Machine Learning for analysis of speckle dynamics: quantification and outlier detection

Tatiana Konstantinova, Lutz Wiegart, Maksim Rakitin, Anthony M. DeGennaro,* and Andi M. Barbour†
Brookhaven National Laboratory Upton, New York 11973, USA

(Received 10 May 2022; accepted 1 August 2022; published 21 September 2022)

X-ray photon correlation spectroscopy (XPCS) provides an understanding of complex dynamics in materials that are tied to their synthesis, properties, and behaviors. Analysis of XPCS data for dynamics that are far from equilibrium is labor intense and often can impede the discovery process, especially in experiments with high collection rates. Moreover, binning and averaging, involved in the analysis for alleviating poor signal-to-noise ratio, leads to a loss of temporal resolution and the accumulation of systematic error for the parameters quantifying the dynamics. Here, we integrate a denoising autoencoder model into workflows for the analysis of nonequilibrium two-time intensity-intensity correlation functions. Noise reduction allows for extracting the parameters that characterize the sample's dynamics with the temporal resolution limited only by frame rates. Not only does it improve the quantitative usage of the data, but it also creates the potential for automating the analytical workflow, which is a key to high-throughput or autonomous XPCS experiments. Various approaches for the uncertainty quantification and extension of the model for anomalies' detection are discussed.

DOI: 10.1103/PhysRevResearch.4.033228

I. INTRODUCTION

Technological advances, such as high-brightness x-ray photon sources [1–3] and high-rate high-sensitivity detectors [4,5], enable new discoveries [6] by means of x-ray photon correlation spectroscopy (XPCS) [7–9]. XPCS experiments provide insights into manufacturing processes, such as crystal growth [10–13], curing reactions in polymer systems [14–16], and three-dimensional (3D) printing [17–20], which are the key to targeted, controlled material design. Such processes typically follow complex far-from-equilibrium dynamics. Additionally, nonequilibrium dynamics are a signature of many functional systems studied with XPCS, including superconductors [21], proteins [22], alloys [23], metallic glasses [24], charge-density-wave materials [25], and strained filled elastomers [26].

Such dynamics are commonly represented via two-time intensity-intensity correlation functions ($C2$s) [27], defined by the expression

$$C2(q;t_1,t_2) = \frac{\langle I(q;t_1)I(q;t_2)\rangle}{\langle I(q;t_1)\rangle\langle I(q;t_2)\rangle}, \tag{1}$$

where $I(q;t)$ is the intensity of a detector pixel corresponding to the wave vector q at time t. The average is taken over pixels with equivalent q values. The quantitative analysis of a nonequilibrium $C2$ [Fig. 1(a)] often starts with its binning along the sample's age axis $t_a = (t_1 + t_2)/2$ into quasiequilibrium regions [28–31] prior to making cuts along the time delay axis $t_d = |t_1 - t_2|$ and averaging them, i.e., obtaining aged one-time intensity-intensity correlation functions ($C1$s) [32] of the delay time

$$C1(q;t_a,t_d) = \langle C2(q;t_a,t_d)\rangle_{t_a}. \tag{2}$$

The reason for binning and averaging is the reduction of noise for subsequent analysis. The resulting $C1$s are fit to a functional form that is characteristic for the dynamics under investigation.

The traditional analysis of experimental data with complex nonequilibrium dynamics presents several challenges. The selection of quasiequilibrium regions is generally a notrivial and time-consuming task, which requires interventions from a researcher for visual inspection of the calculated $C2$ and of intermediate fitting results. With increasing data collection rates [34,35], it often becomes unfeasible to properly do on-the-fly data analysis during XPCS experiments when such an analysis relies on manual inputs. This can lead to inefficient use of beam facilities and misdirect the efforts of scientific staff. The sheer volume of collected data can reach hundreds of thousands of $C2$s a day, making even the offline data processing a demanding task. The analysis of nonequilibrium $C2$s becomes a bottleneck towards understanding the material dynamics and thus limiting scientific discoveries. Moreover, the binning procedure deteriorates the temporal resolution along the t_a-direction. Automating the analytical routines is an inevitable step for conducting measurements at modern light source facilities [34], including high-throughput autonomous XPCS experiments. The ability to perform the analysis with experimental frame-rate resolution is the ultimate target.

Since the motivation behind binning and averaging the data is to curtail the impact of noise, preprocessing a $C2$

*adegennaro@bnl.gov
†abarbour@bnl.gov

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

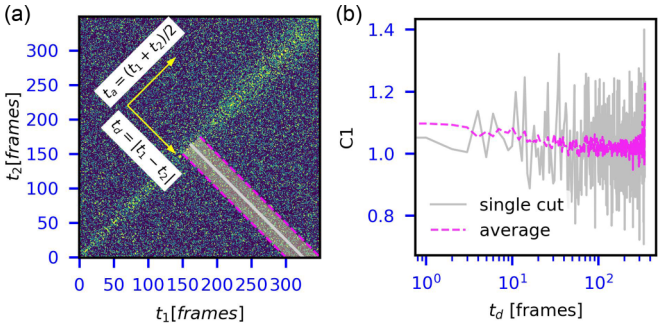

FIG. 1. Examples of intensity-intensity correlation functions. (a) A nonequilibrium $C2$. t_1 and t_2 are experiment time measured in frames. Yellow arrows indicate the directions of age axis t_a and time delay axis t_d, defined according to the formulas in the inset. Pink dashed lines and the grey solid line represent the edges and the bin center, respectively, for the region used for generating the $C1$ in (b). (b) $C1$s for a single cut along the bin center (grey solid) and average within the bin of 50 frames (pink dashed). Here and henceforth, MATPLOTLIB library [33] is used for plotting.

with a noise-reducing algorithm would increase the chance of achieving a frame-rate resolution along t_a during quantitative analysis. Previously, we demonstrated an approach for noise reduction in $C2$s with equilibrium dynamics, which is based on a convolutional neural network encoder-decoder model (CNN–ED) [36]. The method is shown to work well for equilibrium data and has a promising potential for use with slowly ageing dynamics when applied as a sliding window along t_a. However, there are two major concerns regarding the use of the model for general cases of nonequilibrium dynamics. First, if the timescale of the dynamics changes fast, an equilibrium approximation may not be valid, even within narrow age regions. Thus, the application of the model would lead to loss of experimental resolution and the extraction of less accurate results, especially for the cases of automated analysis where the functional form of the dynamics is not known *a priori*. Second, for a $C2$ containing a much larger number of frames than the input size of the model, a considerable portion of it would still contain the original level of noise after the model is applied along t_a.

Here, we build upon the knowledge gained from the previous model to address the reduction of noise in arbitrary-sized $C2$s with nonequilibrium dynamics such as ageing. We further demonstrate how the new model, integrated into a workflow for quantitative analysis of XPCS data, eliminates the need of age binning and allows wider parameter bounds for the fit of $C1$s. This, in turn, enables the sought-after automation of the analysis. The results of such an analysis have a temporal resolution, which is limited only by the experimental acquisition rates. The methods for estimating the credibility of the results—uncertainty scoring for denoised $C2$s and *trust regions* for dynamics' parameters — are discussed. Various options for analysis workflows are demonstrated for several XPCS experiments, which show applicability beyond quantitative analysis of common dynamics. An example of analysis workflow is shared in form of a JUPYTER [37] notebook in the accompanying git repository [38].

II. METHODS

A. Data

For a denoising model to perform well for nonequilibrium $C2$s, data from experiments with these types of dynamics are used for training. The data are collected from 65 XPCS experiments at the CHX beamline of NSLS-II from three different samples, recorded at different conditions and with different detector collection rates. Some of the measurements are repeated multiple times at the same conditions. The samples generally exhibit dynamics, common for many materials, that at various rates monotonically decelerate with age or stay at quasiequilibrium. For all considered dynamics, individual $C1$s can be approximated with the Kohlrausch-Williams-Watts (KWW) form [39]

$$C1(t_d) = C_\infty + \beta e^{-2(\Gamma t_d)^\alpha}, \tag{3}$$

where Γ is the rate of the dynamics, β is the dynamic contrast factor, α is the compression constant, and C_∞ is the baseline. All parameters are functions of t_a and q. Several regions of interest in the reciprocal space are considered for each experiment when calculating $C2$s. This results in a total of 492 $C2$s, ranging from 134 to 2950 frames. The original full-sized $C2$s are split between training (454) and validation (38) sets prior to generating 100×100 frames inputs to the model and augmenting the data. The training set is used for obtaining model weights and the validation set is used for model optimization, i.e., parameters tuning, and the evaluation of its performance as described in the Supplemental Material (SM) [40]. The augmentation of data includes systematic frames subsampling and the inversion of the $C2$ matrix to artificially diversify the distribution of considered XPCS experiments (see SM [40]).

Each input (100×100 frames) is scaled to have zero mean and unit variance prior to passage through the model. A reverse transformation is performed for the model output.

B. Denoising autoecoder model

The CNN–ED model architecture [36], used for equilibrium data, demonstrates several advantages, such as simplicity, control of overfitting, as well as fast training and application. Trying the same model architecture with certain adjustments for nonequilibrium data is a natural choice. The model presented here consists of an encoder with two 10-channel convolutional [41] layers that compress the input into eight encoded variables (the latent vector) and the decoder with two ten-channel transpose-convolutional layers. Low dimensionality of the latent vector helps to only pass essential information through the model, suppressing the noise. Two modifications are implemented during the model training to meet the peculiarities of nonequilibrium dynamics. First, since a "noise-free" version of an input cannot be obtained by averaging multiple cropped inputs collected from the same full-sized $C2$, the model is trained in an autoencoder mode [42] with an input and its target being the same. Second, a nonequilibrium $C2$ cannot be described by a single $C1$ and thus only the mean square loss between the output and the target is chosen as the training loss function. Due to the loss

function being applied directly to the model output, there is no clear interpretation of the latent space components. Applying additional constraints to the latent vector would produce interpretable [43] parameters. Implementing such constraints would require a fixed set of parameters that describe all possible dynamics passed to the model and an expensive manual labeling of the training/testing data, i.e., fitting of all training and validation inputs to a certain functional form. However, an interpretable latent vector is not required for the applications considered in this work and thus is not implemented here.

Despite the increased complexity of the dynamics in the training data, the architecture still appears appropriate for the denoising task based on the model's good performance for the validation set (see SM [40]). The size of the convolutional kernel reflects the local receptive field of the model and thus it is expected that it would affect the model performance. The effectiveness of the denoising autoencoder models (DAE) with convolutional kernel sizes from 1 to 17 is tested and only minor improvements in the smoothness of the output and the level of finite details preserved for larger kernels are discovered. The models with larger kernels take a longer time to train and the computational time during their application also becomes significant. In the following, unless stated otherwise, a model with the kernel size of 1 is used. The fact that the same model architecture works well for both equilibrium and nonequilibrium ageing $C2$s confirms its robustness for different types of dynamics.

Since the model has a fully connected layer (the latent vector), it can only be applied to a fixed-size input, i.e., 100×100 frames. However, the dimensions of an experimental $C2$ can vary considerably. Application of the model in a sliding window fashion along the age axis t_a can help with the noise reduction in certain cases like in our original quasiequilibrium analysis [36]. Nevertheless, this approach is not suitable for general cases of $C2$s, where the speckle signals decorrelate outside of the first 100 delay frames. The employed solution for the model application to inputs of arbitrary size (*down-and-up mapping*) is described in detail in SM [40] so it may be adopted by others.

C. Uncertainty quantification for denoising

When applying a model to unseen data, there are three potential sources of uncertainty: error (noise) of the input, model's bias, and model's variance. The bias of the model primarily originates from selection of the training dataset and the model's architecture. Variance appears due to random initializations of the model's weights and the ordering of data batches during the model's training. Approaches for quantifying uncertainty for a deep learning model often involve random perturbations of the model during training and/or application and obtaining the distribution of respective outputs [44].

We estimate the variance of the DAE by considering deviations of outputs for models trained with different random initializations (see SM [40]). It appears that the model's variance is generally much smaller than the corresponding values of β, thereby its influence on the dynamics quantification is insignificant. Moreover, the fluctuations of values in the neighboring points of a $C2_{\text{denoised}}$, caused by *down-and-up mapping* during the model application, scale linearly with the inherent variance of the model. Thus, it is not necessary to separately estimate the model's variance for experiments with more than 100 frames, as the point-to-point variations in a $C2_{\text{denoised}}$ already influence the quantitative analysis of its $C1$s.

From a practical point of view, the main potential source of uncertainty of the model is its bias. Obviously, a DAE must exhibit *some* bias to remove noise. Hence, to drive the decision about applicability of the model in each case, the quantification of the bias ought to answer the following question: How certain can one be that the model's output is a valid representation of the underlying sample's dynamics in the input? Naturally, a model returns the most trustworthy results for the inputs that are very similar to the examples in the training set. We suggest two quantitative measures — uncertainty scores — for estimating the bias by comparing the new input to the training examples.

The first measure is based on the Euclidean distance between the latent vector representation of an input and the center of the latent representations of all training examples. The analysis of the pairwise distributions of the eith latent coordinates of the training set reveals that they form a single compact cluster (see SM [40]). To ensure an equal contribution of each latent dimension to the uncertainty score, all coordinates are scaled to have a zero mean and unit variance across the training dataset. The shape of the cluster is determined by the distribution of dynamics and the noise level in the training set. The distance between any point and the cluster center is indicative of how different a given input is from most of the training data and thus can be used as an uncertainty score. To improve the interpretability of the bias, the distance to the center is normalized by its median value among the training set examples. The normalization of the distance helps to set a general threshold for an acceptable level of the bias without reference to a particular model or the training set.

The second measure of the denoising bias is less abstract and does not depend on the model architecture. It involves the evaluation of trends in residuals, i.e., the differences between the model's outputs and the corresponding targets. When an output is a valid representation of the sample's dynamics, the residual should resemble random and sometimes correlated noise in the form of vertical and horizontal stripes, without trends along either t_a or t_d. Since the training examples mostly represent monotonically ageing dynamics, the model tends to perform well for similar cases. A trend in a residual can indicate a heterogeneity on top of the ageing dynamics or a completely different type of dynamics. Such situations require additional attention during analysis and are discussed in more details in Sec. III A.

To identify the trends, it is convenient to look at the projection of a two-dimensional residual to the vertical (or horizontal) direction. Without trends, the autocorrelation coefficients (ACCs) of this projection are close to zero. In contrast, when a trend is present, the absolute value of the ACCs grows. We calculate the first ACCs for all the examples of the training set and approximate the probability density distribution using a Gaussian kernel density estimator (see SM [40]). The probability density for the first ACC of the residual for a test example is the second type of uncertainty scoring of the model prediction.

Since both bias measures are based on the distribution of the training examples, they are suitable for detecting anomalies — inputs that are different from the training set. The analysis workflow can flag a new $C2$ as an anomaly if its uncertainty score is above the user-defined threshold.

D. Extracting dynamics parameters

A conventional quantitative analysis of XPCS results with nonequilibrium dynamics, such as ageing, involves taking time-slices along the delay axis t_d of a $C2$ at different sample's ages t_a and fitting the resulting $C1$s with an appropriate model that describes the dynamics, which in many cases is the KWW form [Eq. (3)]. Typically, several adjacent slices within age range Δt_a need to be averaged to target a signal-to-noise ratio that is sufficient for extracting dynamics parameters with reasonable certainty. It is necessary that the dynamics do not change much within Δt_a. An example of selecting a quasiequilibrium bin within a $C2$ is shown in Fig. 1. This approach leads to a loss of temporal resolution for the parameters. Moreover, the procedure of binning $C1$s often requires a repetitive evaluation of the fitting results and adjustment of bins' boundaries by an expert researcher. For a large volume of collected data, it becomes challenging to properly split $C2$s for each experimental region of interest (ROI), leading to a growing uncertainty of the results. The ability to perform an adequate quantitative analysis for a single value of Δt_a across the entire experiment would facilitate the automation of the analysis process. The highest temporal resolution and maximum usage of the experimental data are achieved when considering cuts with $\Delta t_a = 1 \frac{1}{\text{frame rate}}$ (or, fr^{-1}).

An alternative approach to the traditional analysis could be a "global" fit of $C2$ with Eq. (3), parameterized as functions of t_a [$F_p(t_a)$, $p = \Gamma, \beta, \alpha, C_\infty$]. This method reduced the number of parameters per $C2(t_1, t_2)$ point. However, the selection of such functions would require a prior knowledge of the system behavior, which rarely can be available for online data analysis during the experiments. Moreover, nonlinear regression is very sensitive to the initial parameter values, parameter limits, and outlier points (i.e., noise). Thus, the initial parameter values for a $F_p(t_a)$, that are likely q-dependent would also need to be supplied by a researcher prior to analysis. While the method of a "global" fit can be appropriate for offline data analysis for certain systems, it is not a suitable solution for automated first-pass analysis of high-volumes of diverse data.

We compare the possibility of conducting a quantitative analysis of a $C2_{\text{raw}}$ and a corresponding $C2_{\text{denoised}}$ while considering $C1$ cuts with $\Delta t_a = 1\ fr^{-1}$ without a prior knowledge of the dynamic's parameter ranges (Fig. 2). An example of a noisy $C2_{\text{raw}}$ (350 frames) from the validation set [Fig. 2(a)] is passed through the model, resulting in the $C2_{\text{denoised}}$ with a significantly reduced level of noise [Figs. 2(b) and 2(c)]. The latent space representations of the 100×100 frames $C2_{\text{reduced}}$ are close to the center of the training set [Fig. 2(d)], indicating that the model output is likely a valid representation of the dynamics in the input. To extract the parameters, each of the $C1$ cuts is fitted to Eq. (3). Cuts with less than 5 t_d values are not considered. The results of the fit are shown in Figs. 2(e) to 2(h). The fits for the $C1$ cuts from the $C2_{\text{raw}}$ with $\Delta t_a = 35\ fr^{-1}$ are provided as ground truth values. Apparently,

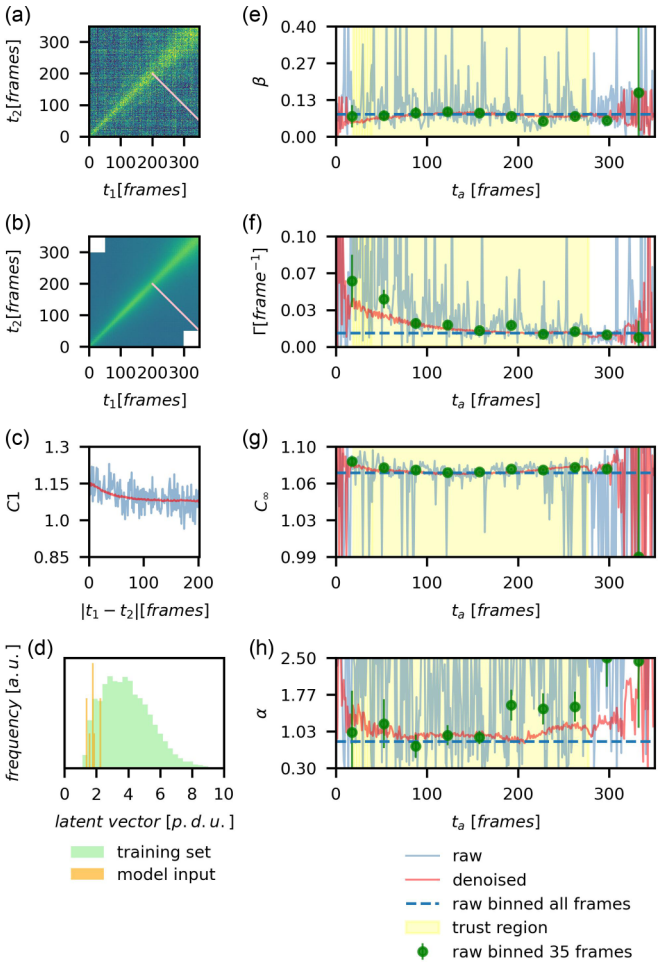

FIG. 2. Comparison of quantitative analysis for $C2_{\text{raw}}$ and $C2_{\text{denoised}}$. (a) $C2_{\text{raw}}$, (b) $C2_{\text{denoised}}$, (c) $C1$s taken along the line in (a,b), (d) distribution of latent vector lengths for the training set (green) and the $C2_{\text{reduced}}$s (orange), (e)–(h) plots for dynamics parameters values according to the legend. The *trust regions* are calculated for fits of $C2_{\text{denoised}}$.

high noise in the $C2_{\text{raw}}$ does not allow extracting meaningful information about the temporal evolution of the dynamics parameters without restricting the fit parameters within narrow regions or increasing the bin width. The $C2_{\text{denoised}}$, on the other hand, produces smooth, slowly evolving trends, matching the ground truth values.

While binning and averaging $C1$s reduces the noise, it is not a universal solution because the material's dynamics can vary considerably within an age bin, and therefore, a single set of parameters cannot describe them. For the example shown in Fig. 2, this becomes apparent when considering averaging all available $C1$s and fitting the result to Eq. (3). Since β and C_∞ are not changing in this experiment, the resulting values for these parameters are close to the ground truth values. However, the values for Γ and α are not close to their corresponding average values.

For automatic analyses to become autonomous, it is important to flag the results that cannot be fully trusted. For example, if the dynamics are not fully captured by the experiment, parameters of Eq. (3) become mutually dependent

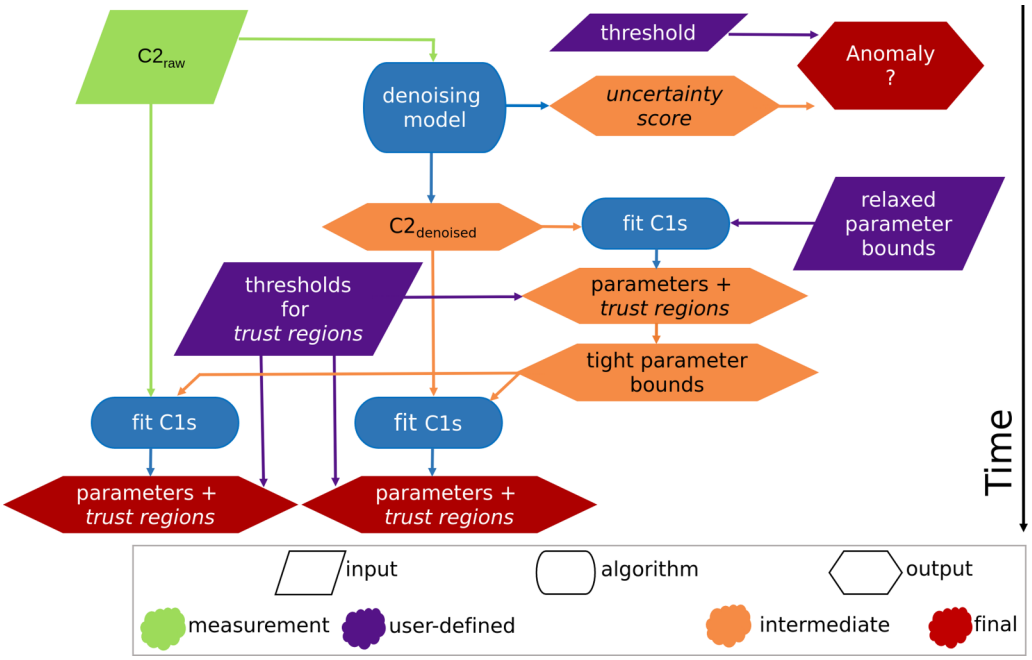

FIG. 3. Smart analysis workflow. Shape keys: rhombus—data supplied to the algorithm, rounded shape—model, diamond shape—result of calculations. Color keys: green—results of XPCS measurements, purple—supplied by a user, orange—intermediate results, red—the final results.

and the outputs of the nonlinear regression can be misleading because several different sets of the dynamics parameters can produce almost equally good fits. Therefore, we introduce the concept of *trust regions* for all fit parameters, noting that this may be generally applied even to traditionally obtained $C1$s for equilibrium dynamics. A *trust region* of a parameter is a binary vector with the length equal to the number of bins along t_a, which indicates whether the parameter is likely to be reliably identified within each bin. The binary values are determined by several criteria including the rate of the dynamics, goodness-of-fit to Eq. (3), correlations between parameters and parameters' relative errors.

When a material's dynamics are slow with respect to the maximum lag t_d covered by the experiment, the baseline C_∞ cannot be reliably identified. Likewise, the true value of β cannot be extracted if the dynamics are not fully captured by the experiment. In our algorithmic approach, the quantitative measure for establishing respective thresholds for the dynamics rate is a half-time $T_{1/2}$, i.e., the time it takes for the contrast to decay by half:

$$T_{1/2} = \left(\frac{ln2}{2}\right)^{\frac{1}{\alpha}} \frac{1}{\Gamma}. \tag{4}$$

For slow dynamics, when the half-time is larger than a user-defined portion of $C1$'s length, the *trust region* values for C_∞ and β are set to zero: not trustable. In the case of fast dynamics, when the half-time is less than a certain number of frames, the *trust region* for β is set to zero. However, the C_∞ can be reliably extracted for fast dynamics and hence its *trust region* values are set to 1. Similarly, threshold-based conditions are defined for the correlation coefficient between parameters, relative errors of the parameters, and R^2 measure of the fit.

III. RESULTS

A. Smart analysis workflow

Depending on the application, there are various ways to design an analysis workflow that includes $C2_{\text{raw}}$ and $C2_{\text{denoised}}$ as well as any prior knowledge about the material's dynamics and parameters of the experimental setup, e.g., maximum contrast factor. The DAE provides the means for constructing modules tackling various aspects of the analysis: anomaly detection, automated fits for dynamics parameters, disentangling average dynamics and fluctuations, and so on. The modules can be arranged within the same workflow.

Here, we present an example of a smart analysis workflow (SAW) for extracting the dynamics parameters from a $C2_{\text{raw}}$ as illustrated in Fig. 3. The word *smart* refers to the fact that the analysis is performed in the contexts of the intermediate results. This approach allows to avoid unnecessary fits of the anomalous data, automatic adjustment of the fit parameters' bounds, and evaluation of trustworthiness of results, which are crucial information for autonomous and high-throughput experiments.

In SAW, a $C2_{\text{denoised}}$ is first used for flagging unusual observations and then for an iterative fit of the dynamics parameters. The last step of the analysis results in two sets of parameters with corresponding errors and *trust regions*: one for the $C2_{\text{raw}}$ and one for the $C2_{\text{denoised}}$. A step-by-step algorithm for extracting the parameters is provided in the SM [40].

Optional flagging of unusual observations is done based on whether the uncertainty score for the $C2_{\text{denoised}}$ is above a user-defined threshold. Upon calculating the uncertainty score for the $C2_{\text{denoised}}$, one may wish to extend the workflow to investigate possible dynamic heterogeneities of the sample by relying on the denoised output and the results of the fits.

The bias of the denoising model provides the opportunity to separate the average "envelope" dynamics and stochastic heterogeneities by subtracting either the $C2_{\text{denoised}}$ or the $C2$, computed based on the fit parameters for the $C2_{\text{denoised}}$, from the original $C2_{\text{raw}}$. The residual $C2$ can be studied with the fourth-order time correlation function [45,46].

According to the demonstrated SAW, $C2_{\text{denoised}}$ is used for automatically narrowing the parameters' boundaries based on the range of its initial fit results within the valid *trust regions*. The new boundaries are then used to fit both the $C2_{\text{raw}}$ and the $C2_{\text{denoised}}$. Naturally, when reporting experimental outcomes, the fitting results of unprocessed $C2_{\text{raw}}$ have a higher priority with respect to the fitting results of $C2_{\text{denoised}}$ because signal processing, such as noise removal, introduces an uncertainty. However, it is also possible to supplement the fitting results for $C2_{\text{raw}}$ with the fitting results for $C2_{\text{denoised}}$ for regions in parameter space where low signal-to-noise ratio in raw data prevents one from obtaining meaningful parameter values. Note that the fitting results for $C2_{\text{denoised}}$ can be tested against $C2_{\text{raw}}$ with goodness-of-fit measures such as R^2. Alternatively, one can use the fitting results for $C2_{\text{denoised}}$ to select quasiequilibrium regions for the analysis of $C2_{\text{raw}}$ or for global fit of a $C2_{\text{raw}}$ (see SM [40]).

The outputs of SAW include the *trust regions*, which help to drive the decisions regarding using the results for reporting or for driving an autonomous experiment. The thresholds for calculating the *trust regions* can be established once and reused between measurements. This would allow for the XPCS facilities users to get scientific insights about complex material dynamics during the experiments even for new samples.

B. Application to new data: Standard analysis

We test SAW for XPCS analysis on previously reported data for 3D printing with an ink containing lithium titanate particles [20]. The experiment is not a part of the training or the validation datasets used in the current work. However, the far-from-equilibrium dynamics of the ink, exhibited during its deposition and recovery, are similar to the types of dynamics used during the models' training. Thus, the experiment presents an intended use case for the DAE.

There are several factors that limit the signal-to-noise ratio in this experiment: operando character of the measurements, beam-sensitivity of the ink, and anisotropy of the dynamics, which requires selection of small ROIs in the reciprocal space for analysis. In the original work, unevenly spaced age bins of various width were selected to obtain good quality "aged" $C1$ slices [20]. Here, we test the advantages of applying the DAE for conducting the same analysis with bin width $\Delta t_a = 1\ fr^{-1}$.

Figure 4 provides an example of the DAE being applied to the $C2_{\text{raw}}$ for one of the ROIs. Since the dynamics is relatively fast, the points near the t_a axis are denoised when the model is applied in a sliding window fashion and the points away from t_a are denoised via *down-and-up mapping* approach. The difference in level of noise for $C2_{\text{raw}}$ and $C2_{\text{denoised}}$ can be seen from corresponding $C1$s taken at different ages. The bin width of 1 frame has been considered for $C2_{\text{denoised}}$ and resulting $C1_{\text{denoised}}(t_a)$ are fitted to Eq. (3) to obtain $C1_{\text{fit}}(t_a)$. $C1_{\text{raw}}(t_a)$ for bin widths 1 and 100 are compared against the same $C1_{\text{fit}}(t_a)$. As expected, the fit to denoised version of correlation function describes the raw data well. As the width of the bins increases, the values of $C1_{\text{raw}}(t_a)$ lie closer to the corresponding $C1_{\text{fit}}(t_a)$ lines. The model captures well not only the characteristic time of the signal decorrelation, but also the change in the contrast factor β, caused by initial fast nonergodic dynamics with timescales outside of the experimental time window.

The sample's dynamics for multiple ROIs is quantified according to the scheme in Fig. 3. We calculate the drift velocity from relaxation times at different wave vectors along deposition direction $\Phi = -90^\circ$ using only denoised $C2$s and compare them to the values obtained in [20], as shown in Fig. 4(g). The values obtained from $C2_{\text{denoised}}$ are in good agreement with the results of the original analysis, confirming the DAE does not distort the data. Instead, it eliminates the need to carefully select the age bins and ensures the ultimate temporal resolution of an acquisition period. The results shown in Fig. 4(g) are obtained from four experiments, measured at differed acquisition rates and beam attenuation level. All four experiments were analyzed with the same threshold settings.

C. Application to new data: Outlier detection

While it is typically easy for a researcher to visually compare two or more $C2$s or identify an unusual observation, the high dimensionality of the data complicates such task for a machine during an automated analysis. Thus, an efficient dimensionality reduction approach must be employed. By using the latent space representation generated by the DAE, it is possible to quantitatively compare $C2$s from a group of measurements even if the model was not exposed to the same types of dynamics during training. This is suitable for clustering or anomaly detection in series of XPCS experiments.

We test the application of the model for instabilities detection during data collection for static dynamics of a LSAT sample (MTI Corporation) at the CSX beamline of NSLS-II. Six scattering time series, 7200 frames each (exposure 1.0875 seconds) are purposefully collected during the synchrotron accelerator testing to introduce a disturbance of experimental conditions. Three of the series are collected in a stable beam regime, one series is recorded with small disturbance (semi-stable regime) of the beam energy and one series is recorded during strong beam energy disturbance (unstable regime). From each series, 72 model inputs (size 100×100 frames) are generated by considering every 72nd frame with different starting points. Since there is no dynamics, each of the *down-mapped* inputs from a single series contains the same information, but different noise. To make the comparison sensitive to the mean value and variance of the contrast, the inputs are not scaled, but are clipped within values of 0 and 3. The values at $t_d = 0$ are extrapolated with the neighboring values as it is done for a routine model application.

Figure 5 shows the representations of stable, semi-stable, and unstable series in the plane of two latent space coordinates. All coordinates were standardized based on the values for the stable series. Inputs formed from the stable scattering series are located in close proximity of each other, forming a tight cluster. This confirms the contractive property of the

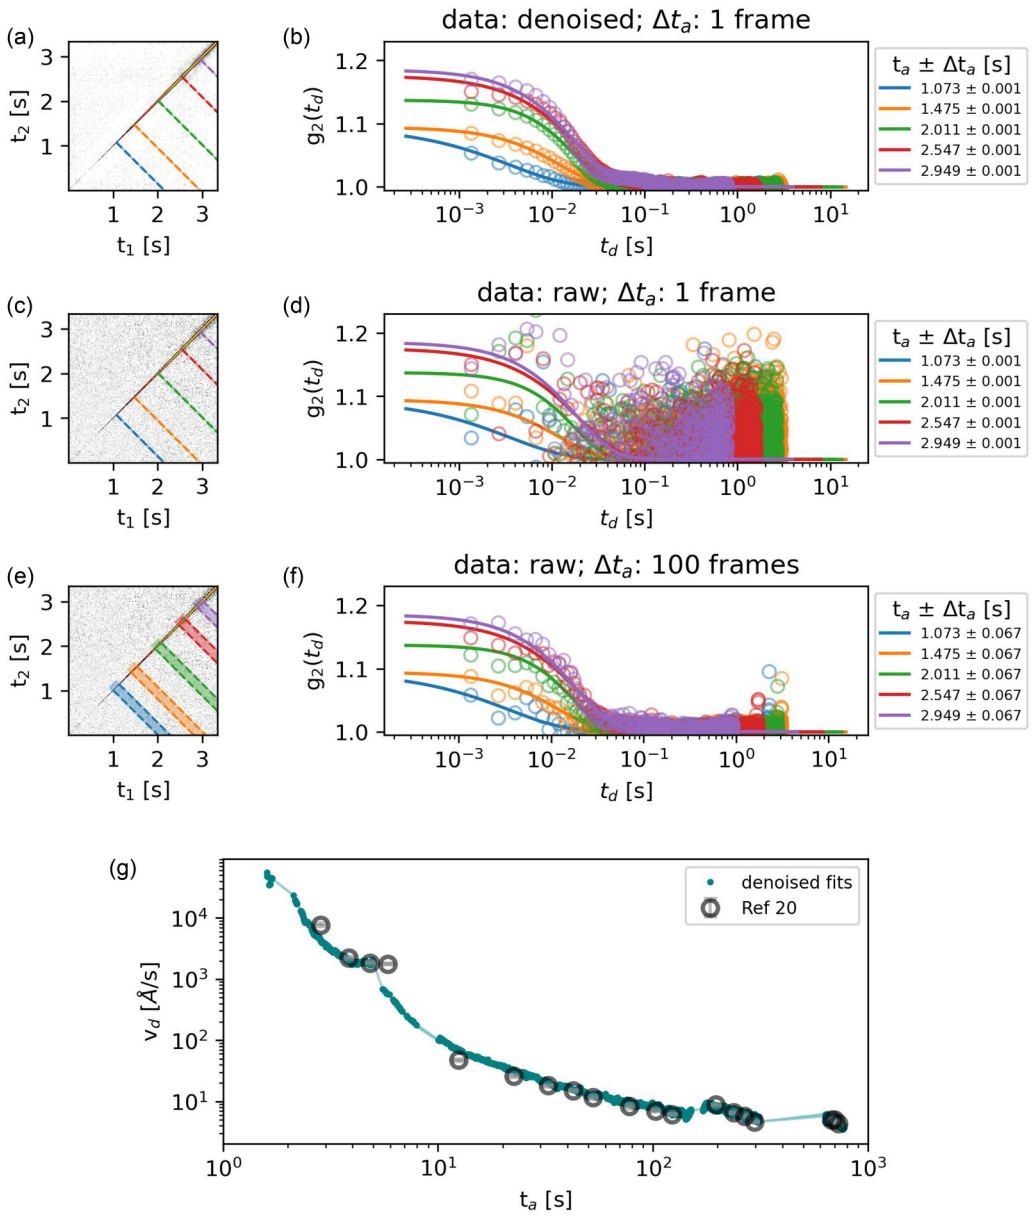

FIG. 4. Comparison of use $C2_{\text{raw}}$ and $C2_{\text{denoised}}$ from 3D printing experiments [20] for quantitative analysis. (a) $C2_{\text{denoised}}$ and bins of width 1 frame (dashed lines). (b) $C1$s corresponding to the bins in (a) (open circles) and their fits to Eq. (3) (solid lines). (c) $C2_{\text{raw}}$ and bins of width 1 frame (dashed lines). (d) $C1$s corresponding to the bins in (c) (open circles) and fits to denoised data from (b) (solid lines). (e) $C2_{\text{raw}}$ and bins of width 100 frames (dashed lines). (f) $C1$s corresponding to the bins in (e) (open circles) and fits to denoised data from (b) (solid lines). (g) Drift velocity obtained from a set of experimental scattering series, calculated using $C2_{\text{denoised}}$ (solid teal circles) and from Ref. [20]. Results from $C2_{\text{denoised}}$ have 7854 points and the previous results have 19 points.

DAE — the fact that similar inputs are close in the latent space. Having multiple $C2$s representing the same dynamics through the *down-mapping* procedure allows estimating the characteristic size of the cluster, to which all other points in the latent space should be compared to.

The disturbance of the beam intensity does not affect the $C2$ as much as long as the beam is not lost. This is expected due to normalization in Eq. (1). In agreement to this expectation, the inputs collected from experiment with a small disturbance (i.e., semi-stable series) are close to the undisturbed (stable) series in the latent space for all eight coordinates (see SM [40]). However, when the disturbance is so strong that the beam is lost for part of the series, the latent space representations of the corresponding $C2$s are much further away from the undisturbed series along some of the directions. In our case, the distance between the stable and unstable series is the largest along the sixth coordinate. The unstable series forms a separate cluster, where points are much more spread out than for stable and semi-stable series. Selecting a threshold distance between points filters out the measurements taken at unstable experimental conditions. Thus, it is possible to identify unusual observations in a set of XPCS series by looking into the DAE's latent representation of $C2$s and comparing the distance between series using algo-

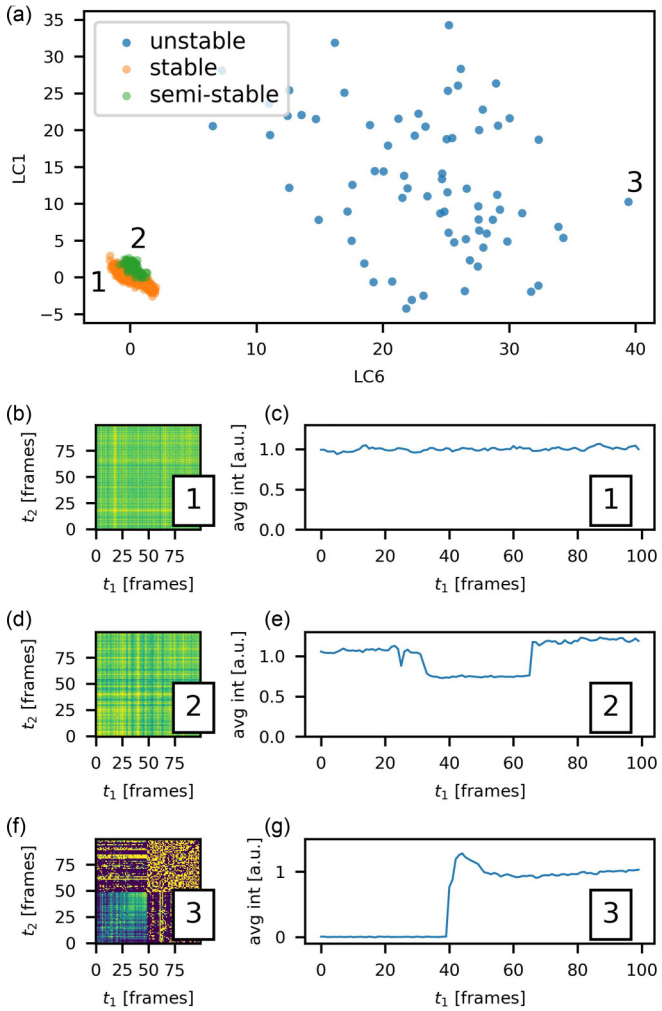

FIG. 5. Detection of anomalies in static experiments by encoding data with the DAE model. (a) Points in the latent space corresponding to inputs generated from the experimental series. The points are shown in the plane formed by the first and sixth coordinates. $C2$s for (b) stable, (d) semi-stable, and (f) unstable experimental conditions. (c), (e), and (f) Average intensities on the detector for the respective experiments on the left.

rithms, such as local outlier detection [47] or isolation forest [48].

Importantly, the above approach allows to compare the experiments directly based on their observations and not merely on the experimental conditions, which may have little effect on the photon correlation values. Here, the semi-stable experiment in Fig. 5(e) could be marked as an anomaly based on the variation of the average intensity that, in fact, results in only insignificant noise in the $C2$.

IV. DISCUSSION

Increased data collection rates at light source facilities challenge [31,34] established approaches to handling the experimental results. Traditional analysis of XCPS data is a multistep, often iterative, process that requires continuous evaluation of intermediate results by a domain expert. Large volume of observations and a researcher's subjectivity in alleviating low signal-to-noise ratio complicate proper extraction of valuable information from experiments. Concerningly, this creates uncertainty during data collection, preventing efficient experimental steering. Computationally reducing noise in $C2$s as a data processing step helps to achieve several goals: automating the analysis workflow, improving temporal resolution of parameters that quantify the system's dynamics, and increasing quantitative usage of data with high cost of collection or from radiation-sensitive samples [49].

Here, we demonstrate how a DAE model can be included into the smart analysis workflow of XPCS data for experiments with nonequilibrium dynamics. Quantification of the model's bias helps driving the decision about the use of its outputs as well as flagging unusual observations, such as heterogeneities or dynamics that are very different from ageing. The denoised correlation function can be used for optimizing and complementing the fit of the raw data. The concept of *trust regions* combines the assessment of the fit quality as well as domain expertise, which helps to not only report the most reliable results, but also to automate sequential narrowing of the fit parameter boundaries.

The model's performance for unseen 3D printing data (not included in the training/validation datasets) demonstrates the advantages of its application to complex XPCS experiments. For measurements where the material's dynamics are similar enough to the model's training dataset, the analysis does not require a human-in-the-loop after all the thresholds are selected prior to analysis. Moreover, the analysis performs the known required tasks to enable autonomous data acquisition. It provides values of dynamics parameters, which are important for making decisions about adapting data acquisition parameters such as acquisition rate, exposure time, duration of data acquisition, or about changing the sample and/or the processing parameters.

We further demonstrate how encoded representations of $C2$s can be used for quantitative comparison of two or more scattering time series, which can be useful for identifying "anomalies" such as experimental instabilities or changes with scientific meaning (e.g., phase transitions). The comparison can even be done for the types of dynamics that are not present in the training set.

In conclusion, in this work we demonstrate how a CNN-based DAE can be used for $C2$s with nonequilibrium dynamics for experimental scattering series of arbitrary size. The addition of the DAE to XPCS analysis along with quantifying uncertainty helps with automating the analysis and improves the temporal resolution of extracted parameters. Besides, analysis of residuals and latent space representations of the inputs helps to detect anomalous dynamics that go beyond monotonic ageing. This property can be employed for recognizing heterogeneities or phase transitions. Several examples of incorporating the DAE in analysis workflows for experimental data, collected at NSLS-II, demonstrate its effectiveness for unseen data and the diversity of its applications. Denoising and encoding properties of the model are promising for various online and offline data analysis tasks for speckle dynamics, and they are likely an essential tool to enable autonomous experiments at XPCS facilities. The seemingly wide applicability of these types of DAE models further illustrates the their importance in the field.

At the time of publication, the code used to produce the analysis presented in this work as well as a pretrained model was found at [38].

ACKNOWLEDGMENTS

We thank Daniel Allan for insightful discussions related to parallel computing with DASK. This research used the CHX and CSX beamlines and resources of the National Synchrotron Light Source II, a U.S. Department of Energy (DOE) Office of Science User Facility operated for the DOE Office of Science by Brookhaven National Laboratory(BNL) under Contract No. DE-SC0012704 and under a BNL Laboratory Directed Research and Development (LDRD) Project No. 20-038 "Machine Learning for Real-Time Data Fidelity, Healing, and Analysis for Coherent X-ray Synchrotron Data."

[1] U. Bergmann, J. Corlett, S. Dierker, R. Falcone, J. Galayda, M. Gibson, J. Hastings, B. Hettel, J. Hill, Z. Hussain *et al.*, Science and Technology of Future Light Sources: A White Paper, Tech. Rep. Stanford Linear Accelerator Center (SLAC), 2009 (Reports No. ANL-08/39, No. BNL-81895-2008, No. LBNL-1090E-2009, and No. SLAC-R-917); see also https://www.osti.gov/biblio/950452.

[2] M. E. Couprie, New generation of light sources: Present and future, J. Electron Spectrosc. Relat. Phenom. **196**, 3 (2014).

[3] I. Vartanyants and A. Singer, in *Synchrotron Light Sources and Free-Electron Lasers*, edited by E. J. Jaeschke, S. Khan, J. R. Schneider, and J. B. Hastings (Springer, New York, 2016).

[4] P. Grybos, P. Kmon, P. Maj, and R. Szczygiel, 32k Channels readout IC for single photon counting detectors with 75 μm pitch, ENC of 123 e- rms, 9 e- rms offset spread and 2% rms gain spread, in *2015 IEEE Biomedical Circuits and Systems Conference (BioCAS)* (IEEE, Piscataway, NJ, 2015), pp. 1–4.

[5] X. Llopart, M. Campbell, R. Dinapoli, D. San Segundo, and E. Pernigotti, Medipix2: A 64-k pixel readout chip with 55-/spl mu/m square elements working in single photon counting mode, IEEE Trans. Nucl. Sci. **49**, 2279 (2002).

[6] Q. Zhang, E. M. Dufresne, and A. R. Sandy, Dynamics in hard condensed matter probed by x-ray photon correlation spectroscopy: Present and beyond, Curr. Opin. Solid State Mater. Sci. **22**, 202 (2018).

[7] A. Madsen, A. Fluerasu, and B. Ruta, Structural dynamics of materials probed by x-ray photon correlation spectroscopy, in *Synchrotron Light Sources and Free-Electron Lasers: Accelerator Physics, Instrumentation and Science Applications* (Springer International Publishing, New York, 2016), pp. 1617–1641.

[8] O. G. Shpyrko, X-ray photon correlation spectroscopy, J. Synchrotron Radiat. **21**, 1057 (2014).

[9] S. K. Sinha, Z. Jiang, and L. B. Lurio, X-ray photon correlation spectroscopy studies of surfaces and thin films, Adv. Mater. **26**, 7764 (2014).

[10] M. Takahasi, X-ray diffraction study of crystal growth dynamics during molecular-beam epitaxy of III-V semiconductors, J. Phys. Soc. Jpn. **82**, 021011 (2013).

[11] J. G. Ulbrandt, M. G. Rainville, C. Wagenbach, S. Narayanan, A. R. Sandy, H. Zhou, K. F. Ludwig, Jr., and R. L. Headrick, Direct measurement of the propagation velocity of defects using coherent x-rays, Nat. Phys. **12**, 794 (2016).

[12] R. L. Headrick, J. G. Ulbrandt, P. Myint, J. Wan, Y. Li, A. Fluerasu, Y. Zhang, L. Wiegart, and K. F. Ludwig, Coherent x-ray measurement of step-flow propagation during growth on polycrystalline thin film surfaces, Nat. Commun. **10**, 2638 (2019).

[13] G. Ju, D. Xu, M. J. Highland, C. Thompson, H. Zhou, J. A. Eastman, P. H. Fuoss, P. Zapol, H. Kim, and G. B. Stephenson, Coherent x-ray spectroscopy reveals the persistence of island arrangements during layer-by-layer growth, Nat. Phys. **15**, 589 (2019).

[14] E. B. Trigg, L. Wiegart, A. Fluerasu, and H. Koerner, Dynamics of polymerization and gelation in epoxy nanocomposites via x-ray photon correlation spectroscopy, Macromolecules **54**, 6575 (2021).

[15] B. M. Yavitt, D. Salatto, Z. Huang, Y. T. Koga, M. K. Endoh, L. Wiegart, S. Poeller, S. Petrash, and T. Koga, Revealing nanoscale dynamics during an epoxy curing reaction with x-ray photon correlation spectroscopy, J. Appl. Phys. **127**, 114701 (2020).

[16] T. Hoshino, Y. Okamoto, A. Yamamoto, and H. Masunaga, Heterogeneous dynamics in the curing process of epoxy resins, Sci. Rep. **11**, 9767 (2021).

[17] L. Wiegart, G. Doerk, M. Fukuto, S. Lee, R. Li, G. Marom, M. Noack, C. Osuji, M. Rafailovich, J. Sethian *et al.*, Instrumentation for in situ/operando x-ray scattering studies of polymer additive manufacturing processes, Synchrotron Radiation News **32**, 20 (2019).

[18] K. J. Johnson, L. Wiegart, A. C. Abbott, E. B. Johnson, J. W. Baur, and H. Koerner, In operando monitoring of dynamic recovery in 3D-printed thermoset nanocomposites by XPCS, Langmuir **35**, 8758 (2019).

[19] M. T. Arango, Y. Zhang, C. Zhao, R. Li, G. Doerk, D. Nykypanchuk, Y. K. Chen-Wiegart, A. Fluerasu, and L. Wiegart, Ink-substrate interactions during 3D printing revealed by time-resolved coherent x-ray scattering, Mater. Today Phys. **14**, 100220 (2020).

[20] C.-H. Lin, K. Dyro, O. Chen, D. Yen, B. Zheng, M. T. Arango, S. Bhatia, K. Sun, Q. Meng, L. Wiegart *et al.*, Revealing mesostructure dynamics in additive manufacturing of energy storage via operando coherent x-ray scattering, Appl. Mater. Today **24**, 101075 (2021).

[21] A. Ricci, G. Campi, B. Joseph, N. Poccia, D. Innocenti, C. Gutt, M. Tanaka, H. Takeya, Y. Takano, T. Mizokawa *et al.*, Intermittent dynamics of antiferromagnetic phase in inhomogeneous iron-based chalcogenide superconductor, Phys. Rev. B **101**, 020508(R) (2020).

[22] N. Begam, A. Ragulskaya, A. Girelli, H. Rahmann, S. Chandran, F. Westermeier, M. Reiser, M. Sprung, F. Zhang, C. Gutt *et al.*, Kinetics of Network Formation and Heterogeneous Dynamics of an Egg White Gel Revealed by Coherent X-Ray Scattering, Phys. Rev. Lett. **126**, 098001 (2021).

[23] K. Ludwig, F. Livet, F. Bley, J.-P. Simon, R. Caudron, D. Le Bolloc'h, and A. Moussaid, X-ray intensity fluctuation spectroscopy studies of ordering kinetics in a cu-pd alloy, Phys. Rev. B **72**, 144201 (2005).

[24] Z. Evenson, B. Ruta, S. Hechler, M. Stolpe, E. Pineda, I. Gallino, and R. Busch, X-Ray Photon Correlation Spectroscopy Reveals Intermittent Aging Dynamics in a Metallic Glass, Phys. Rev. Lett. **115**, 175701 (2015).

[25] J.-D. Su, A. R. Sandy, J. Mohanty, O. G. Shpyrko, and M. Sutton, Collective pinning dynamics of charge-density waves in $1T$-TaS_2, Phys. Rev. B **86**, 205105 (2012).

[26] F. Ehrburger-Dolle, I. Morfin, F. Bley, F. Livet, G. Heinrich, S. Richter, L. Piche, and M. Sutton, XPCS investigation of the dynamics of filler particles in stretched filled elastomers, Macromolecules **45**, 8691 (2012).

[27] G. Brown, P. A. Rikvold, M. Sutton, and M. Grant, Speckle from phase-ordering systems, Phys. Rev. E **56**, 6601 (1997).

[28] A. Madsen, A. Fluerasu, and B. Ruta, Structural dynamics of materials probed by x-ray photon correlation spectroscopy, in *Synchrotron Light Sources and Free-Electron Lasers: Accelerator Physics, Instrumentation and Science Application* (Springer, New York, 2020).

[29] O. Bikondoa, On the use of two-time correlation functions for x-ray photon correlation spectroscopy data analysis, J. Appl. Crystallogr. **50**, 357 (2017).

[30] F. Khan, S. Narayanan, R. Sersted, N. Schwarz, and A. Sandy, Distributed x-ray photon correlation spectroscopy data reduction using hadoop mapreduce, J. Synchrotron Radiat. **25**, 1135 (2018).

[31] Q. Zhang, E. M. Dufresne, Y. Nakaye, P. R. Jemian, T. Sakumura, Y. Sakuma, J. D. Ferrara, P. Maj, A. Hassan, D. Bahadur *et al.*, 20 μs-resolved high-throughput x-ray photon correlation spectroscopy on a 500k pixel detector enabled by data-management workflow, J. Synchrotron Radiat. **28**, 259 (2021).

[32] A. Madsen, R. L. Leheny, H. Guo, M. Sprung, and O. Czakkel, Beyond simple exponential correlation functions and equilibrium dynamics in x-ray photon correlation spectroscopy, New J. Phys. **12**, 055001 (2010).

[33] J. D. Hunter, Matplotlib: A 2D graphics environment, Comput. Sci. Eng. **9**, 90 (2007).

[34] C. Wang, U. Steiner, and A. Sepe, Synchrotron big data science, Small **14**, 1802291 (2018).

[35] N. Schwarz, S. Campbell, A. Hexemer, A. Mehta, and J. Thayer, Enabling scientific discovery at next-generation light sources with advanced ai and hpc, in *Smoky Mountains Computational Sciences and Engineering Conference* (Springer, New York, 2020), pp. 145–156.

[36] T. Konstantinova, L. Wiegart, M. Rakitin, A. M. DeGennaro, and A. M. Barbour, Noise reduction in x-ray photon correlation spectroscopy with convolutional neural networks encoder-decoder models, Sci. Rep. **11**, 14756 (2021).

[37] T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley, J. Hamrick, J. Grout, S. Corlay, P. Ivanov, D. Avila, S. Abdalla, and C. Willing, Jupyter notebooks – a publishing format for reproducible computational workflows, in *Positioning and Power in Academic Publishing: Players, Agents and Agendas*, edited by F. Loizides and B. Schmidt (IOS, Amsterdam, 2016), pp. 87–90.

[38] Accompanying code, https://www.osti.gov/doecode/biblio/78420.

[39] G. Williams and D. C. Watts, Non-symmetrical dielectric relaxation behaviour arising from a simple empirical decay function, Trans. Faraday Soc. **66**, 80 (1970).

[40] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevResearch.4.033228 for additional details.

[41] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, Gradient-based learning applied to document recognition, Proc. IEEE **86**, 2278 (1998).

[42] M. A. Kramer, Nonlinear principal component analysis using autoassociative neural networks, AIChE J. **37**, 233 (1991).

[43] M. Neumeier, M. Betsch, A. Tollkühn, and T. Berberich, Variational autoencoder-based vehicle trajectory prediction with an interpretable latent space, in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)* (IEEE, New York, 2021), pp. 820–827.

[44] M. Abdar, F. Pourpanah, S. Hussain, D. Rezazadegan, L. Liu, M. Ghavamzadeh, P. Fieguth, X. Cao, A. Khosravi, U. R. Acharya *et al.*, A review of uncertainty quantification in deep learning: Techniques, applications and challenges, Information Fusion **76**, 243 (2021).

[45] D. Orsi, L. Cristofolini, G. Baldi, and A. Madsen, Heterogeneous and Anisotropic Dynamics of a 2D Gel, Phys. Rev. Lett. **108**, 105701 (2012).

[46] S. Maccarrone, G. Brambilla, O. Pravaz, A. Duri, M. Ciccotti, J.-M. Fromental, E. Pashkovski, A. Lips, D. Sessoms, V. Trappe *et al.*, Ultra-long range correlations of the dynamics of jammed soft matter, Soft Matter **6**, 5514 (2010).

[47] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander, Lof: identifying density-based local outliers, in *Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data* (ACM, New York, 2000), pp. 93–104.

[48] F. T. Liu, K. M. Ting, and Z.-H. Zhou, Isolation forest, in *2008 Eighth IEEE International Conference on Data Mining, Pisa, Italy* (IEEE, Piscataway, NJ, 2008), pp. 413–422.

[49] J. Möller, M. Sprung, A. Madsen, and C. Gutt, X-ray photon correlation spectroscopy of protein dynamics at nearly diffraction-limited storage rings, IUCrJ **6**, 794 (2019).