binding energy of core electrons of a specific element in the material resulting in an absorption discontinuity, known as an absorption edge. Hence, each element exhibits multiple absorption edges representing the binding energies of the element's electrons. The sharp rise in the absorption coefficients is called the *K*-edge (for 1s electrons) and *L*₁, *L*₂, and *L*₃-edges (for 2s and 2p_{1/2} and 2p_{3/2} electrons). The absorption spectrum of an element in materials is altered by factors such as the oxidation and bonding symmetry of the element. Since the absorption edge is a direct measure of the x-ray energy

required to excite a core electron to unoccupied electronic states, the absorption edge shifts depending on the oxidation state of the element, and pre-edge peaks appear when transitions are allowed to unoccupied electronic states below the Fermi level (unoccupied bound states). Strong absorption just above the absorption edge, due to transitions into unoccupied states above the Fermi level, gives rise to the so-called "white line". Fine structure oscillations beyond the absorption edge are due to the interference of the excited photo electron with the electrons of the neighboring atoms. Long multiple scattering photo-electron paths are possible for low energy photo-electrons which can give rise to resonant features just above the absorption edge^{5,6}. Measuring these fine features therefore reveals the phase and chemical state of elements within the sample.

XANES spectroscopy is traditionally performed by sequentially scanning the x-ray energy across the absorption edge using an energy-adjustable monochromator (such as a double crystal monochromator). During the scan the spectra are recorded by collecting the incident and transmitted x-ray intensity for each x-ray energy point. This is called a transmission mode measurement. The absorption coefficient is obtained by taking the log ratio of the incident to transmitted x-ray intensity following Beer's law. The

¹Advanced Photon Source, Argonne National Laboratory, Lemont, IL, 60439, USA. ²These authors contributed equally: Mark Wolfman, Chengjun Sun.✉ e-mail: mingdu@anl.gov; skelly@anl.gov; mcherukara@anl.gov

absorption coefficient is also proportional to the fluorescent x-rays that are produced as a result of the core-hole produced by exciting the core-electron. Such that the absorption coefficient can also be obtained by measuring the ratio of fluorescent x-ray intensity normalized by the incident x-ray intensity called fluorescence-mode measurement. The time consumed by the measurement is thus proportional to the number of energy points scanned and the time spent at each of those points to collect sufficient measurement statistics. There are two common ways in which the monochromator is scanned in an XANES measurement. Historically, the monochromator is stepped from one energy point to the next energy point. The step size in x-ray energy between the points is varied depending on the part of the x-ray spectrum being measured. Typical step sizes are 5 to 10 eV in the smoothly varying pre-edge region, 0.1 to 1 eV in the fast changing absorption edge region, and 1 to 5 eV in the more slowly varying post edge region. One common failure mode of this style of data collection is the improper definition of the boundaries of each region, which can happen without precise knowledge about the structure of the sample's spectrum. The high-level division of step sizes or scan velocities is also coarse and unable to adapt to localized fine features within each region. The 5 eV step size used for the pre-edge region can cause a pre-edge feature to be completely missed; errors also occur if the step size in the absorption edge region is too small (the monochromator does not actually move) or too large (features are not well defined)⁷. More recently, monochromators have become capable of slew scanning and a uniform grid of energy points becomes the simplest data collection mode. This results in over-sampling in some regions (pre- and post-edge) of the spectra necessitating post processing such as rebinning of data points. In either collection mode, if the scan is set up properly, the density of measured points are restricted to the highest point density required for a rather large region of the entire spectrum, which often results in more points than what is needed to sample the essential information.

The number of points to sample is a even greater concern for dynamic XANES experiments. In a dynamic process where the sample's chemical properties undergo continuous variation often under external stimuli such as voltage⁴ and temperature⁸, XANES can be done repeatedly during the process to characterize the sample's states at different time points. For example, for a chemical process that involves changes in the oxidation state of a certain element, the position of the absorption edge and the fine structure near it often differ in XANES spectra taken before and after the process, and measuring XANES spectra during the transition thus allows one to find out the progress of the process at the time when that spectrum is measured. In turn, plotting the transition progress against time reveals the kinetics of the transition. The time resolution of dynamic XANES is as important as its energy resolution in each spectrum. A sampling strategy that shortens the per-spectrum acquisition time by reducing the number of points while maintaining the energy resolution of the spectra is thus highly desired.

Adaptive sampling strategies have been developed to break the stalemate between resolution and experiment time. Adaptive sampling algorithms run during an experiment, receive the measured values or instrument readouts in real-time, analyze the updated data, and suggest the points to measure or actions to take in order to maximize the information gain from the measured object. The points measured under the guidance of adaptive sampling algorithms are often sparse and do not necessarily lie on a regular grid. Given the measured values, one may reconstruct the measured signal in an arbitrarily dense grid through interpolation assuming that the missing data lies smoothly between the measured points. A successful adaptive sampling algorithm is distinguished by its ability to suggest a small number of measurement points and yield a spectrum capturing maximal spectral features.

A survey of existing adaptive sampling methods leads us to categorize them into data-driven approaches and Bayesian optimization approaches. Data-driven approaches refer to those that use a trained model to predict the best action to take based on measurements made in the past and the current environment (such as sample conditions or probe positions). One representative technique is reinforcement learning⁹, where a model is gradually

trained to make the best move (e.g., to suggest the next location to sample) using the reward (e.g., the reduction of the reconstructed signal's error after taking a certain measurement) from past actions. Another variant is represented by algorithms named SLADS-Net¹⁰ and FaST¹¹. Used in scanning microscopy to sample a 2D image, these methods employ a neural network trained to predict the reduction of the reconstructed image's error after measuring at a point, with the positional local contextual information of that point as the input to the network. During an experiment, the error reduction is predicted on a set of points, and the ones with the highest predicted error reduction are sampled in the next step. While they have demonstrated success in reported use cases, a common concern about data-driven approaches, as noted in ref. 12, is that they need to be trained on data from previous experiments, but the transferrability of the trained models to test samples is unclear due to the lack of interpretability.

Bayesian optimization (BO) approaches on the other hand use a random process, which is often a Gaussian process (GP), to model the expected values and the uncertainties of the signal being measured. GP-based BO has been used for the adaptive sampling in various characterization techniques. Noack et al.¹² introduced the applications of such a Bayesian optimization method in the 2D space-domain sampling of spatially resolved x-ray scattering (also reported in ref. 13), Fourier transform Infrared absorption spectroscopy (FTIR) (also reported in ref. 14), angle-resolved photoemission spectroscopy (ARPES), and the reciprocal space-energy-domain sampling of inelastic neutron scattering. In 1D spectroscopy where sampling is done in the energy domain, Bayesian optimization has also been used for x-ray magnetic circular dichroism spectroscopy¹⁵ and x-ray absorption spectroscopy (XAS)¹⁶. Closer to our intended application is¹⁷, where Bayesian optimization is applied for grazing-exit XANES.

For a GP, uncertainty quantification is realized through the kernel function $k_{\theta}(x_1, x_2)$ that estimates the correlation between two points given their positions in the measured space, x_1 and x_2 . Commonly, a stationary kernel function such as the radial basis function or Matérn kernel¹⁸ is used, which estimates the correlation purely by their distance, i.e., $k_{\theta}(x_1, x_2) = k_{\theta}(|x_1 - x_2|)$, where θ is the set of hyperparameters defining the kernel's behavior. If point 1 is a measured point and point 2 is unmeasured, a larger $|x_1 - x_2|$ generally results in a smaller correlation and thus higher uncertainty for the value of point 2, and the decay rate of correlation is set by the length scale of the kernel which is a part of θ and can be estimated using a sparse collection of initial measurements. A GP still needs to be constructed with data, but the amount of data needed is far less than data-driven approaches and they can come from a sparse set of initial measurements at the beginning of the current experiment. Once the GP is constructed, it then yields the estimated value and uncertainty for any point x in the sampled space. In statistics terms, these are the posterior mean $m(x)$ and posterior standard deviation $s(x)$. These quantities are used to formulate an acquisition function $a(x)$ that quantifies the benefit of taking a measurement at x . The next measurement is done at the maximum point of $a(x)$, i.e., $x^* = \text{argmax}_x a(x)$. The measured value y^* along with x^* is used to update the GP, which then yields a new suggestion. This process is repeated until a stopping criterion is met.

Compared to purely data-driven approaches, BO has the advantages of not requiring a large amount of training data, and is more resistant to drift from training set distribution. Moreover, BO provides uncertainty quantification naturally, and it is straightforward to define and regulate the behavior of a BO algorithm analytical using physical and statistical knowledge. These features make BO a favorable candidate for adaptive sampling.

From the perspective of adaptive sampling algorithm design, a unique characteristic of spectroscopic techniques like XANES is that we often have good prior knowledge about the spectra being measured. XANES spectra are typically composed of a pre-edge region, pre-edge peaks a few eVs below the edge, the absorption edge itself, and the fine structure above the edge. We roughly know the forms they take: the pre-edge region should have a smooth slope, the absorption edge is characterized by a large discontinuity, while the post-edge features are wavy undulations. Injecting this knowledge to a adaptive sampling algorithm and instructing it to sample more densely

around the varying regions can greatly enhance the efficiency of sampling as it reduces the unnecessary exploration in the less interesting slowly varying regions.

Moreover, an XANES spectrum in a dynamic process can be well-approximated by the linear combination of at least 2 spectra collected at different times of the process^{19,20}. It follows that if the spectrum reconstructed during an experiment is fitted with reference spectra, higher fitting residue should indicate regions that need more measurements.

However, GP-BO with standard components do not automatically take advantage of these knowledge and properties of XANES. In ref. 15, the use of an acquisition function only containing the posterior variance makes the algorithm unaware of the spectrum's intensity as the posterior variance is only dependent on the positions (energies) but not the observed values of measurements (see Eq. (8)).¹⁷ uses the upper confidence bound (UCB) of $m(x) + \kappa\sigma(x)$ as the acquisition function. While UCB is a common choice for solving *maximization problems* with Bayesian optimization, we should note that the objective of adaptive sampling is not to find the maximum of the function modeled by the GP (i.e., the spectrum) but to map or reconstruct the function as accurately possible with limited samples. In an adaptive sampling setting, the posterior mean term of the UCB would favor the sampling in regions where the observations are estimated to be high. Although this does encourage the algorithm to find high peaks (especially the top of the white line) which are indeed important in XANES, our concern is that the sampling of pre-edge peaks and the onset of the absorption edge will be delayed until the uncertainty of high-absorption regions drops low enough. Additionally, at the early stage of the experiment, the position of the white line estimated by GP may be inaccurate, potentially leading the algorithm to overly exploit areas that are not the true white line. We believe these limitations can be overcome by developing a method that is injected with the structural knowledge of XANES and leverages it to optimize the sampling distribution in the energy domain.

How can the structural and data-derived knowledge be incorporated in a GP-BO algorithm? Reference¹² listed 3 strategies for domain knowledge injection, namely (1) using a (potentially non-stationary) kernel function that better reflects the length scale of signal variations, (2) constraining the fitting of kernel hyperparameters, and (3) designing an acquisition function that better reflects the benefits and costs of measuring at given locations. We consider (3) as the most straightforward approach because the maximum of the acquisition function is what directly dictates the suggestion of a BO-based algorithm. Also, since the acquisition function is defined in the same domain as the input variable of the measured signal (energy in the case of XANES), we can easily apply our knowledge about the energy-domain structure of XANES spectra. A number of previous works have adopted the acquisition function engineering strategy for other characterization techniques: in an adaptive sampling algorithm designed for x-ray scattering²¹, an acquisition function is designed to take into account the posterior uncertainty, the value and spatial gradient of the estimated (i.e., densely reconstructed) signal, and the cost of taking a particular measurement (e.g., the time consumed by moving the motor). In ref. 16, the authors use a GP to model the difference between the XAS spectrum predicted with a physics model and the observed values, and the acquisition function for sampling is the UCB of this difference; the hyperparameters used in the physics model are searched by another BO. It should be noted that an acquisition function designed specifically for a certain characterization technique is not guaranteed to be applicable to another technique due to the difference in the prior knowledge and assumptions involved.

Here we introduce our adaptive sampling algorithm where an acquisition function is designed and engineered with the prior structural knowledge about XANES spectra. As described in section 4, the main design features of our acquisition function include:

- A first- and second-order gradient component that returns higher values at peaks, values and slopes of the reconstructed spectrum;
- A fitting residue component that responds to the difference between the reconstructed spectrum and the spectrum linearly fitted using two reference spectra;
- An acquisition reweighting function that automatically detects the location and width of the absorption edge, based on which it reduces acquisition function values in the pre-edge smoothly varying region and boosts its values in the high-variation region above the edge.

We test our algorithm in three simulated case studies and one real-world experiment. We first assess our algorithm on the sampling of single XANES spectra and closely examine the behavior of the algorithm during data collection. Comparing the reconstructed spectra with the densely sampled ground truths, we show that the overall root-mean-squared (RMS) errors of the normalized spectra are generally below 0.005; moreover, in most cases, our method requires only 15–20% of the number of points measured in the ground truth. The error reduction with our method is faster than with either uniform grid sampling or a BO algorithm that uses alternative acquisition functions. We then use our algorithm to collect spectra in dynamic XANES experiments tracking the evolution of the chemical states of certain elements in materials during dynamic processes. The percentage of transition of the chemical state is calculated for each spectrum during the process through reference spectrum fitting, and the curve of the transition progress is compared with that obtained using traditionally densely sampled spectra to show that our method reveals information about the kinetics of the process as accurately as traditional sampling while requiring a much smaller number of measurements. Additionally, we show that the adaptively sampled spectra allow one to determine the energy of the maxima of the white line with an error of at most 0.03 eV, and the absorption edge with an error of at most 0.1 eV. Finally, we also demonstrate the use of our algorithm in a live experiment at a synchrotron beamline where it communicates with beamline instruments in real-time to guide the in-situ XANES measurements of a battery electrode. An illustration of our method's workflow using this experiment is shown in Fig. 1.

Results

Simulated sampling of a single spectrum from the YBCO sample

We first apply our method in sampling the XANES spectrum of yttrium barium copper oxide (YBCO), a compound exhibiting high-temperature superconductivity. The data for this simulated test case was collected in a previous experiment with manually chosen scan grid; in the range between 8920 and 9080 eV, the original spectrum was acquired with a step size of 10 eV below 8960 eV, 0.3 eV between 8960 eV and 9012 eV, and 1.5 eV above 9012 eV. This constitutes a total of 218 points in the ground truth data. For our test, we get the virtually measured values at the energies suggested by the algorithm through linearly interpolating the ground truth. The data used for sampling are unprocessed, meaning their values may contain sample-specific scaling in edge height, and a slope in the energy domain. This simulates the data condition that our sampling algorithm may encounter in an actual experiment. As mentioned in section 4.3, the algorithm is allowed to stop before the set number of measurements is reached if the maximum weighted uncertainty drops below a threshold. The maximum posterior standard deviation threshold for early stopping was set to $t = 0.05$. We used 10 uniformly positioned initial measurements to build the GP; the length scale was fit on reference spectra and was found to be 8.8 eV. The stopping criterion was triggered after 40 additional measurements giving a total of 50 measurements. The reconstructed spectrum given by the GP's posterior mean, the posterior standard deviation (represented by half of the vertical length of the shaded area at each point), the measured points, and the ground truth are shown in Fig. 2. These quantities are plotted for 10, 25, and 40 measured points (including the 10 initial samples). With the constraint of value-aware terms in the comprehensive acquisition function and the acquisition reweighting function, the algorithm barely sampled points in the pre-edge smoothly varying region below 8960 eV during the first 15 guided measurements. Rather, it allocated most of the measurements on the absorption edge at 8080 – 9000 eV and the post-edge features above 9000 eV. In particular, the fast post-edge undulations between 9000 and 9030 eV are most densely sampled in the first 15 guided measurements. The algorithm started to explore the pre-edge and post-edge smoothly varying

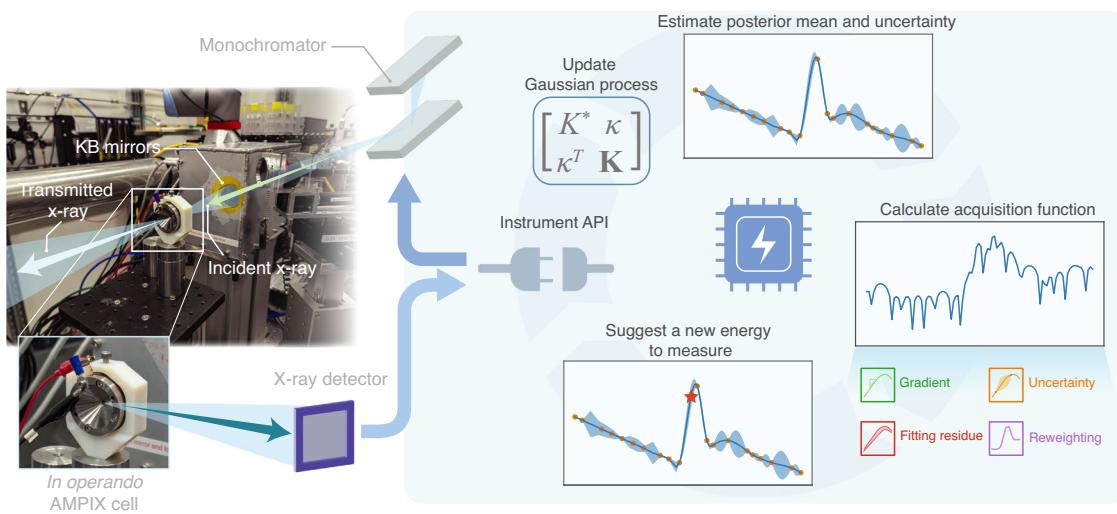
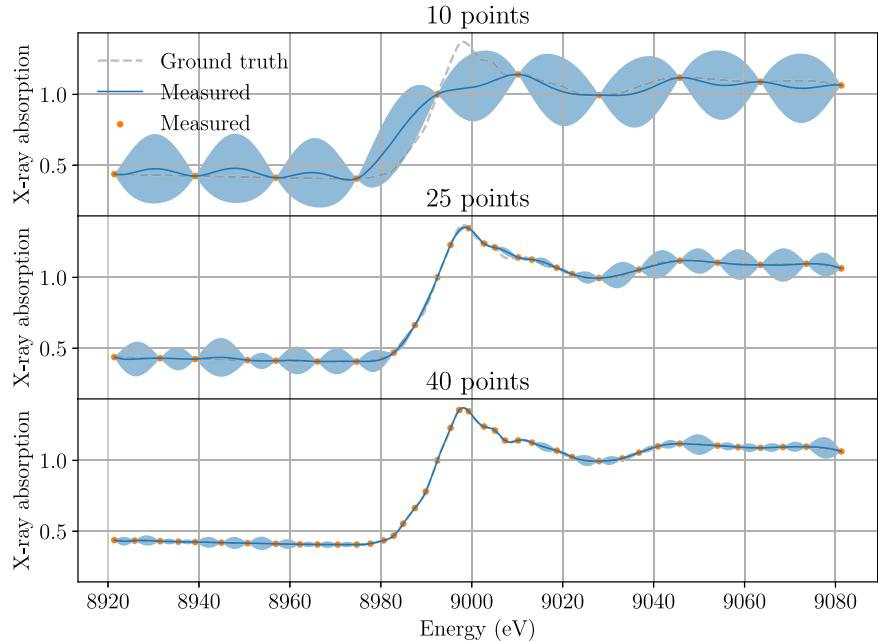


Fig. 1 | Diagram showing the workflow of our adaptive sampling method. The workflow is illustrated using the real-world battery cell experiment demonstrated in this paper. The backend algorithm receives data from the x-ray detector and sends data to the monochromator through the instrument's application programming

interface. With each measured energy, the algorithm updates its internal state, computes the comprehensive acquisition function and directs the monochromator to measure at a new energy. This cycle is repeated until the stopping conditions are met.

Fig. 2 | Intermediate reconstructed spectra, posterior standard deviation, measured data points, and true spectrum for the YBCO data. Data are plotted without normalization. The posterior standard deviation at each point is represented by half of the vertical length of the shaded area.



regions after about 20 guided measurements, at which point uncertainty became a more important driving force.

To quantitatively evaluate the accuracy of the spectrum sampled with our algorithm, we calculate the RMS error between the reconstructed spectrum and the ground truth after each new point is measured. For result interpretability, both spectra are normalized for RMS error calculation, such that the linear slope background is subtracted from the whole spectra, and the curvature in the post-edge region is corrected. This is done following the approach documented for the Athena package²²: first, a straight line $f = F_1(x)$ is fitted using pre-edge data and a quadratic function $f = F_2(x)$ is fitted using post-edge data. The fitting ranges are chosen manually. The fitted line is extrapolated to the entire energy range and subtracted from the spectrum. The difference between the two fitted functions at the edge energy $\Delta f = F_2(x_0) - F_1(x_0)$ is estimated to be the edge height, and the spectrum is scaled by $1/\Delta f$. Subsequently, the spectrum is flattened by subtracting $[F_2(x) - F_1(x) - \Delta f]/\Delta f$ for $x > x_0$. This completes the normalization, after which the

edge height becomes 1 and the pre-edge region lies on zero and the post edge oscillates about one. We use the full raw spectra that includes the data outside the adaptively sampled range (< 8920 eV and > 9080 eV) to fit F_1 and F_2 because they contain more smoothly varying segments. The other important feature of the XANES data is the absorption edge position in x-ray energy. This is often defined as the maximum of the first derivative of the absorption edge or the first zero crossing of the second derivative.

We compare the RMS values of our method with 5 reference cases served as the baseline: (1) a GP-BO algorithm that is otherwise the same as our proposed method but without the acquisition reweighting; (2) a GP-BO algorithm that only uses posterior standard deviation as the acquisition function; (3) a GP-BO algorithm using the upper confidence bound (UCB) acquisition function; (4) a GP-BO algorithm using the expected improvement (EI) acquisition function²³; and (5) uniform sampling. In uniform sampling, we always regenerate a uniformly distributed grid of points for a different number of measurements: for example, in a normalized space, we

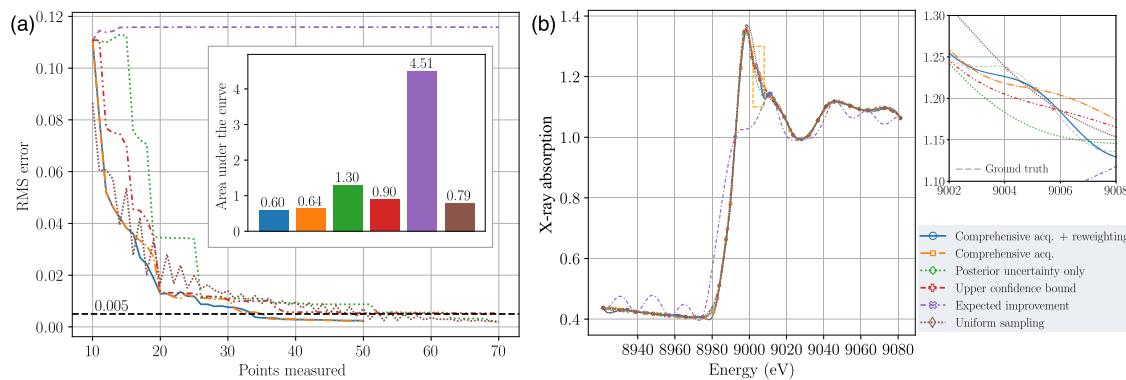


Fig. 3 | Results of single-spectrum sampling of the YBCO sample. Legend at the bottom right corner applies to both sub-figures. **a** Convergence of RMS error with our comprehensive acquisition function and acquisition reweighting, and its comparison with cases (i) with comprehensive acquisition but without reweighting, (ii) with the posterior uncertainty-only acquisition function, (iii) with the upper confidence bound acquisition function, (iv) with the expected improvement acquisition function, and (v) with uniform sampling. The inset bar chart on the side shows the

areas under the curve before the 50th measurement for each case with the same color coding to the RMS error curves. **b** Reconstructed spectra using these methods after 30 points are sampled (initial measurements included). The inset shows the magnified spectra within the energy range between 9002 and 9008 eV. This offers a direct comparison between the comprehensive acquisition method with and without reweighting, with the latter exhibiting larger errors in the plotted region.

sample at $[0, 0.5, 1]$ for 3 measurements, and at $[0, 0.33, 0.67, 1]$ for 4 measurements. As such, uniform sampling is not sequentially consecutive as in other cases and should not be examined for different numbers of measurements; it should only be compared with other cases for the same numbers of measurements. For UCB, the value of the trade-off parameter κ (Eq. (20)) was determined by running the experiment with a range of κ values from 1 to 15, and choosing the one resulting in the smallest area under the curve of RMS error plotted against the number of measurements. More details about the definition of area under the curve (AUC) will be introduced shortly.

Figure 3a plots the RMS error compared with the ground truth against the number of measurements for all cases. Cases with our acquisition function (with and without acquisition reweighting) stopped at 50 measurements when the stopping criterion was triggered; the rest were run until 70 total measurements were performed. GP-BO with the EI acquisition function performs the worst: the error first increases due to badly conditioned interpolation of the reconstructed spectra, and then remains high. This is not surprising because EI is designed for maximization problems instead of adaptive sampling; during the run, it quickly located the top of the white line, but then kept sampling at that point, ignoring the rest of the spectrum. The posterior uncertainty-only acquisition eventually led to a low RMS error below 0.005, but it took over 60 total measurements to achieve, making it worse than uniform sampling that took 45 measurements to reach an RMS error of 0.005. By introducing the posterior mean term that turns high-absorption features more favorable, UCB performed better and exhibited faster error reduction. However, our acquisition function, with or without reweighting, led to faster error reduction at the early stage. The case without reweighting reached 0.005 RMS error with 35 points and the one with reweighting function reached the target error with 33 points.

An alternative metric to compare the error reduction speed of these methods is the area under the curve (AUC) of the RMS error convergence plots. This method is also employed in ref. 24. Smaller AUCs indicate faster error reduction. For each method, we calculated the AUC before the 50th measurement and show them in the inset of Fig. 3. The AUCs of methods with the comprehensive acquisition function are obviously smaller than other approaches. Comprehensive acquisition function with reweighting has the smallest average AUC of 0.60, which marks the best performance among the tested methods. Without reweighting, the method attained an AUC of 0.64. The larger AUC is due to measurements that are unnecessarily spent in the pre-edge region. The performance improvement brought by the reweighting function can be more clearly seen in Fig. 3(b) which shows reconstructed spectra using the listed methods after 30 points (including

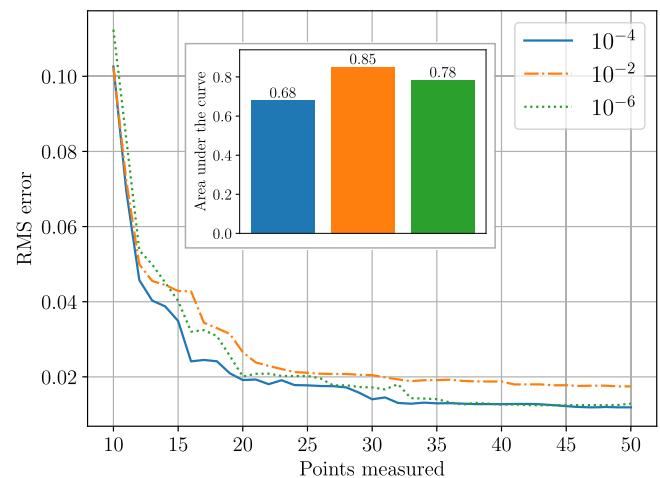


Fig. 4 | The RMS errors of reconstructed spectra over the number of measurements with 3 different noise variances set to the GP. The true observation noise variance is 10^{-4} . The inset shows the AUCs of the three experiments, color coded to match the RMS plot.

initial points) is measured along with the ground truth. Particularly, in the range between 9002 and 9008 eV which is magnified in the side plot, the comprehensive acquisition function with reweighting better captures the local structure than all other methods including the no-reweighting counterpart.

While the data used for this simulated study has a rather high signal-to-noise ratio, data from real experiments are occasionally noisy. When the GP is not sufficiently regularized, its posterior mean can overfit, leading to spurious fluctuations. To investigate the effect of the noise prior, we artificially added isotropic and independent Gaussian noise with a variance of 10^{-4} to the observations and ran adaptive sampling experiments with 3 variance values for the noise prior (σ_o^2): 10^{-2} , 10^{-4} , and 10^{-6} . We also enabled adaptive noise variance described in Section 4.1 to maintain reconstruction accuracy of the white line. The prior variances are incorporated in the GP by setting the noise covariance matrix as $N = \sigma_o^2 I$. Figure 4 shows the RMS error curves and AUCs of the 3 runs. The fastest error reduction and smallest AUC is attained when the prior noise variance matches the actual variance of 10^{-4} . A smaller prior variance resulted in larger error in the early stage of the experiment due to fluctuations of the posterior mean as a

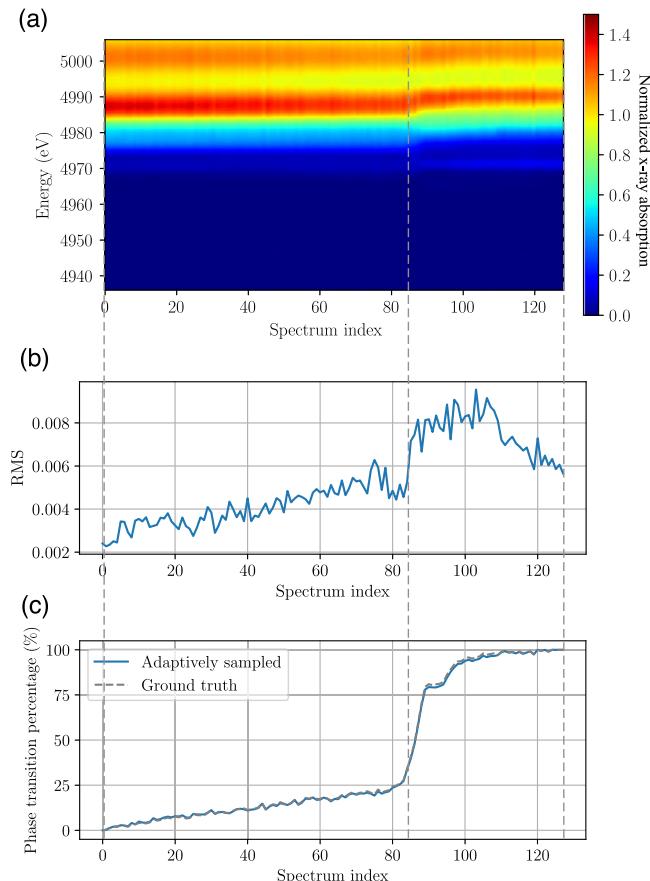


Fig. 5 | Results of applying our adaptive sampling method in a dynamic XANES experiment that tracks the phase transition progress of the LTO sample. **a** The stack of spectra at 50°C sampled using our method. Data plotted and normalized. **b** The RMS error of normalized spectra plotted against spectrum index. **c** The phase transition percentages calculated through linear fitting with data collected. Percentages were calculated using normalized data. **b** and **c** are aligned with **a** in the horizontal axis. The gray dashed lines marks the beginning, end of the experiment and the onset of the phase transition.

consequence of overfitting, which are particularly severe in the pre-edge region. A prior variance that is too large on the other hand results in even slower error reduction and higher final error because the posterior mean is over-smoothed and unable to capture fast varying features.

Simulated sampling in a dynamic XANES experiment on the LTO sample

In this section, we apply our adaptive sampling method to a dynamic Ti *K*-edge XANES experiment intended for tracking the progress and studying the kinetics of the phase transition of lithium titanium oxide (LTO), a battery material. Raw XANES data were collected during the phase transition of the material at 2 different temperatures, 50 °C and 70 °C. In the range between 4936 and 5006 eV, the ground truth was sampled with a uniform step size of 0.5 eV, resulting in 141 points per spectrum. 128 spectra were collected for the phase transition at 50 °C, while 14 were collected at 70 °C due to the faster transition rate. Since the start and end phase of both temperatures are the same, we use the first and last spectra collected at 70 °C as the reference spectra used for the fitting residue term of the comprehensive acquisition function, and use the one at 50 °C for sampling test. Line plots of the normalized ground truth spectra at the beginning and end of the phase transition at both temperatures are shown in Fig. S1 of the Supplemental Document.

Since the XANES spectrum evolves during a dynamic experiment, we run adaptive sampling independently for every measured spectrum.

For each spectrum, we used 10 uniformly sampled initial points to build the GP. The maximum number of total measurements was set to 40. With these settings, we sampled the 128 spectra in the test set. The stack of sampled data is plotted in Fig. 5a. For presentation purposes, the data plotted are normalized. The RMS error between each reconstructed spectrum is calculated with the corresponding ground truth after normalization, and is plotted in Fig. 5b. The RMS error remains around or below 0.005 before spectrum 85, after which it starts to increase, but stays below 0.01. This inflection point is marked by the gray vertical dashed line across Fig. 5; the horizontal axes of all the plots are aligned. In (b), the marked position shows an obvious change of spectral structure, indicating the onset of phase transition. The phase transition results in a sharper and higher pre-edge peak at 4972 eV which is the main source of the increased RMS error, but overall this peak is still well captured in the reconstructed spectrum. The reconstructed spectra before and after this transition (spectrum 80 and 90), before normalization, are shown in Fig. S2 in the Supplemental Document; the intermediate estimated spectrum and uncertainties of spectrum 90, which has an obvious pre-edge peak, is shown in Fig. S3.

Using the first and last spectra of the reference set as bases, we linearly fit each reconstructed spectrum, and the obtained coefficients were used to estimate the percentage of phase transition when that spectrum was measured. The trajectories of the transition percentage calculated using the adaptively sampled spectra as well as the ground truth data are plotted in Fig. 5c. Despite the slightly elevated RMS error when the phase transition begins, the percentages calculated using adaptively sampled spectra are very close to the ones obtained with ground truth data: the maximum error in the percentage is only 1.7%, while the linear combination fitting analysis typically has an uncertainty of a few percents. Considering that a maximum of 40 points are sampled for each spectrum whereas the ground truth data were collected with 141 points per spectrum, our method reduced the number of measurements by 72% without losing accuracy in the phase transition's kinetic information.

Simulated sampling in a dynamic XANES experiment on the Pt/γ-Al₂O₃ sample

In this section, we present a case study on sampling the XANES spectra of a 0.35 wt% Pt on γ-Al₂O₃ catalyst sample during an in-situ reduction process. During the collection of the original data, the sample was heated from 26 °C to 497 °C in H₂, resulting in continuous reduction of the Pt oxidation state. Pt XANES measurements were recorded during this process. Details regarding the sample preparation and reduction conditions can be found in ref. 25. The strong peak at the absorption edge, also known as the white line, corresponds to the empty 5d electrons of Pt⁴⁺. As the Pt becomes reduced and these empty states become occupied, the XANES white-line intensity decreases²⁶. This sample is particularly interesting for this study due to the low Pt loading resulting in a small edge step and the large background absorption from the Al₂O₃ creating a large slope in the raw x-ray absorption spectra. 22 ground truth spectra were collected throughout the reduction process. In the range between 11400 eV and 11700 eV, 238 equally spaced energy points were measured.

Compared to the YBCO sample, Pt's spectra at room temperature has a sharp and high white line, i.e., the absorption peak immediately following the absorption edge. This poses a major challenge to adaptive sampling as the accurate reconstruction of the white line requires at least one measurement close enough to the peak's maximum, and any deviation would cause the white line's height or position to be misrepresented in the reconstructed spectrum. The second-order derivative and fitting residue terms in our comprehensive acquisition function are particularly helpful in locating the peak's maximum, and as such, we specify relatively high weight values for both terms ($\phi_{g'} = 3 \times 10^{-2}$, $\phi_r = 10^3$). We take the first (26°C) and last (497°C) spectrum as the reference data, and the remaining 20 spectra are used for adaptive sampling. 10 initial measurements are collected uniformly at the beginning of each run and the length scale parameter is fit on the reference data.

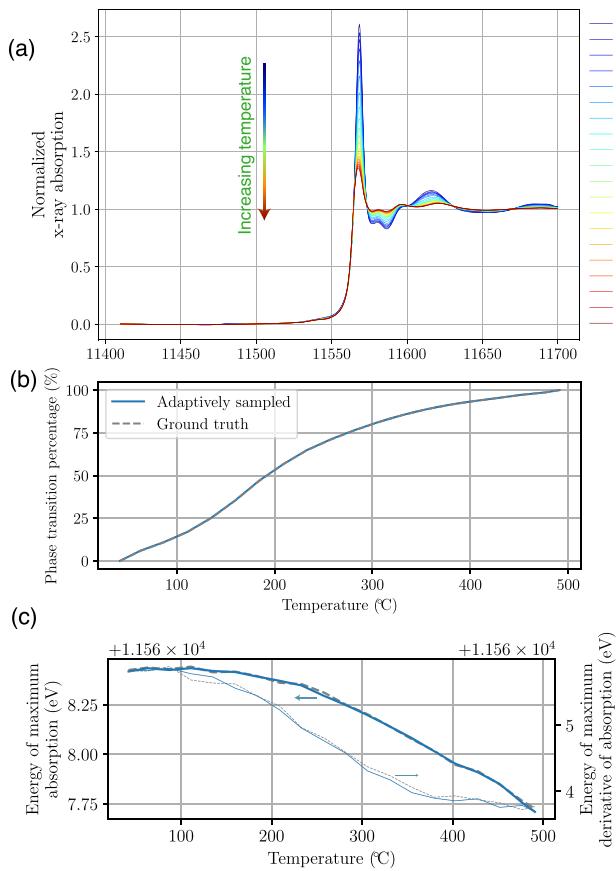


Fig. 6 | Dynamic experiment results of the Pt sample. **a** All normalized spectra during the reduction process, sampled using our method. **b** The transition percentages calculated using normalized data sampled with our method. **c** The energies of the maxima of the white lines (bold lines), as well as those of the first derivative of the absorption edge (thin lines), plotted for all spectra over the reduction process. Arrows beside curves point to the y-axis that should be referred to.

We start by examining the performance of our algorithm on a single spectrum. We choose the spectrum measured at 42 °C which features a high white line. Fig. S4 in the Supplemental Document shows the sampling process at 10, 30, and 50 total measured points. In addition to the high white line, Fig. S4 also reveals a large slope in the background due to the large absorption of the Al_2O_3 substrate, but our algorithm still works robustly, managing to capture the maximum of the white line before the 30th measurement, with other features also decently sampled in the mean time. We also compare the RMS error convergence of our method (with and without acquisition function reweighting) with other acquisition functions (posterior uncertainty and UCB) and uniform sampling (Fig. S5a). Like the YBCO case, the trade-off factor κ of UCB was chosen to be the one giving the lowest AUC. The methods employing the comprehensive acquisition function clearly outperform the rest: an RMS error (on normalized spectra) of 0.005 is reached with about 46 total measurements, or 20% of the number of ground truth points in the sampled range. On the other hand, the baseline methods still have much higher RMS errors after 70 total measurements. The AUCs of the RMS convergence curves reveal the same trend, where the AUCs of methods with the comprehensive acquisition function is at least around 3 times lower than the others. The high errors of posterior uncertainty acquisition function and uniform sampling mainly originate from their failure to precisely capture the white line, as shown in Fig. S5(b) which compares the sampling process of all methods. On the other hand, UCB is the first to capture the maximum of the white line among all methods, but the sampling of other features is severely deficient despite of the tuned exploration-exploitation trade-off factor.

We then test our method on data of the dynamic XANES experiment. Figure 6a plots the 20 spectra (excluding the reference spectra) sampled using our method, which are normalized for showing purpose. The percentages of state evolution of the sample, calculated through the linear fitting to the normalized reconstructed spectra using the two reference spectra as bases, are shown in Fig. 6b. Again, the percentage trajectories computed using adaptively sampled spectra agree well with that calculated with the ground truth data, with the maximum error in the percentage being 0.2%.

The energy of the maximum of the white line, as well as the location of the absorption edge quantified by maximum point of the first derivative of the edge, are also important quantities that characterize the sample's chemical state. The $\text{Pt}/\gamma\text{-Al}_2\text{O}_3$ sample features a sharp white line and steep edge which make it an ideal case to assess the accuracy of our algorithm in determining both quantities. For each spectrum during the heating process, we fit a quadratic function to the local absorption values in a small neighborhood around the maximum of the spectrum to determine its location in energy. The same is done for the first derivative of the spectrum. Figure 6c plots the evolution of the energy of the white line's maximum (thick lines) and that of the edge derivative's maximum (thin lines). The data for the adaptively sampled spectra (blue) and the ground truths (gray) are plotted in the same graph. Both quantities show good agreement, with a maximum error of less than 0.02 eV for the white line and less than 0.1 eV for the edge. The standard monochromator selects the incident x-ray energy with a FWHM of 1.3 eV at 10 KeV, so an error of less than 0.1 eV is smaller than our measurement resolution²⁷.

Validation in a real experiment

We have deployed our algorithm at the 25-ID-C beamline at the Advanced Photon Source and used it to guide a real-world experiment. The Bluesky infrastructure²⁸ allows the adaptive sampling algorithm to receive the measured intensities (and thereby the absorption coefficients), and send suggestions of the next energy to measure. With this control setup, we performed a smartly controlled in-situ XANES experiment that measures the absorptions of a lithium nickel manganese cobalt oxide ($\text{LiNi}_{1/3}\text{Mn}_{1/3}\text{Co}_{1/3}\text{O}_2$) battery electrode during its discharging process. We first collected 2 reference spectra on a pristine sample and a fully charged (4.5 V) sample in the range between -100 eV to $+100$ eV relative to the Ni K-edge of 8333 eV. We then sequentially measured 8 spectra under the guidance of our algorithm as the in-situ cell was discharged from 4.5 V. To make the case more challenging for our algorithm, the adaptive spectra were collected in a larger range of -100 eV to $+150$ eV relative to the Ni edge, i.e., no reference data is available for the last 50 eV. This allows us to create a realistic scenario where the reference spectra are not completely available for the entire measured energy range, and to examine the balance of the fitting residue component of the acquisition function with other components, as an over-weighted fitting residue component would draw most of the samples to the energy range where the reference spectra are present. However, our algorithm still performed robustly. The normalized traditional spectra are plotted together in Fig. 7a, with the sampled points indicated by markers. In view of the noise in the pre-edge region, we used a noise variance of 10^{-3} in the reconstruction of the spectra. The first spectrum collected in the process is plotted in violet, and the last is in red. We also calculated the sampling density of all energies sampled among the 8 spectra using kernel density estimation, and plot it as the shaded area at the bottom of the figure. The sampling density is highest at the absorption edge and the white line, which are the most important features to track the change of oxidation state of an element during a dynamic process. Fast-varying features, like peaks and valleys in the post-edge region, are also more densely sampled. For example, the spectrum peak at 8405 eV is well aligned with a peak in the sampling density. Smoothly varying regions away from the edge are sampled less, reducing the number of measurements.

In terms of experiment time, the traditional dense collection of a 401-point reference spectrum took 620 seconds on average. Considering that the energy range of the reference spectrum is 4/5 of the adaptively sampled range, we scale it to estimate the time for densely sampling the 250 eV's

Fig. 7 | Dynamic experiment results of the NMC111 sample. **a** The 8 normalized spectra during the discharging process sampled using our method, plotted together and with the sampled points indicated by dots. From violet to red, the color of the lines and dots indicate the time sequence of the data. The densities of sampled points aggregated from all the 8 spectra are shown by the shaded region. **b** The transition percentages calculated using normalized data sampled with our method as comparison of the Ni^{2+} and Ni^{3+} standards. The gray dashed line is the regression line.

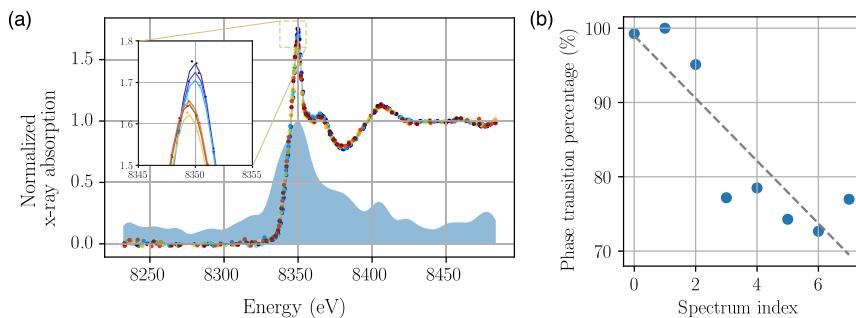


Table 1 | The number of measured points of adaptive and dense sampling, and their ratios, for cases shown in this paper

Sample	Type	Range (eV)	# points (adaptive)	# points (traditional)	Adaptive sampling ratio (%)
YBCO	Simulated	160	33	218	15.1
LTO	Simulated	70	40	141	28.4
Pt	Simulated	300	46	238	19.3
NMC111	Real	250	60	501 (scaled)	12.0

Since the reference spectra of the NMC111 case were collected in a smaller range (200 eV) than the adaptively sampled spectra (250 eV), we scale the actual number of points in reference spectra (401) to match the range of adaptively sampled spectra.

range, which gives 775 seconds. The adaptive sampling of 60 points, including the 20 initial measurements and the subsequent 40 adaptive measurements, took an average of 180 seconds, or 23.2% of dense sampling. This ratio is higher than that of the numbers of points (12.0%), but is still in a reasonable range considering the overhead of computation and the longer average travel distance of the motors to position the monochromator from the current energy to the next.

Using the reference spectra as bases, we estimate the percentage of Ni's oxidation as it reduces from around +4 to +2, as shown in Fig. 7(b). The plot reveals the decreasing trend in the oxidation state, expected for a discharging process; it also reveals that the reduction of Ni was not complete, which plateaus at about 75% or approximately +3.5. This observation provides insights about the electrode's cycling mechanisms for battery researchers.

We summarize the 4 test cases by listing the numbers of points measured adaptively (for the 3 simulated cases, they are the numbers of points to reach an RMS error of 0.005 compared to the ground truths) and densely sampled in Table 1. The ratios of the numbers of points sampled adaptively and densely sampled are also reported.

Discussion

Our adaptive sampling algorithm is based on GP-BO and leverages domain-knowledge specific to the structure of XANES spectra to guide the experiment such that important features are sampled more frequently. This knowledge is injected by designing an acquisition functions that returns higher values where (1) fast variations, peaks, or valleys are present, (2) the reconstructed spectrum is not well explained by the linear combination of reference spectra, and (3) the location is away from the smoothly varying region below the absorption edge. Artificially designing the acquisition function for XANES is straightforward because the measured space is only one dimensional, and assumption that an XANES spectrum sequentially contains a smooth pre-edge region, pre-edge peaks, an absorption edge, and post-edge features generally holds true. Additionally, the fact that an XANES spectrum is an average of the element-of-interest oxidation states, it can be represented by the linear combination of a few reference spectra

providing an effective analytical approach to assess the accuracy of the reconstructed spectrum. Therefore, our method works well without the need of extensive training data as in the case of data-driven approaches like reinforcement learning⁹ or SLADS-Net¹⁰, and is less susceptible to the generalization gap that data-driven methods often encounter when working with unseen test data. Our method also provides an explicit uncertainty quantification and the ability to guide the acquisition using it compared to the two data-driven approaches listed above. However, we would like to note that uncertainty quantification is also possible with data-driven approaches. One such example is deep kernel learning (DKL)²⁹. In DKL, the joint distribution of past and new observations is also modeled as a GP that uses a common (and possibly stationary) kernel function, but a neural network g_Φ is trained project the input features to a different space before they are processed by the kernel. In other words, the evaluation of the kernel is modified to $k_\theta[g_\Phi(x_1), g_\Phi(x_2)]$. The non-linear projection of g_Φ allows the kernel to adapt to different feature length scales in the data: points in the fast varying parts of the data can be projected to have larger spacing which is equivalent to reducing the length scale specifically in that region. However, g_Φ is not necessarily shift-equivalent; when applied to XANES, that means if the absorption edge's location in a test spectrum is different from those in the training data, g_Φ does not recognize this difference by itself and would still project input features based on their absolute positions from what it learns during training. This issue can be partly solved by shifting the training and test data in such a way that their absorption edges are all at a standardized position. This is easy to do for training data, but for test data the only way to find the correct standardizing shift is to detect the edge's location using a sparse set of initial measurements, which may be inaccurate. Additionally, since the structures and locations of fine features vary drastically across different samples, it is hard for g_Φ to learn a projection with sub-feature level precision. On the other hand, our method adaptively detects peaks, valleys and high-gradient areas using derivatives and under-described areas using fitting residues and escalates the acquisition function precisely in those areas. It might be possible to adapt g_Φ in DKL to also take these computed features into account instead of just the positions in the input space, but this is beyond the scope of this paper.

We would also like to discuss the relative merits of our method compared to Quick-XAFS (QXAFS), a fast acquisition technique for x-ray absorption spectroscopy backed by instrumentation instead of a steering algorithm. In QXAFS, the monochromator constantly moves in the energy space without stopping to acquire x-ray intensities. This continuous scanning approach allows QXAFS to measure an XANES spectrum in a few seconds. However, even compared with QXAFS, our adaptive sampling approach still offers two major advantages:

1. QXAFS often has more stringent requirements on instrumentation. For example, QXAFS uses a specialized monochromator that oscillates the crystals to sweep through energies rapidly and requires high-speed motors to support the fast and smooth motion of the monochromator; it also needs detectors that can respond to fast variations of x-ray intensity without signal smearing, for which a gridded ionization chambers has to be devised³⁰. Without these advanced instruments,

QXAFS can suffer from loss of energy resolution and accuracy due to motion blur and mechanical instability. On the other hand, while our method also benefits from advanced monochromator motors, it is more tolerant and more easily usable by endstations without these high-end instruments.

2. Although QXAFS is fast, it still exposes the sample under radiation for the entire energy range scanned. With an adaptive sampling method like ours, one can use the same effective exposure time for each energy point as QXAFS so as to obtain the same signal-to-noise ratio, while reducing the number of measurements and decreases the total radiation dose of the sample. This mitigates the radiation-induced degradation of radiation-sensitive specimens.

A general concern about GP-BO is the $O(n^3)$ scaling with the number of measured points in inverting the covariance matrix for posterior calculation. However, the one dimensional measurement space of XANES eases this concern as the size of the covariance matrix is well under control: the size of the matrix equals the number of measurements, and in our test cases shown above, measuring an energy range of 100–300 eV typically takes just 40–60 measurements. On a CPU, the time taken by our algorithm to make a prediction is below 0.1 second and the speed did not change significantly during the acquisition.

A limitation of our method, which is also common to GP-BO algorithms using common kernel functions, is that when the actual feature length scale in the data is extremely non-uniform, the posterior mean in sparsely sampled smoothly varying regions may be inaccurate. This is explained by Eq. (7) which shows that the posterior mean's residue from the prior mean is an average of that residue of past measurements weighted by $\kappa^T(K+N)^{-1}$. If the distance of a test point x^* to nearby past measurements is well above the kernel's length scale, its covariance with these points is underestimated and points lying farther away, including those with very different observed values, become more strongly weighted and the posterior mean biases towards the weighted global average. As a result, it is likely for the posterior mean in undersampled pre-edge smoothly varying regions to oscillate. We note, however, that this typically occurs when the measured energy range is very large and contains a substantial portion of smoothly varying regions (more than 100 eV); it is also more pronounced at the beginning of the experiment when only a few points are measured. One way to address this is to design a projection function to map the input features to another space that better describes the spatially varying length scale, in a way similar to DKL but with the projection function designed analytically. Additionally, the range of feature-rich regions in an XANES spectrum is typically known for a specific type of sample, so it is possible to confine the sampling in that region without including an exceedingly large portion of plateau. Even if the feature-rich range is not known, we may automatically determine that during an experiment by first running a series of coarse measurements across the energy range, and then estimating the absorption edge's location using the energy-domain gradient of the measured values. Following that, the range for sampling can be chosen to be a few tens or hundreds of eVs below and above the edge energy.

Meanwhile, while our current algorithm is designed to maximize the information gain from each measurement, we should acknowledge that in a real experiment, the overhead of moving the instrument—particularly the monochromator in the case of XANES—can be equally important, in view that a monochromator with the Bragg axis driven by a stepping motor can generate considerable overhead if directed to make frequent changes in its moving direction. An algorithmic way to solve this problem is to also incorporate a term quantifying the cost of moving the monochromator in this acquisition function, so that the suggested energy is one that gives a good balance between information gain and movement cost. The modified acquisition function may take the form of $a(x) = a(x)/c(x, x)$ ²¹, where $c(x, x)$ is the cost function taking into account both the energy being assessed and the energies measured in the past; the latter may be needed to calculate the additional cost due to the change of moving momenta or directions. An analytical model of the cost can introduce additional hyperparameters that

take extra efforts to tune, but it may also be constructed as a data-driven model that learns from experimentally collected movement-cost data and predicts the cost in an autoregressive manner³¹.

Beyond overcoming the current limitations, our future prospect also includes exploring the broader applications of features in the algorithm. Notably, the linear fitting of the reconstructed spectrum with reference spectra, used in the fitting residue component of the comprehensive acquisition function, is also a technique commonly used in the analysis of XANES spectra, and the fitted coefficients reveals the progress of a dynamic process. At the end of measuring each spectrum, this progress information can be computed and reported to the user (or a higher-level experiment control algorithm), so that the whole experiment can be stopped immediately when the process finishes, avoiding taking redundant spectra. Moreover, the fitting also indicates how well the current spectrum is represented by the references. If the fitting residue of certain regions remains high even after sufficient measurements, it most likely indicates that the reference spectra are improper for the sample being measured, and the user or the experiment control algorithm can be prompted to look for better reference spectra from the database. To summarize, beyond automating the collection of individual spectra, our algorithm could also become the core of a larger autonomous system that streamlines data collection and analysis.

To conclude, we have developed an adaptive sampling algorithm for XANES spectroscopy based on Bayesian optimization. Domain knowledge about XANES is incorporated into the algorithm through a designed acquisition function that makes the algorithm aware of regions of higher interest in the energy domain in real time. With this, the algorithm can direct experimental instruments to sample more densely at the absorption edge and other regions with fast varying features than areas that are more uniform. Through simulated studies and a real-world demonstration at the Advanced Photon Source, we show that our algorithm yields spectra as accurate as those sampled using traditional methods while only taking 10–30% of the measurements. The algorithm shortens the acquisition time of XANES, allowing for better time resolution in dynamic experiments, and avoids the over/undersampling problems in traditional XANES acquisition with pre-defined sampling grid. It signifies a step towards future x-ray experiment endstations that are more efficient, more autonomous, and more intelligent.

Methods

Gaussian process and Bayesian optimization

Standard Gaussian process (GP) models and Bayesian optimization (BO) methods have been extensively studied in the literature^{23,32}. We reiterate the fundamental theories here in the context of the adaptive sampling problem under concern. In a spectroscopy experiment like XANES, one measures a function $f(x)$ in the one-dimensional energy domain. For a series of n past measurements collectively denoted by \mathbf{f} at locations \mathbf{x} , the probability density function of \mathbf{f} is

$$p(\mathbf{f}) = \frac{1}{\sqrt{(2\pi)^n |\mathbf{K}|}} \exp \left[-\frac{1}{2} (\mathbf{f} - \boldsymbol{\mu})^T \mathbf{K}^{-1} (\mathbf{f} - \boldsymbol{\mu}) \right] \quad (1)$$

$$= \text{GP}(\mathbf{f}; \boldsymbol{\mu}, \mathbf{K}). \quad (2)$$

Eq. (2) is a GP describing the distribution of \mathbf{f} , where $\boldsymbol{\mu}$ is the expectation that often takes a simple form such as a constant. \mathbf{K} is the covariance matrix with each element K_{ij} representing the covariance between the observations at x_i and x_j . This covariance is estimated by the kernel function $k = k_{\theta}(x_i, x_j)$, where θ is the set of hyperparameters of the kernel. In this work, we use a Matérn kernel which takes the following form:

$$k_{\theta}(x_i, x_j) = \frac{1}{(\nu)2^{\nu-1}} \left(\frac{\sqrt{2\nu}}{l} |x_i - x_j| \right)^{\nu} K_{\nu} \left(\frac{\sqrt{2\nu}}{l} |x_i - x_j| \right) \quad (3)$$

where $K_{\nu}(\cdot)$ is a modified Bessel function and $\Gamma(\cdot)$ is the gamma function. We use $\nu = 2.5$ because it yields overall the best performance on the data

tested. l is the length scale and is the only tunable hyperparameter of the Matérn kernel, i.e., $\theta = \{l\}$.

We further assume that the measurement is noisy so the actual past observations \mathbf{y} follows $p(\mathbf{y}|\mathbf{f}) = \mathcal{N}(\mathbf{f}, \mathbf{N})$, where \mathcal{N} is a Gaussian distribution. For isotropic and independent noise, $\mathbf{N} = \sigma_o^2 \mathbf{I}$, where σ_o^2 is the noise variance and \mathbf{I} is the identity matrix; in a real experiment, σ_o^2 can be obtained through calibration. This allows us to formulate the log-likelihood function as

$$\begin{aligned} \log p(\mathbf{y}|\mathbf{x}, \theta) &= \log \int p(\mathbf{y}|\mathbf{f}, \mathbf{x}, \theta) p(\mathbf{f}|\mathbf{x}, \theta) d\mathbf{f} \\ &= -\frac{1}{2} [(\mathbf{y} - \boldsymbol{\mu})^T (\mathbf{K}(\theta) + \mathbf{N})^{-1} (\mathbf{y} - \boldsymbol{\mu}) + \log |\mathbf{K}(\theta) + \mathbf{N}| \\ &\quad + n \log(2\pi)]. \end{aligned} \quad (4)$$

By maximizing Eq. (4) with regards to θ , one can fit the kernel function's hyperparameter set when constructing a GP model using an initial set of observations and their locations. However, if one has good knowledge about the sample being measured, it is also possible to set strong priors or fixed values for the hyperparameters. This is applicable to the length scale parameter of the Matérn kernel, where we have found setting the length scale with a good estimate sometimes leads to better performance than fitting it on initially measured data when the initial points are scarce. A better alternative is to fit the hyperparameters using densely sampled reference spectra previously collected on similar materials, provided that they have comparable feature length scales as the spectra of the measured sample. The reference spectra are distinct from the spectrum being adaptively sampled, so using them for hyperparameter fitting does not violate the principle of train/test data isolation.

When a new measurement is made at point x^* , the joint distribution of the measured value f^* and previous measurements \mathbf{f} is

$$p(f^*, \mathbf{f}) = GP\left(\begin{bmatrix} f^* \\ \mathbf{f} \end{bmatrix}; \begin{bmatrix} \boldsymbol{\mu}^* \\ \boldsymbol{\mu} \end{bmatrix}, \begin{bmatrix} K^* & \boldsymbol{\kappa} \\ \boldsymbol{\kappa}^T & \mathbf{K} \end{bmatrix}\right) \quad (5)$$

where $K^* = k_\theta(x^*, x^*)$, $\boldsymbol{\kappa} = k_\theta(x^*, \mathbf{x})$.

Taking observation noise into account, the posterior predictive distribution of f^* at x^* is

$$p(f^*|\mathbf{y}) \propto \mathcal{N}[m(x^*), \sigma^2(x^*)] \quad (6)$$

with

$$m(x^*) = \boldsymbol{\mu} + \boldsymbol{\kappa}^T (\mathbf{K} + \mathbf{N})^{-1} (\mathbf{y} - \boldsymbol{\mu}) \quad (7)$$

and

$$\sigma^2(x^*) = k_\theta(x^*, x^*) + \boldsymbol{\kappa}^T (\mathbf{K} + \mathbf{N})^{-1} \boldsymbol{\kappa}. \quad (8)$$

Eq. (6)–(8) provides a value estimation (through the posterior mean $m(x^*)$) and uncertainty quantification (through the posterior variance $\sigma^2(x^*)$) for f^* . Based on these quantities, one may create an acquisition function $a(x)$ that evaluates the “worth” of taking a measurement at any point x . The formulation of $a(x)$ can be designed to reflect the objective of the experiment. For example, if the objective is to reduce the overall uncertainty of the data, then $a(x)$ could simply be the posterior standard deviation, i.e., $a(x) = \sigma(x)$.

With the acquisition function defined, one can guide subsequent data acquisition through a Bayesian optimization workflow: one finds the maximum point of $a(x)$, i.e., $x^* = \text{argmax}_x a(x)$, and take the next measurement at x^* . x^* and the measured value f^* are used to update \mathbf{y} , $\boldsymbol{\mu}$, $\boldsymbol{\kappa}$, and \mathbf{K} , yielding the new posterior predictive distribution parameters (Eq. (7) and (8)). This process is repeated until a stopping criterion is met.

In Eq. (7) and (8), the noise covariance \mathbf{N} regularizes the solution of the posterior mean and posterior variance, preventing them from overfitting to the noise in the observations. With $\mathbf{N} = \sigma_o^2 \mathbf{I}$, the noise covariance effectively

smooths the posterior mean, but we have observed that if the observation noise is very large, even when the set variance matches the actual observation noise's variance, it sometimes oversmooths sharp features in the posterior mean, causing important peaks such as the white line to be underestimated. Although this problem can be mitigated by sampling more around the peaks, we should note that since the adaptive sampling algorithm needs to capture all features across the entire spectrum, the number of samples allocated for each peak is still limited even though our method is designed to favor the sampling of peaks and fast-varying features. As such, we designed an optional feature of adaptive noise variance, where the noise variance associated with each observation is reduced linearly with the difference between the observed value and the minimum of the initial data (y_{floor}), i.e.,

$$n = n_0 - \frac{d}{y_{\text{cutoff}}} (1 - \xi) n_0 \quad (9)$$

where $d = \min[\max(y - y_{\text{floor}}, 0), y_{\text{cutoff}}]$. y_{cutoff} and ξ are the difference cutoff and decay factor in the standardized observation data space, which are adjustable but typically set to 1.5 and 0.01. This scheme preserves the noise regularization for low-absorption areas and relaxes the regularization for high-absorption features.

Acquisition function design for XANES adaptive sampling

From Eq. (8), one can find that the posterior variance is only a function of the *locations* of past measurements, but not their values; furthermore, a stationary kernel like Matérn is only dependent on the distance rather than absolute positions of the two input points, and so is the posterior variance. It is then obvious that simply using the posterior variance in the acquisition function $a(x)$ has the drawback of not being aware of the absolute positions and values of the measured or GP-estimated spectrum (i.e., its posterior mean). To allow the adaptive sampling algorithm to make better decisions in the points to measure, we inject our domain knowledge about XANES spectra by designing a acquisition function with the following components:

Gradient component. In an XANES spectrum, the most important features are the x-ray energy of the absorption edge, the pre-edge features and undulations around the absorption edge. The edge itself and the inclining and declining parts of wavy features can be identified by the large magnitude of their gradient with regards to energy. The peaks and valleys of near-edge undulations on the other hand are highlighted by their relatively large second-order gradient. Thus, we take both the magnitudes of the first- and second-order gradient of the posterior mean $m(x)$ in formulating the gradient component of the acquisition function:

$$a_g(x) = \phi_g \left\| \frac{dm(x)}{dx} \right\| + \phi_{g'} \left\| \frac{d^2 m(x)}{dx^2} \right\| \quad (10)$$

where ϕ_g and $\phi_{g'}$ are weight coefficients. The values of these coefficients can be scheduled to gradually reduce throughout the experiment, so that the algorithm focuses on regions with high first or second-order derivatives at first but slowly shifts towards exploration on uncertainty regions. We use an exponential decaying schedule for these coefficients. Each time the GP model is updated with new measurement, the weight coefficients are updated as

$$\phi_i \leftarrow \beta \phi_i \quad (i \in \{g, g'\}) \quad (11)$$

The decay constant β is heuristically set to be 0.999.

We also allow the initial values of ϕ_g and $\phi_{g'}$ to be estimated automatically as

$$\phi_i = \alpha_i \frac{\max \sigma(x)}{\max a_{g,i}(x, \phi_i = 1)} \quad (i \in \{g, g'\}) \quad (12)$$

where α_i is the importance relative to the posterior uncertainty for term i (first-order or second-order), and $a_{g,i}$ refers to the corresponding term in a_g .

We set both α_i 's to 0.5. Eq. (12) scales the acquisition function component's values so that their maxima are comparable with the posterior standard deviation.

Lastly, considering that the XANES spectra measured by a realistic instrument sometimes contain a background slope which may bias the first-order gradient, we detect the range of the pre-edge region using the method described in Section 4.2 and estimate the gradient in that region. This allows us to subtract the bias from the first-order gradient.

Fitting residue component. In dynamic XANES where multiple spectra of the sample are measured at different time points during a evolving process (e.g., phase transition), the XANES spectra of the equilibrium-state phases before and after the transition are often available as reference spectra. A spectrum measured during phase transition can be accurately described by the linear combination of the reference spectra¹⁹. Therefore, during a guided experiment, if one fits the posterior mean $m(x)$ (i.e., the densely reconstructed spectrum) with the reference spectra, a high fitting residue would indicate that values estimated for that local region might have deviated from the true values, and more points should be measured there. We thereby formulate the fitting residue component of the acquisition function as

$$a_r(x) = \phi_r \|\tilde{f}(x) - m(x)\|^2 \quad (13)$$

where $\tilde{f}(x)$ is spectrum linearly fitted to $m(x)$ using the reference spectra as bases. Like in the case of the gradient component, the weight coefficient ϕ_r can also be scheduled to decay exponentially. The initial values of ϕ_r may also be automatically determined using a method similar to the gradient component.

Acquisition function reweighting. Both the gradient and fitting residue components are calculated based on $m(x)$, which may deviate from the actual data. For example, a low-gradient region rendering a small a_g might in fact contain a small feature not detected yet. Therefore, we still need to keep the posterior standard deviation σ in the acquisition function to allow the algorithm to explore uncertain regions even if a_g and a_r are low there. This illicits another issue that the algorithm might allocate more-than-enough points in the smoothly varying pre-edge region driven by its high uncertainty. Hence, we create another modifier for $a(x)$ to reweight its value across the measured space.

Given the $m(x)$ computed after the GP model is updated with a few measurements, we find the location x_e and width u of the absorption edge in the posterior estimate by taking its finite-difference derivative with regards to energy, and identifying the largest peak of the derivative. This allows us to estimate the upper bound of the smoothly varying pre-edge region as $x_{pe} = x_e - c_e u$, where c_e is an adjustable but widely applicable constant that is set to be 1.6 in all tests performed. A shifted, scaled, and elevated sigmoid function is then constructed as

$$w_s(x) = \left(\frac{1}{1 + \exp\left[-\frac{3200}{u}(x - x_{pe})\right]} \right) (1 - w_{s,floor}) + w_{s,floor}. \quad (14)$$

The units of x , x_{pe} and u are eV. $w_{s,floor}$ is the floor or infimum value of the sigmoid, which is set to a small but non-zero number. When $w_{s,floor}$ is multiplied to the $a(x)$, it scales down the acquisition function's value in the pre-edge region by $w_{s,floor}$, but without completely zeroing it out. This allows the algorithm to sample more frequently in the edge and post-edge region while maintaining the possibility of uncertainty-driven exploration in the pre-edge region when the edge and post-edge acquisition function become low enough after sufficient amount of points are sampled. The factor of 3200 is determined empirically.

Optionally, one may also boost up the acquisition function in the region closely after the absorption edge, which contains most of the

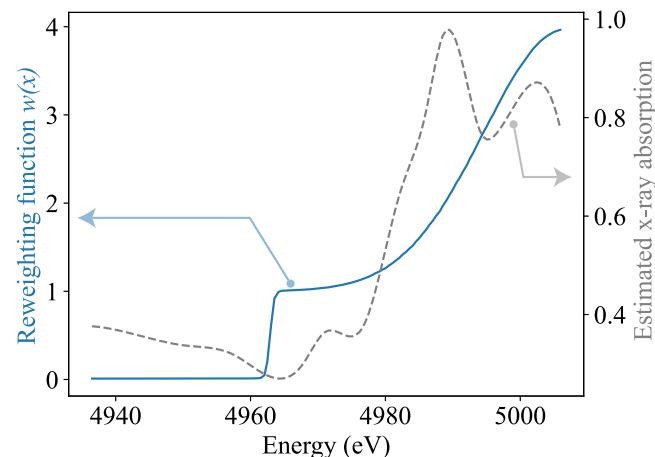


Fig. 8 | An example acquisition reweighting function. The reweighting function (blue, left axis) is plotted along with the x-ray absorption estimated through posterior mean (gray, right axis).

informative features. This is done by adding a Gaussian function of

$$w_g(x) = g \exp\left(-\frac{[x - (x_e + u)]^2}{0.5u^2}\right) \quad (15)$$

where g is a gain value kept as 3.0. Together, the acquisition reweighting function is

$$w(x) = w_s(x) + w_g(x). \quad (16)$$

An example reweighting function is shown together with the posterior mean used to generate it in Fig. 8.

The acquisition function is modified by this reweighting function as the blending between the original acquisition function and the weighted version of it. If the unmodified acquisition function is $a_0(x)$, then the modified function is given as

$$a(x) = \gamma w(x)a_0(x) + (1 - \gamma)a_0(x) \quad (17)$$

where γ is the blending coefficient which can also be scheduled to decay from the initial value of 1.0 so that the effect of reweighting gradually fades out, letting uncertainty play a more significant role at later stages. We use an exponential decay schedule similar to the weight coefficients of the acquisition function components and the decay coefficient is set to 0.95.

The heuristically determined values for the parameters involved in these parts of the algorithm, namely β , g , the decay coefficient of γ , and the parameters of $w(x)$, are widely applicable and normally do not need to be tuned. For the 4 experiments shown in this paper, these parameters are kept the same.

With the features introduced above, we create a comprehensive acquisition function that is value-aware and incorporates our domain knowledge about XANES spectra. The acquisition function is formulated as the following:

$$a(x) = w(x)\sigma(x) \max(a_g + a_r, \tau) \quad (18)$$

where τ is a lower bound that is set to be 0.03 throughout this work unless otherwise mentioned. The posterior variance $\sigma(x)$ and domain knowledge components ($a_g(x)$ and $a_r(x)$) are multiplied instead of added. This is based on the fact that when a point is measured, the $\sigma(x)$ in its close proximity drops towards 0. We want this behavior to control the entire acquisition function, so that the algorithm will not make repeated measurement in a neighborhood where enough points have already been measured. Also, the sum of domain knowledge components $a_g + a_r$ are constrained with a lower

bound τ , so that a seemingly smoothly varying or low-residue region does not totally zero out the acquisition function, and the algorithm preserves the possibility to take measurements there driven by uncertainty. These two designs prevent the algorithm from either over-sampling in a region or missing regions containing undetected features.

Other design features

The acquisition can be stopped when the maximum posterior standard deviation weighted by the sigmoid term of the reweighting function drops below a threshold, i.e.,

$$\max [\sigma(x)w_s(x)] < t. \quad (19)$$

Many BO algorithms find the maxima of the acquisition function through numerical optimization. However, since the input space of XANES is one-dimensional, it is in fact more convenient and effective to numerically find the maximum of the acquisition function by evaluating it on a fine grid of points, and then simply taking the maximum.

Baseline acquisition functions

We compare our method with a series of baselines that include three Bayesian optimization methods employing different acquisition functions: posterior standard deviation, upper confidence bound (UCB), and expected improvement (EI).

The posterior standard deviation acquisition function simply uses $a(x) = \sigma(x)$, where $\sigma(x)$ is the posterior standard deviation coming from Eq. (8). It is an exploration-only acquisition function that is purely uncertainty-driven.

UCB is expressed as

$$a_{UCB}(x) = m(x) + \kappa\sigma(x) \quad (20)$$

where $m(x)$ is the posterior mean and κ is a trade-off factor. It is designed for maximization problems, where the $m(x)$ term drives the algorithm to sample where the estimated function is high; however, it can also explore areas where the estimated function is not the highest but has large uncertainty. While we should note that maximization problems are distinct from adaptive sampling problems, the $m(x)$ could still be useful for XANES sampling as many important features have relatively high absorption. As such, we include UCB as a baseline.

EI is formulated as

$$a_{EI}(x) = \int \max(y - \hat{y}, 0) \mathcal{N}[y; m(x), \sigma^2(x)] dy \quad (21)$$

where \hat{y} is the highest function value observed so far. It gives the expectation of the increase of y compared to \hat{y} at a certain x , given that the probability density function of y at that point is a Gaussian distribution parameterized by the GP's posterior mean and standard deviation. EI is widely used in solving maximization problems, but its performance in adaptive sampling problems is yet to be seen.

Experimental data collection

Our algorithm is tested with data collected on 4 samples, namely a superconductor material yttrium barium copper oxide (YBCO), a battery material lithium titanium oxide (LTO), a catalytic material with Pt particles deposited on $\gamma\text{-Al}_2\text{O}_3$ substrate, and finally a battery material lithium nickel manganese cobalt oxide (NMC111). The demonstrations of our algorithm on the first three samples are simulated, meaning that ground truth spectra have been collected using a dense traditional sampling grid beforehand, and the algorithm runs with data interpolated from the truths. The test on the NMC111 sample was done in the real world, with the algorithm communicating bidirectionally with the monochromator and photon detector. This section documents the experimental collection of the data.

The $\text{Li}_4\text{Ti}_7\text{O}_{12}$ (LTO) samples were commercially purchased from EnerDel. The preparation of orientation-dependent $\text{YBa}_2\text{Cu}_3\text{O}_7$ (YBCO)

samples involves synthesizing YBCO powder and orienting the powders with a high magnetic field, the details of which can be found in³³.

The XANES spectra of both YBCO and LTO were measured at the 20-BM beamline³⁴ of the Advanced Photon Source (APS). In both experiments, a harmonic rejection mirror and a 10-15% detuning of x-ray intensity using the 2nd crystal of the Si (111) monochromator were applied to avoid high-energy x-ray harmonics. These measurements were performed in transmission mode. Two LTO samples were respectively measured at 2 different temperatures, 50°C and 70°C. At each temperature, multiple XANES spectra were collected throughout the sample's phase transition process. The Linkam THMS600 heating stage was used for temperature-dependent measurements.

For the data collection of the $\text{Pt}/\gamma\text{-Al}_2\text{O}_3$ material, a sample was prepared following the procedures in ref. 25. XANES spectra were continuously measured on the sample while it undergoes reduction in hydrogen atmosphere and the reduction conditions are also documented in the same paper.

The Pt on alumina catalyst was prepared by impregnation of platinum-chloride precursor followed by calcination at 525 °C in air²⁵. The powder sample was ground and pressed into a metal sample holder for the transmission measurements at the MRCAT³⁵ at Advanced Photon Source. The x-ray beam size on the sample was approximately 1 mm by 1 mm as defined by slits. The double crystal Si(111) monochromator was scanned to select the x-ray energy and a Rh-coated mirror was used to remove x-rays with higher harmonic energies. The X-ray energy was calibrated by using the Pt absorption edge of Pt foil. The custom designed in-situ cell is described elsewhere³⁶. The sample was heated from room temperature to 500°C in 100% H_2 , resulting in continuous reduction of the Pt oxidation state. In-situ Pt XANES measurements were recorded in transmission during this process.

For the NMC111 sample, an NMC111 cathode laminate was prepared as 90 % (w/w) $\text{LiNi}_{1/3}\text{Mn}_{1/3}\text{Co}_{1/3}\text{O}_2$ (BASF TODA Battery Materials LLC NM-3101), 5 % (w/w) Timcal C-45 carbon, and 5 % (w/w) Solvay 5130 PVDF Binder. The slurry was coated onto a 20 μm Al foil current collector for a combined thickness of 61 μm . 12.7 mm diameter electrodes were punched from the dried laminate for cycling. For in-situ measurements, the electrode was placed in an AMPiX cell³⁷ with a Li metal anode and 1.0 mol L^{-1} LiPF_6 in ethylene carbonate/dimethyl carbonate 50/50 (v/v) electrolyte.

The data collection with the NMC111 sample was done at the 25-ID-C beamline of the Advanced Photon Source, where a KB mirror pair is used for focusing x-ray on the sample and a Si(111) monochromator is used for energy selection. The monochromator is driven by a step motor without air-bearing. Our algorithm was integrated with beamline instruments such that it could receive the measured absorption value at a certain energy, and steers the monochromator to the next energy that it suggests to measure at. More details about this control flow is introduced in section 4.6.

During the experiment, the battery cell was charged then discharged galvanostatically at 0.205 mA (0.1 h^{-1}) between 3.5 V and 4.5 V using a Maccor 4300. 8 XANES spectra were collected sequentially with equal time interval during the process under the guidance of our algorithm. Additionally, 2 reference spectra were measured on the uncharged (pristine) and fully charged (4.5 V) samples. The reference spectra were measured over a range of ± 100 eV relative to the K absorption edge of Ni (8333 eV), or 8233 eV to 8433 eV, with a uniform step size of 0.5 eV. On the other hand, adaptive sampling was done on a larger range of -100 eV to 150 eV relative to the Ni edge.

Implementation of beamline control

The operation of our algorithm at 25-ID-C uses Bluesky²⁸ to establish the communication and control flow. The software system involves the following components:

- The backend containing the described algorithm, exposing methods update and suggest for updating algorithm internal states with new measurements and outputting new energy suggestions, respectively;
- A recommender class (XANESSamplingRecommender) exposing a tell method to receive measured absorption and pass it to the

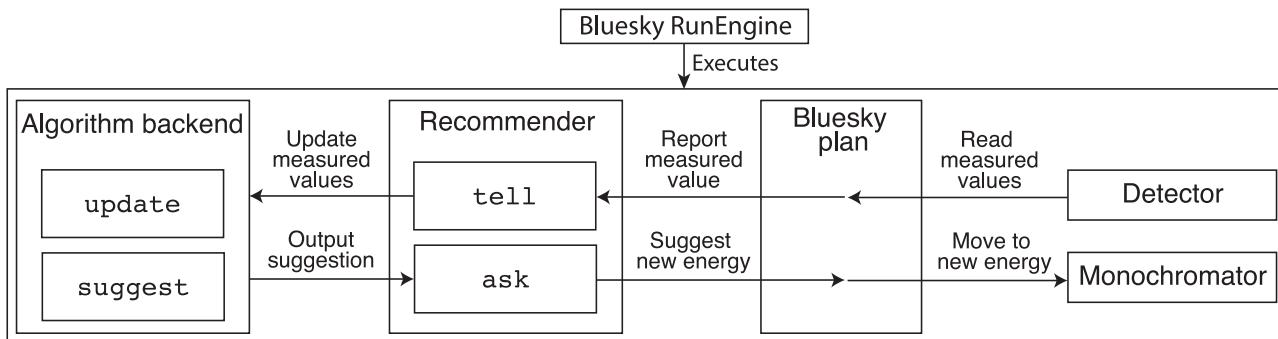


Fig. 9 | The control and communication flow between the backend adaptive sampling algorithm and beamline instruments.

backend, and an ask method to query the algorithm for the next suggested energy point and replay it to the monochromator;

- Bluesky Ophyd³⁸Device objects that symbolize the monochromator and detector;
- A Bluesky plan streamlining the process using the Device objects and the recommender;
- A high-level script that moves the sample stage to the right location, controls the shutters, and executes the Bluesky plan using a Bluesky RunEngine instance.

Figure 9 illustrates the interactions among these components. The experiment workflow defined in the Bluesky plan (and as shown in Fig. 1) is conceptualized in the following pseudocode:

```

IO: Device # Ophyd Device object of reference
intensity
It: Device # Ophyd Device object of test intensity
energyPositioner: Device # Ophyd Device object of
monochromator
recommender: XANESSamplingRecommender stopping-
ConditionTriggered: bool = False
measureInitialPointsAndInitializeBackend()
# In the actual implementation, the following loop
is executed
# by passing the recommender to the
# bluesky_adaptive.per_event.adaptive_plan
function, yet
# we expand it here.
while True:
    try:
        suggestedEnergy = recommender.ask(n=1)
    except NoRecommendation:
        # Stopping condition triggered
        break
    energyPositioner.move(suggestedEnergy)
    IOVal = IO.read()
    ItVal = It.read()
    recommender.tell(suggestedEnergy, -log(ItVal
    / IOVal))
  
```

Data availability

The data that supports the findings of this study is available on U.S. DOE OSTI at <https://doi.org/10.17038/XSD/2583950>.

Code availability

The code that implements the proposed algorithm is available at https://github.com/AdvancedPhotonSource/auto_beamline_ops.

Received: 18 February 2025; Accepted: 7 August 2025;

Published online: 27 October 2025

References

1. Silva, C. L. et al. On the origin of low-valent uranium oxidation state. *Nat. Commun.* **15**, 6861 (2024).
2. Ade, H. et al. Chemical contrast in X-ray microscopy and spatially resolved XANES spectroscopy of organic specimens. *Science* **258**, 972–975 (1992).
3. Wan, G. et al. Solvent-mediated oxide hydrogenation in layered cathodes. *Science* **385**, 1230–1236 (2024).
4. Yang, T. et al. Ultrahigh-nickel layered cathode with cycling stability for sustainable lithium-ion batteries. *Nat. Sustain.* (2024).
5. Zabinsky, S., Rehr, J., Ankudinov, A., Albers, R. & Eller, M. Multiple-scattering calculations of x-ray-absorption spectra. *Physical Review B* **52**, 2995 (1995).
6. Jacobsen, C. *X-ray Microscopy. Advances in Microscopy and Microanalysis* (Cambridge University Press, Cambridge, 2019).
7. Meyer, R. J. et al. Recommendations to standardize reporting, execution, and interpretation of x-ray absorption spectroscopy measurements. *Journal of Catalysis* **432**, 115369 (2024).
8. Pascarelli, S. et al. Materials under extreme conditions using large X-ray facilities. *Nat. Rev. Methods Primers* **3**, 82 (2023).
9. Betterton, J.-R., Ratner, D., Webb, S. & Kochenderfer, M. Reinforcement learning for adaptive illumination with x-rays. *2020 IEEE International Conference on Robotics and Automation (ICRA)* **00**, 328–334 (2020).
10. Zhang, Y. et al. SLADS-Net: Supervised Learning Approach for Dynamic Sampling using Deep Neural Networks. *Electronic Imaging* **2018**, 131–1–1316 (2018).
11. Kandel, S. et al. Demonstration of an ai-driven workflow for autonomous high-resolution scanning microscopy. *Nature Communications* **14**, 5501 (2023).
12. Noack, M. M. et al. Gaussian processes for autonomous data acquisition at large-scale synchrotron and neutron facilities. *Nature Reviews Physics* **3**, 685–697 (2021).
13. Noack, M. M. et al. A Kriging-Based Approach to Autonomous Experimentation with Applications to X-Ray Scattering. *Scientific Reports* **9**, 11809 (2019).
14. Holman, E. A. et al. Autonomous adaptive data acquisition for scanning hyperspectral imaging. *Communications Biology* **3**, 684 (2020).
15. Ueno, T. et al. Adaptive design of an x-ray magnetic circular dichroism spectroscopy experiment with gaussian process modelling. *npj Computational Materials* **4**, 4 (2018).
16. Zhang, Y. et al. Autonomous atomic hamiltonian construction and active sampling of x-ray absorption spectroscopy by adversarial bayesian optimization. *npj Computational Materials* **9**, 46 (2023).
17. Cakir, C. T. et al. Machine learning for efficient grazing-exit x-ray absorption near edge structure spectroscopy analysis: Bayesian optimization approach. *Machine Learning: Science and Technology* **5**, 025037 (2024).

18. Matérn, B. *Spatial Variation*. Lecture Notes in Statistics (Springer, New York, 1986), 2nd edition edn.
19. Kelly, S. D., Hesterberg, D. & Ravel, B. Analysis of soils and minerals using x-ray absorption spectroscopy. In *Methods of Soil Analysis Part 5—Mineralogical Methods*, 387–463 (American Society of Agronomy and Soil Science Society of America, Madison, WI, USA, 2008).
20. Calvin, S. *XAFS for Everyone* (CRC press, 2024).
21. Noack, M. M., Doerk, G. S., Li, R., Fukuto, M. & Yager, K. G. Advances in Kriging-Based Autonomous X-Ray Scattering Experiments. *Scientific Reports* **10**, 1325 (2020).
22. Ravel, B. & Newville, M. ATHENA and ARTEMIS. In *International Tables for Crystallography*, vol. 1 (International Union of Crystallography, Chester, England, 2020).
23. Garnett, R. *Bayesian Optimization* (Cambridge University Press, 2023).
24. Godaliyadda, G. M. D. P. et al. A framework for dynamic image sampling based on supervised learning. *IEEE Trans. Comput. Imaging* **4**, 1–16 (2018).
25. Kelly, S. et al. Advanced characterization for industrial catalysis applications. *Chin. J. Catalysis* **40**, 1637–1654 (2019).
26. Chen, J., Finfrock, Y. Z., Wang, Z. & Sham, T.-K. High energy resolution fluorescence detection of the Pt L_{3,2}-edge whitelines of Pt-based bimetallic systems: Implications for the Pt 5d_{5/2,3/2} density of states. *J. Phys. Chem. C* **125**, 2327–2333 (2021).
27. Matsushita, T. & Hashizume, H. *Handbook on Synchrotron Radiation*. No. vol. 1 in *Handbook on Synchrotron Radiation* (North Holland Publishing Company, 1983). <https://books.google.com/books?id=xF1SrgEACAAJ>.
28. Allan, D., Caswell, T., Campbell, S. & Rakitin, M. Bluesky's ahead: A multi-facility collaboration for an a la carte software project for data acquisition and management. *Synchrotron Radiation News* **32**, 19–22 (2019).
29. Wilson, A. G., Hu, Z., Salakhutdinov, R. & Xing, E. P. Deep Kernel Learning. *arXiv* 1511.02222 (2015).
30. Müller, O., Nachtegaal, M., Just, J., Lützenkirchen-Hecht, D. & Frahm, R. Quick-EXAFS setup at the SuperXAS beamline for *in situ* X-ray absorption spectroscopy with 10ms time resolution. *Journal of Synchrotron Radiation* **23**, 260–266 (2016).
31. Box, G. E., Jenkins, G. M., Reinsel, G. C. & Ljung, G. M. *Time series analysis: forecasting and control* (John Wiley & Sons, 2015).
32. Berger, J. O. *Statistical Decision Theory and Bayesian Analysis*. Springer Series in Statistics (Springer, New York, 1985).
33. Heald, S. M., Tranquada, J. M., Moodenbaugh, A. R. & Xu, Y. Orientation-dependent x-ray-absorption near-edge studies of high-*T_c* superconductors. *Phys. Rev. B Condens. Matter* **38**, 761–764 (1988).
34. Heald, S. et al. Xafs and micro-xafs at the pnc-cat beamlines. *Journal of Synchrotron Radiation* **6**, 347–349 (1999).
35. Segre, C. U. et al. The MRCAT insertion device beamline at the Advanced Photon Source. *AIP Conference Proceedings* **521**, 419–422 (2000).
36. Bare, S. R., Michelson, G. E., Modica, F. S., Ringwelski, A. Z. & Yang, N. Simple flow-through reaction cells for *in situ* transmission and fluorescence x-ray absorption spectroscopy of heterogeneous catalysts. *Rev. Sci. Instrum.* **77**, 023105 (2006).
37. Borkiewicz, O. J. et al. The AMPiX electrochemical cell: a versatile apparatus for *in situ* X-ray scattering and spectroscopic measurements. *Journal of Applied Crystallography* **45**, 1261–1269 (2012).
38. Ophyd contributors. Ophyd. <https://github.com/bluesky/ophyd>.

Acknowledgements

This research was performed on APS beam time award (<https://doi.org/10.46936/APS-187700/60013005>) from the Advanced Photon Source, a U.S. Department of Energy (DOE) Office of Science user facility at Argonne National Laboratory, and is based on research supported by the U.S. DOE Office of Science-Basic Energy Sciences, under Contract No. DE-AC02-06CH11357. The research was also supported by the Canadian Light Source and its funding partners. Data for demonstrations were collected at beamlines 25-ID, 20-BM, and 10-ID (MRCAT) of the Advanced Photon Source. The NMC-111 electrodes were supplied by the U.S. Department of Energy's (DOE) CAMP (Cell Analysis, Modeling and Prototyping) Facility, Argonne National Laboratory. The CAMP Facility is fully supported by the DOE Vehicle Technologies Office (VTO).

Author contributions

M.D., C.S., S.K. and M.C. proposed the research and discussed through the course of the method's development. M.D. developed the code that implements the method, conducted computational experiments, and analyzed the data. C.S. contributed the LTO and YBCO data. S.K. contributed the Pt data. M.W. prepared the NMC111 sample and conducted beamline experiment with M.D.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41524-025-01771-7>.

Correspondence and requests for materials should be addressed to Ming Du, Shelly D. Kelly or Mathew J. Cherukara.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© UChicago Argonne, LLC, Operator of Argonne National Laboratory under exclusive licence to Springer Nature Limited 2025