

Unsupervised learning approaches to characterize heterogeneous samples using X-ray single particle imaging

YULONG ZHUANG,^{a,b} SALAH AWEL,^c ANTON BARTY,^c RICHARD BEAN,^d

JOHAN BIELECKI,^d MARTIN BERGEMANN,^d BENEDIKT J. DAURER,^e

TOMAS EKEBERG,^f ARMANDO D. ESTILLORE,^c HANS FANGOHR,^{a,b,d,r}

KLAUS GIEWEKEMEYER,^d MARK S. HUNTER,^g MIKHAIL KARNEVSKIY,^d RICHARD

A. KIRIAN,^h HENRY KIRKWOOD,^d YOONHEE KIM,^d JAYANATH KOLIYADU,^d

HOLGER LANGE,^{i,j} ROMAIN LETRUN,^d JANNIK LÜBKE,^{c,i,k} ABHISHEK MALL,^{a,b}

THOMAS MICHELAT,^d ANDREW J. MORGAN,^l NILS ROTH,^{c,k} AMIT K. SAMANTA,^c

TOKUSHI SATO,^d ZHOU SHEN,^e MARCIN SIKORSKI,^c FLORIAN SCHULZ,^{i,s} JOHN

C. H. SPENCE,^h PATRIK VAGOVIC,^{c,d} TAMME WOLLWEBER,^{a,b,i} LENA WORBS,^{c,k}

P. LOURDU XAVIER,^{a,c,i} OLEKSANDR YEFANOV,^c FILIPE R. N. C. MAIA,^{f,m}

DANIEL A. HORKE,^{c,i,n} JOCHEN KÜPPER,^{c,i,k,o} N. DUANE LOH,^{e,p} ADRIAN

P. MANCUSO,^{d,q} HENRY N. CHAPMAN^{c,i,k} AND KARTIK AYYER^{a,b,i*}

^aMax Planck Institute for the Structure and Dynamics of Matter, 22761 Hamburg,

Germany, ^bCenter for Free-Electron Laser Science, 22761 Hamburg, Germany,

^cCenter for Free-Electron Laser Science, DESY, 22607 Hamburg, Germany,

^dEuropean XFEL, 22869 Schenefeld, Germany, ^eCenter for BioImaging Sciences,

National University of Singapore, Singapore 117557, ^fDept. of Cell and Molecular
Biology, Uppsala University, 75124 Uppsala, Sweden, ^gLinac Coherent Light Source,

SLAC National Accelerator Laboratory, Menlo Park, 94025, USA, ^hDepartment of

Physics, Arizona State University, Tempe, AZ 85287, USA, ⁱThe Hamburg Center

for Ultrafast Imaging, Universität Hamburg, 22761 Hamburg, Germany, ^jInstitute of

Physical Chemistry, Universität Hamburg, 20146 Hamburg, Germany, ^kDepartment

of Physics, Universität Hamburg, 22761 Hamburg, Germany, ^lUniv. of Melbourne, Physics, Victoria, 3010, Australia, ^mNERSC, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA, ⁿInstitute for Molecules and Materials, Radboud University, 6525 AJ Nijmegen, Netherlands, ^oDepartment of Chemistry, Universität Hamburg, 20146 Hamburg, Germany, ^pDepartment of Physics, National University of Singapore, Singapore 117551, ^qDepartment of Chemistry and Physics, La Trobe Institute for Molecular Science, La Trobe University, Melbourne, Victoria 3086, Australia, ^rUniversity of Southampton, SO17 1BJ, Southampton, United Kingdom, and ^sInstitute of Nanostructure and Solid State Physics, University of Hamburg, Luruper Chaussee 149, 22761 Hamburg, Germany.

E-mail: kartik.ayyer@mpsd.mpg.de

Abstract

One of the outstanding analytical problems in X-ray single particle imaging (SPI) is the classification of structural heterogeneity, which is especially difficult given the low signal-to-noise ratios of individual patterns and that even identical objects can yield patterns that vary greatly when orientation is taken into consideration. We propose two methods which explicitly account for this orientation-induced variation and can robustly determine the structural landscape of a sample ensemble. The first, termed common-line principal component analysis (PCA) provides a rough classification which is essentially parameter-free and can be run automatically on any SPI dataset. The second method, utilizing variation auto-encoders (VAEs) can generate 3D structures of the objects at any point in the structural landscape. We implement both these methods in combination with the noise-tolerant expand-maximize-compress

(EMC) algorithm and demonstrate its utility by applying it to an experimental dataset from gold nanoparticles with only a few thousand photons per pattern and recover both discrete structural classes as well as continuous deformations. These developments diverge from previous approaches of extracting reproducible subsets of patterns from a dataset and open up the possibility to move beyond studying homogeneous sample sets and study open questions on topics such as nanocrystal growth and dynamics as well as phase transitions which have not been externally triggered.

1. Introduction

X-ray single particle imaging (SPI) is a method to reconstruct 3D structures of isolated nanoscale objects by collecting a large number of diffraction patterns using bright X-ray pulses. The diffraction patterns sample the object's three-dimensional Fourier transform along randomly oriented spherical slices, which enables a Fourier synthesis of the 3D structure as long as the orientations can be determined (Neutze *et al.*, 2000). However, the X-ray pulses are bright enough to destroy the sample after each shot, and so each pattern is collected from a different particle. If the particles are reproducible up to the resolution of interest, then determination of 3D structures is fairly straightforward, involving the determination of the orientation and incident X-ray fluence for each shot. Various algorithms have been proposed and implemented for this purpose including the expand–maximize–compress (EMC) algorithm (Loh & Elser, 2009), which has been used for a number of experimental demonstrations (Ekeberg *et al.*, 2015; Rose *et al.*, 2018; Shi *et al.*, 2018; Ayyer *et al.*, 2019; Ayyer *et al.*, 2021). Other methods utilizing intensity correlations have also been experimentally demonstrated (Kurta *et al.*, 2017; von Ardenne *et al.*, 2018).

However, one of the challenges in analyzing a serial dataset like that produced in an SPI experiment is the proper classification of patterns in terms of their struc-

tures. This is a necessary step in order to obtain a high resolution structure since the underlying assumption of the above mentioned algorithms is that the particles are identical, which is never true in practice. On the other hand, the framework used for this classification can be used to study datasets where the heterogeneity is not just a drawback, uncovering the landscape of structural variations in the sample. This requires the detection of discrete classes of object shapes representing contaminants, aggregates etc. and even the detection of continuous deformations depending on the scientific problem being studied.

In past years, different machine learning algorithms have been developed for use in this classification of structural variability. These algorithms are usually applied to the patterns themselves using methods including spectral clustering (Yoon *et al.*, 2011), support vector machines (Bobkov *et al.*, 2015), and convolutional neural networks (Zimmermann *et al.*, 2019; Ignatenko *et al.*, 2021).

In this work, we take an unsupervised learning approach to analyze an experimental dataset of more than 2.4 million diffraction patterns of nominally 42 nm cubic gold nanoparticles collected at the European XFEL (Ayyer *et al.*, 2021). The goal is to classify the entire structural ensemble, including separating discrete contaminants as well as to study any continuous structural variations that may be present. We discuss two approaches, the first of which produces an embedding of the diffraction patterns based on their 3D structures. This approach can be applied without modification to most SPI datasets. The second is a generative method using variational auto-encoders (VAEs) which enables us to visualize the 3D structure of the particle at any point along its landscape. Both methods are applied to the dataset to study a continuous, XFEL-induced deformation of the cubic nanoparticles to spheres.

2. Results

2.1. Experiment and dataset information

The dataset discussed in the rest of this work was collected as part of the experiment described in Ayer *et al.* (2021), which we review in brief: The SPI experiment was performed with the MHz-rate European XFEL (Decking *et al.*, 2020) at the SPB/SFX (single particles, biomolecules, and clusters/serial femtosecond crystallography) instrument (Mancuso *et al.*, 2019) with 6 keV photons in pulses with an average energy of 2.5 mJ (2.6×10^{12} photons) measured upstream of the focusing optics. The Adaptive Gain Integrating Pixel Detector (AGIPD) (Allahgholi *et al.*, 2019), was placed 705 mm downstream of the interaction region to collect each diffraction pattern up to a scattering angle of 8.3° . In this analysis, we will use data from the central region of the detector up to 1.8° .

The samples were nominally gold cubes with edge lengths 42 nm and were injected into the X-ray beam via an electrospray aerosolization aerodynamic-lens-stack sample delivery system. For this sample, 34,197,950 frames were recorded of which 2,451,068 were judged to contain sample diffraction (hit ratio $\sim 7.17\%$). Part of the dataset was collected with the European XFEL running at an intra-train repetition rate of 1.1 MHz, during which a high fraction of the diffraction patterns originated from spherical particles. This was found to be due to the pre-exposure of particles in the wings of the previous XFEL pulse in the train, which seemed to lead to melting. This effect disappeared when reducing the repetition rate to 0.5 MHz. As a result, the sample contains diffraction patterns of cubes, spheres and potentially cubes at different stages of melting/softening.

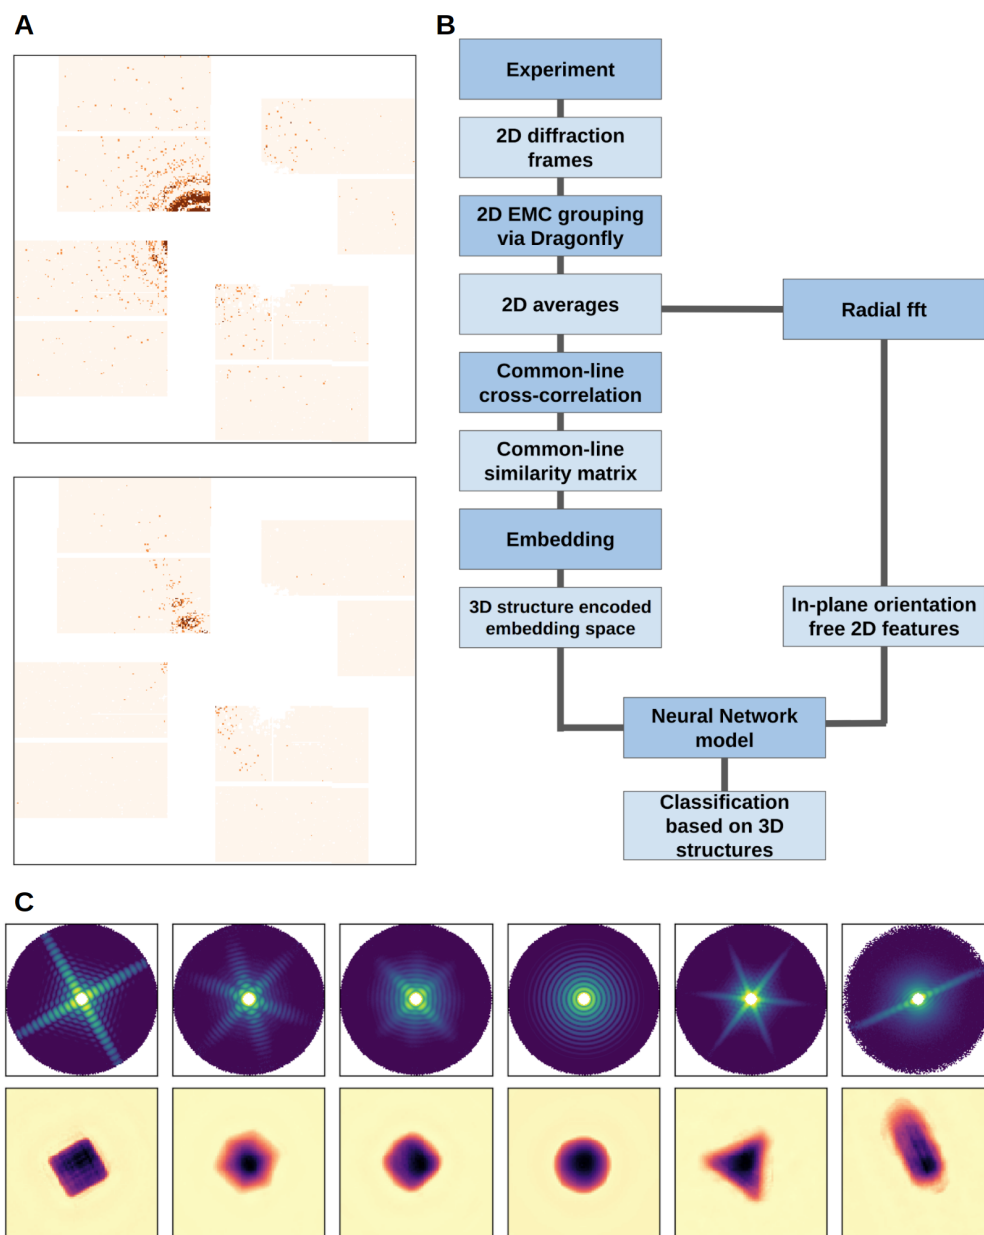


Fig. 1. (A) Examples of diffraction patterns in the data set. The color scale maximizes at 4 photons. (B) Workflow of the CLPCA method. (C) Upper row shows some typical frame averages on logarithmic scale in the dataset, while the lower row shows their corresponding real space 2D density projections via phase retrieval. The real-space field of view is 132.4 nm.

2.2. Classification of the entire ensemble

There are two major challenges in understanding the sample ensemble from the diffraction data. First, many of the patterns are too weak (10^{2-3} photons/pattern) to apply single shot analysis, as illustrated in Fig. 1(A). Second, the dataset is not structurally homogeneous, not just in that the sample set includes non-cubic particles, but also that a fraction of the cubic particles suffered varied levels of melting/softening due to pre-exposure by the previous pulse in the train.

We first developed a common-line principal component analysis (CLPCA) method to understand SPI datasets from an arbitrary ensemble of particles. Fig. 1(B) shows our workflow, which we discuss in detail below.

2.2.1. Generation of 2D averages via bootstrapping Due to the weak scattering signal in single diffraction frames, the first step is to average similar 2D detector frames together (accounting for in-plane rotations), which was done with the 2D classification procedure implemented in *Dragonfly* (Ayyer *et al.*, 2016). The code implements a modification of the EMC algorithm that classifies all 2D frames into a given number of averages (models, or classes in *Dragonfly*). This averaging improves the signal-to-noise ratio and fills in detector panel gaps, at the potential cost of washing out some structural variations. In order to mitigate the latter effect, one can use a very large number of models, but this can lead to instabilities in the iterative reconstruction and reduced signal-to-noise in individual class averages.

In order to get more pattern averages, a bootstrapping method was used by running the reconstruction 5 times with 200 models, each time with a random subset of 80% of the frames. In this way, each of the 1000 models are composed of a different group of similar 2D frames. In Fig. 1(C), we show an example of some typical 2D diffraction averages and corresponding projected electron density maps after phase retrieval using

a combination of the difference map and error reduction algorithms similar to that employed in Ayyer *et al.* (2019). This highlights the variety both in the samples, but also in the diffraction patterns from the same samples in different orientations, such as the rotated versions of identical cubes in the first two columns of Fig. 1(C).

2.2.2. Common line 3D classification The 2D EMC method can only group similar 2D frames together, but as mentioned above, diffraction patterns of the same object can look very different depending on the orientation. In order to understand the variations of structures in the sample, we need to further classify the averages through their 3D features, rather than considering them just as 2D images.

At small scattering angles, each average is a Fourier transform of a projection of the target object at a given orientation. According to the Fourier slice theorem, this means that each average represents a slice through the 3D Fourier transform of the object. Any two patterns of different orientations from the same object should share a common intersection line (at larger angles, these lines become arcs), as shown in the illustration of Fig. 2(A). The similarity of the diffraction intensities along the best “common-lines” between two patterns should be correlated to the similarity of their 3D structures. For each pair of 2D averages we defined the similarity between their common lines as the cross-correlation coefficient of the respective diffraction signals along those lines. The angles of the common lines are chosen to be those that yield the greatest degree of similarity. This yields a common line similarity matrix (CC-matrix) for all pairs of 2D intensity averages. Common lines have been used previously for orientation determination (Shneerson *et al.*, 2008; Singer & Shkolnisky, 2011; Yefanov & Vartanyants, 2013), where the optimal angles tell one how to fit the two slices in the 3D Fourier space. Here, we are interested in the similarity index of the common lines, as a tool for structural similarity analysis, rather than the relative angles at

which they are maximized.

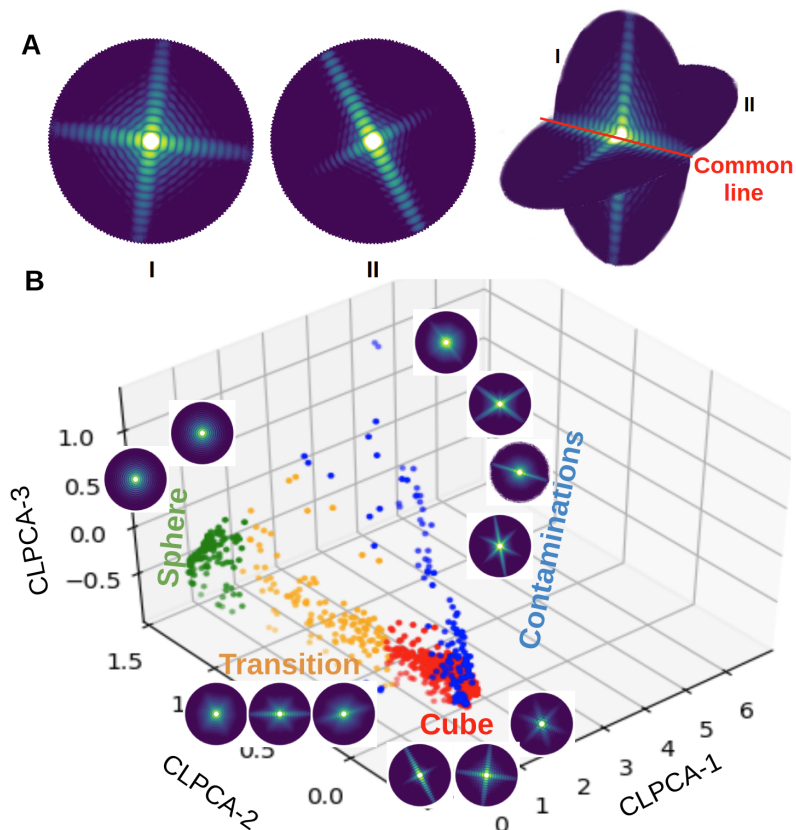


Fig. 2. (A) An illustration of the common-line between two patterns of similar shaped objects but different orientations. (B) The distribution of 1000 averages in the 3D CLPCA space. Different colors are manually divided groups from the first two components: contaminants (blue); cubes (red); transition (yellow); spheres (green). Typical patterns for averages in each group are also shown on logarithmic scale.

To get the landscape of the 3D structures in the sample, we performed dimensionality reduction on this CC-matrix and generate an embedding. Fig. 2(B) shows a 3D embedding of the 1000 averages calculated using principal component analysis. In the rest of this article, we refer to this 3D embedding as ‘CLPCA space’. Until this point, the entire procedure is fully unsupervised and can be run automatically for any SPI dataset.

By comparing with 2D patterns and their locations in CLPCA space, we then qualitatively divided the CLPCA space into 4 groups as shown in Fig. 2(B), red - cubes, green - spheres, blue - contaminants and yellow - rounded cubes. In the embedded space, CLPCA-1 separates these assorted patterns from those originating from spherical/cubic particles, CLPCA-2 tracks the transition between spherical and cubic particles and for spherical particles, CLPCA-3 is associated with their diameter. In addition, we see a sequence where particles transition from cubes to spheres. This is a clear trace of the pre-melting processes observed in the experiment. Without using any *a priori* knowledge about the sample set, we are able to obtain a rough classification of the dataset according to their 3D structures.

We also note from Fig. 2(B) that the dense cluster of cube patterns are correctly identified as being from the same shaped particles even though the patterns themselves vary extensively at different orientations. For the sake of simplicity we only perform the embedding with the PCA method, but other dimensionality reduction methods could also be used interchangeably. In Appendix A, we show a comparison between PCA and other embedding methods, showing no strong preference in terms of separating structural classes for this dataset.

2.2.3. Absolute Embedding of Images The CLPCA method provides a way to classify unknown sample frames according to their 3D structures. But the ‘similarity’ used for generating the landscape is a relative concept i.e. the distributions in the embedded space depend on the sample one chooses. This limits the classification and comparison of certain sub-samples. For example, one might want to look into some subgroup in detail while still be able to relate them to the whole-sample embedding, or generate new sets of averages with bootstrapping. In addition, the time complexity of calculating the similarity matrix scales as the square of number of models, which imposes a

computational hurdle to using too many models at once.

In order to get a quicker and more universal measure of frame features, once we obtain the embedding of a typical set of averages, a neural network-based regression was used to map any previously unseen averages into the defined embedding space. Firstly, we extracted the relevant features of the patterns: for each of our 1000 averages we Fourier-transformed its azimuthal intensity variation at every radius and kept the absolute values to make the pattern invariant to in-plane rotations. We then kept frequency signals within the spatial resolution at each radius as the training features, as shown in Fig. 3(A-C). We used the coordinates of the averages in CLPCA space as training labels, as shown in Fig. 3(E).

The neural network used for fitting the relation between the pattern of 2D averages and their coordinates in CLPCA space has four fully connected hidden layers with 512, 128, 64, 32 nodes per layer. From the 1000 patterns used to calculate the similarity matrix, 800 were used to train the model, which was then validated with the rest of the 200 patterns with mean square error of 0.088, 0.059 and 0.042 for components 1, 2 and 3. With this model, we were able to quickly find the absolute embedding of any single 2D intensity model and zoom into arbitrary subsets of patterns while still retaining a reference to the full data set. An example is shown in Fig. 3(D) where we use CLPCA on a manually selected subset of averages in the “cube” region of CLPCA space. The frames which contributed to these selected averages were classified using *Dragonfly*. The embedding of these new averages is shown in red with reference to the whole-dataset CLPCA embedding.




Fig. 3. Brief illustration of the absolute embedding neural network (NN) model. (A) Average pattern from *Dragonfly*. (B) Polar representation of pattern. (C) Stack of 1D Fourier transform magnitudes along the angular axis for each radial bin. The odd frequency components (due to inversion symmetry) as well as the higher frequencies for signals at smaller radii are removed. These represent the feature vectors for the neural network. (E) Training labels from CLPCA. (D) Example of using absolute CLPCA on a selected cubic subset of frames (red dots). The gray dots represent the embedding of the pattern averages from the whole dataset.

In summary, the CLPCA method provides a robust and mostly parameter-free

framework to classify and visualize the structural landscape of an arbitrary set of coherent diffraction patterns.

2.3. 3D reconstruction of heterogeneous models

In the previous section we found a sequence where the shape of the particles transitioned from cubic to spherical. To further understand this transition, we would like to be able to reconstruct the 3D models along the sequence. Given a set of diffraction patterns from a discrete set of reproducible objects, it is relatively straightforward to generate multiple 3D models using the EMC algorithm (Ayyer *et al.*, 2021; Cho *et al.*, 2021). Here, this approach is difficult due to two reasons, 1.) it needs to assume a number of discrete heterogeneity models, which does not qualitatively capture the continuous cube-sphere transition and 2.) we do not have enough patterns located in the transition region to reconstruct 3D structures via EMC.

Fortunately, previous studies in Cryo-EM single-particle analysis provide us a good way of modeling the continuous heterogeneity. Zhong *et al.* (2021) developed cryoDRGN, a variational auto-encoder (VAE) to efficiently reconstruct heterogeneous complexes and continuous trajectories of protein motions (Kingma & Welling, 2019). Inspired by their paper, we developed a similar deep learning model by combining a VAE and convolutional neural networks (CNNs) to model the continuous shape transition along this “melting” sequence.

2.3.1. Variational Auto-encoders Figure 4(A) shows the architecture of the VAE neural network, consisting of a 2D CNN pattern-encoder to encode 2D patterns along with their orientation estimates into distributions of latent parameters, and a 3D transposed convolution network as volume-decoder to generate 3D intensity volumes from latent numbers. This setting allows the neural networks to learn the 3D-heterogeneity-

structure-encoded latent numbers from the diffraction patterns.

The inputs of the VAE network are the 2D intensity averages used as input for the CLPCA method in the previous section along with their associated relative orientations. The former are obtained from the *Dragonfly* output discussed in Section 2.2.1, while the latter are calculated as described below. We start with the 3D intensity volume of an ideal 42-nm cubic particle, slicing it with 16407 orientations uniformly distributed in quaternion space within the O_h sub-group to account for the symmetry of the object. The cross-correlation coefficient (CC) of each 2D model is calculated with all the orientations. The orientation with the highest CC is recorded for each model. In addition, another parameter which helps in training is termed *gain* referring to the ratio of the best to the average CC over all orientations. This parameter can also indicate the 'cubicness' of the particle since patterns from cubes fit very well to the synthetic model.

The VAE works in the following manner: For each input average X , its orientation, Ω and gain, \mathcal{G} , the encoder network generates a Gaussian distribution of its latent variable values $N(\mu, \sigma)$. The fact that this is not just a single latent vector z , but a distribution makes the auto-encoder variational and enforces the smoothness of the latent space. It also enables the network to use information from neighbouring regions in the space to update regions with limited data. A latent vector z is sampled using this distribution and used by the decoder to generate a 3D intensity volume V_{3d} , which is then symmetrized by the octahedral point group. The known orientation is then used to slice the generated 3D volume to get a reconstructed 2D pattern X' . The goal of training is to minimize the difference between X and X' (further details in Appendix B). 3D information is obtained since the same latent space region is sampled by averages in a variety of orientations. Note, that although the \mathcal{G} is not strictly necessary for the VAE network to reconstruction 3D volumes, it helps to

regularize the latent space of z .




Fig. 4. (A) Schematic of the variational auto-encoder (VAE). It consists of two convolutional neural networks as encoder and decoder, and the model is fit by comparing the similarity between input and decoder generated 2D slices with additional regularization by assuming the latent parameter follows a $N(0, 1)$ distribution. (B) The distribution of the bootstrapped 10000 2D average patterns in CLPCA space (gray). Red dots are selected average patterns along the melting sequence used for the VAE analysis. (C) CLPCA-2 against g for the selected patterns from (B) along the melting sequence.

When the model is trained, we can encode the 2D diffraction patterns into the latent space via the encoder network. The VAE network not only reconstructs the 3D structure of any given input 2D diffraction pattern but can also do so for any chosen location in the latent space. In this way, it can be used to study the continuous transition between cube and sphere in the melting sequence, including in regions where

there aren't sufficient patterns to isolate and obtain a conventional reconstruction.

2.3.2. Tracing the melting sequence As mentioned before, using *Dragonfly* to generate 2D averages comes at the cost of potentially averaging out structural variations in individual frames due to the limited number of classes. To generate the CLPCA landscape, 1000 averages were deemed to be enough to cover most of major shapes in the sample, but in order to trace the continuous shape variation along the melting sequence more detailed minor variation information needs to be retained. To keep more of those variation information along the melting sequence we repeat the same bootstrapping plus *Dragonfly* method described in Section 2.2.1 to generate 10,000 average patterns.

Using the CLPCA method in combination with the absolute embedding approach discussed in Section 2.2.3, we selected 5965 of these averages located along the melting sequence. Fig. 4(B) shows the distribution of these averages and the selected melting sequence averages in CLPCA space. As discussed above, the CLPCA-2 is roughly aligned with the melting sequence, which can be used as a coarse estimator of the melting process. Fig. 4(C) shows the comparison between CLPCA-2 and the gain \mathcal{G} for the selected melting sequence models. Although correlated with the melting sequence at a certain level, by visual inspection \mathcal{G} seems to be good at tracing the cubes while CLPCA-2 is better in isolating the spheres.

In order to reduce computational cost, the input 2D average intensities were down-sampled and slightly truncated at high q . Since the intensity distribution of a compact object is heavily weighted to low q , the input data was normalized by the azimuthally averaged intensity over the whole dataset. Details of the VAE network structure as well as various pre-processing steps are given in Appendix B.

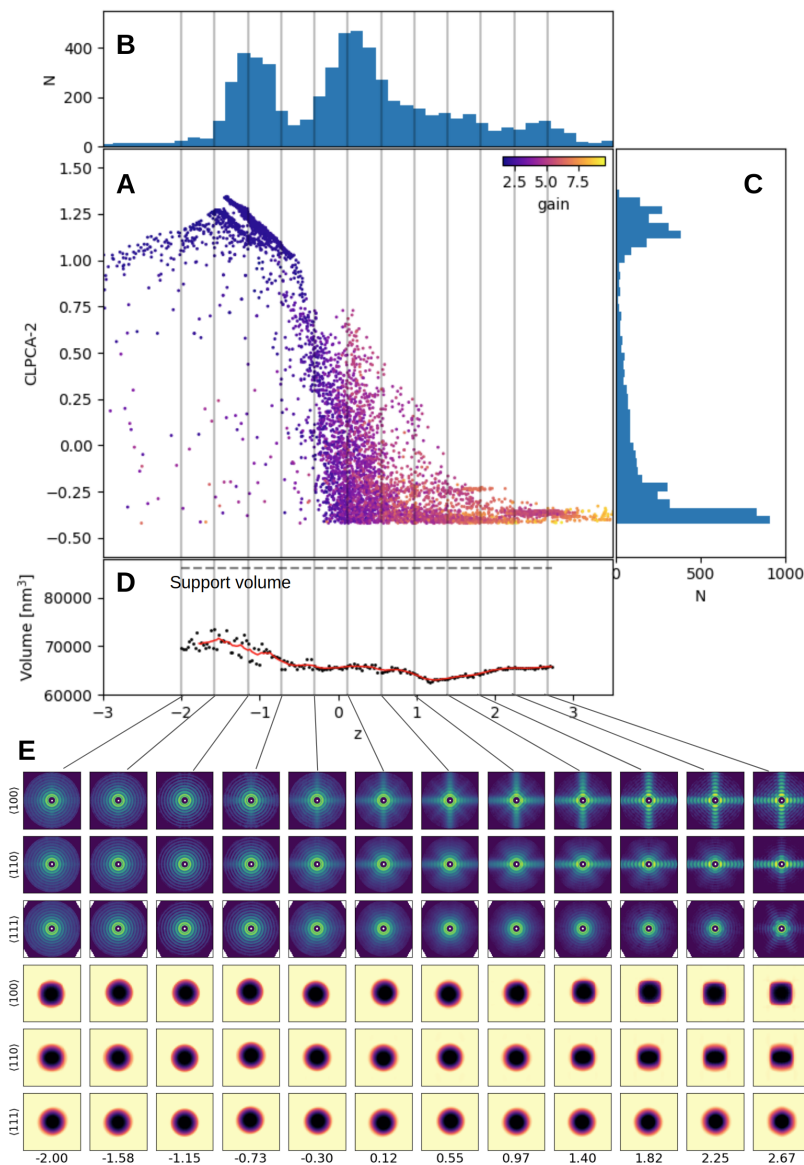


Fig. 5. (A) VAE encoded 1D latent number z against CLPCA-2 for the input sample patterns, color coding is the *gain* parameter of each pattern. (B) Histogram for latent number z . (C) Histogram for CLPCA-2. (D) Volume “evolution” along z , volumes calculated with density threshold equal to 10^{-4} total mass. Dashed horizontal line is the size of the support volume in phase retrieval. Vertical gray lines shows the location of 12 selected z numbers. (E) Top three rows are the VAE-generated intensity volumes from the 12 selected z numbers on logarithmic scale, three rows shows slices in (from row 1 to row 3) $\langle 100 \rangle$, $\langle 110 \rangle$ and $\langle 111 \rangle$ directions respectively. Bottom three rows show their corresponding density projections in real space.

For simplicity, we start by modeling with a one-dimensional latent space, so the

VAE model is forced to trace the strongest variation along the melting sequence/cubic-spherical transition. Fig. 5(A) shows the VAE encoded 1D latent parameter z against CLPCA-2 for the input models, color coded by the \mathcal{G} parameter: brighter yellow indicates the particles are more anisotropic. Fig. 5(B) and (C) are the histograms of z and CLPCA-2 respectively. Though there exists a strong correlation between latent parameter z and CLPCA-2, they show different distributions in tracing the cubic-spherical transition. First, the VAE further separates the cubic side cluster classified by CLPCA (CLPCA-2 > 0.5). Secondly, at lower than $z \approx -1.5$ another sequence is visible which is not clearly seen in CLPCA space.

One advantage of the VAE network is that it allows us to reconstruct/generate the 3D intensity volume for any given z . Here, we selected 12 z numbers of equal steps between -2.00 and +2.67. In Fig. 5(E), the upper three rows show three special ($\langle 100 \rangle$, $\langle 110 \rangle$ and $\langle 111 \rangle$) slices of the VAE-generated 12 intensity volumes on logarithmic scale. We see a clear and smooth cubic-spherical transition from $z \approx 2.7 \rightarrow -1.5$, which shows that the VAE network is able to trace the early melting stages much better than the CLPCA. And for the second ‘sequence’ at $z < 1.5$, the lower contrast in the intensities and less spherical structures indicate they are likely to be the contaminated particles at late melting stages, the low contrast being due to the fact that they contain various shapes and lack of accurate orientation estimates. Compared with the CLPCA method that roughly classifies cubic and spherical particles, the VAE is able to provide a detailed, smooth modeling of the entire transition process.

The bottom three rows of Fig. 5(E) show projections of phase-retrieved density maps in real space along the same three directions. We also find the volume increases by around 10% from cubes to spheres along the melting sequence by calculating the volume evolution of 3D reconstructed volumes at finely sampled 200 different z -values as black dots shown in Fig. 5(D). This is consistent with the density difference between

crystalline gold (19.32 g/cc) and molten/randomly close packed atoms (17.31 g/cc). This is also in agreement with the direct size fitting on 2D average patterns of spherical and cubic particles, shown in the Fig. 9.

This volume analysis allows us to observe an additional feature, namely that the melting sequence seems to show two different stages: Stage 1 from $z \approx 3 \rightarrow -0.3$, where most of the shape change occurs, but with no significant volume changes. Stage 2 from $z \approx -0.3 \rightarrow -1.5$ where all particles are mostly round but the volume increases. This is to be expected since edges and vertices require lower energies to be disrupted than for complete melting, where the bulk lattice structure is lost (Cahn, 1986; Chen *et al.*, 2021). As discussed below, the melting sequence here is only based on the 3D structure of particles. Thus, the “two stages” could only suggest that the shape changing happened “before” the size expansion along z (exposure intensity).

We note that the melting sequence here is purely based on the 3D structure of particles. Since we observe a continuous transition, we assume a monotonic relationship between the latent variable z and the incident fluence in the previous pulse. However, without additional information about the expected distribution of particles with a given incident fluence, we cannot obtain the precise mapping between the two. As all pre-exposures happened 880 ns before the imaging pulse, one should not view the melting sequence as a time-dependent shape variation, but rather an incident fluence- or temperature-dependent behavior.

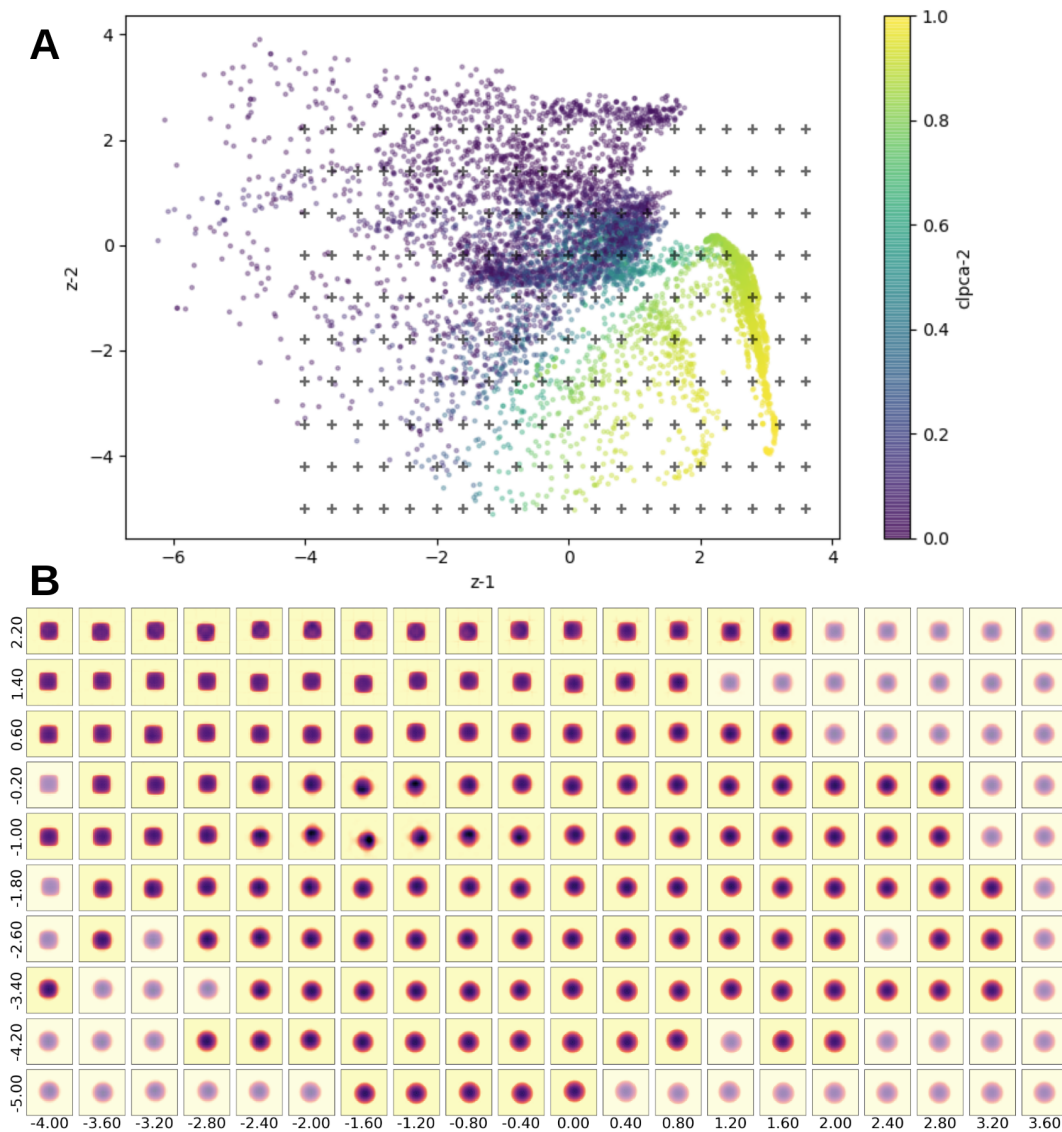


Fig. 6. (A) The distribution of the dataset in 2D latent space, color coded by CLPCA-2. Black crosses in the plot are select 200 latent numbers z . (B) The $\langle 100 \rangle$ direction projection of 200 real space density retrieved from the logarithmic intensity volumes generated from the 200 selected latent numbers z (black crosses in (A)). Transparent slices are volumes generated from latent regions without data points.

We have shown that the one dimensional latent space VAE network is capable of modeling the major transition, but it is not the only variation in the dataset. To trace

more variations we now model it with a higher dimensional latent space. Fig. 6(A) shows the VAE-encoded 2D latent parameters of the sample. Neither the first latent nor the second latent parameter traces the cubic - spherical transition independently. This shows the encoding of more secondary variations, e.g. contaminant, size etc.

We selected 200 evenly distributed points in the latent space and generate their 3D intensity volumes. Fig. 6(B) shows the $\langle 100 \rangle$ direction projection maps of the phase retrieved VAE generated intensity volumes of the 200 selected latent numbers z . One can clearly notice the transition between spherical to cubic features from right to the left is roughly encoded in the first z component. However, the trend is twisted in 2D space, we got multiple transition sequences along the second component. In regions with no or very few data points (half transparent slices), the results are less reliable as the VAE network could not generate reliable intensity volumes.

3. Conclusion

In this work, we demonstrate two methods to study heterogeneous ensembles of samples using X-ray single particle imaging by analyzing the relationships between the 3D structures of the samples rather than the diffraction patterns directly. The first, termed CLPCA, uses the fact that diffraction patterns from ideal particles share a common line/arc from the Ewald construction. The similarity of the best common lines allows us to visualize a low-dimensional embedding of the dataset which roughly represents 3D similarity regardless of orientation. A fully-connected neural network was trained against the embedding in order to embed arbitrary 2D averages with respect to the rest of the data set.

This was applied to an experimental dataset from gold nanoparticles, containing 2.4 million patterns with only a few thousand photons per frame. In order to improve the signal-to-noise ratio and fill in detector panel gaps, 2D classification in *Dragonfly*

was used to combine similar patterns up to in-plane rotations, and a bootstrapping approach was used to increase the number of 2D averages with subsets of the data.

The second method involves a generative VAE network that models continuous structural deformations, and can be used to generate 3D structures at arbitrary points along the landscape. In the studied dataset, this network was able to recover a continuous melting sequence induced by pre-exposure of the particles by the previous pulse in the European XFEL pulse train. We observe an initial rounding of the particles followed by a slight expansion caused by atomic-scale disordering at higher temperatures. Our two methods open up the possibility of studying samples of heterogeneous particles, and potentially tracing dynamic motions of particles in the SPI experiments.

Appendix A Different embeddings of CC-matrix

In Fig. 7 we show a comparison between PCA and other standard embedding methods implemented in the `scikit-learn` Python package (Pedregosa *et al.*, 2011). For the sample set considered, all four embeddings are able to separate the 4 basic shape groups (cube, sphere, transition and contaminant), although this may not be true for other, more complex continuous structural variations.

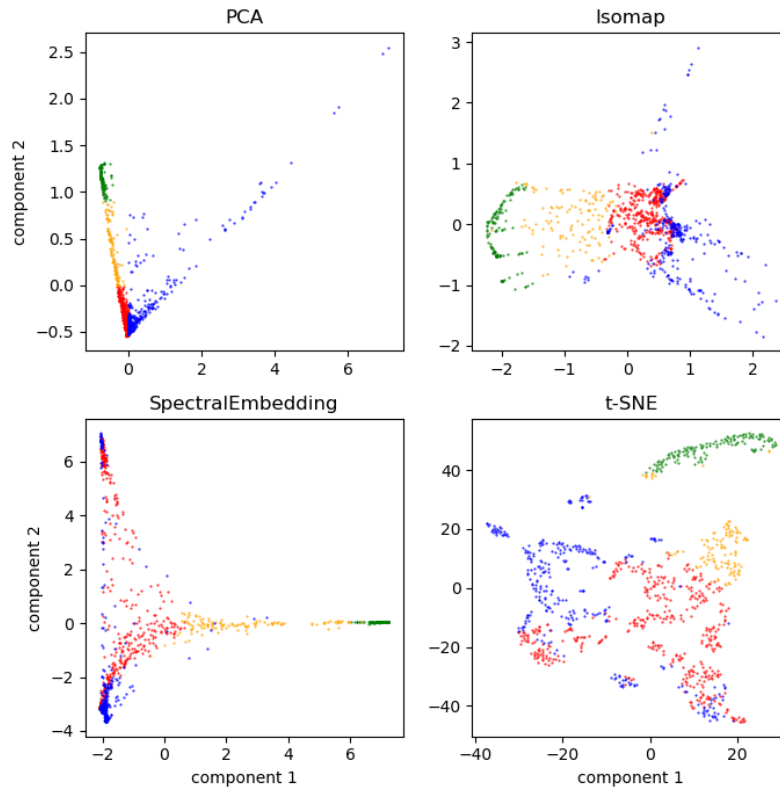


Fig. 7. Comparison between four different common embedding methods, top left: PCA, top right: Isomap, bottom left: Spectral, and bottom right: t-SNE. Color coding is same as in Fig. 2 by manually divided groups from the first two components of PCA space.

Appendix B Details of the VAE model

B.1. Preprocessing the input patterns

The 2D averages from *Dragonfly* have an initial size of 441×441 pixels which is computationally redundant considering the data is highly oversampled. For the results shown in this paper, we used the following procedures to reduce the size. We firstly

downsampled the size to 243×243 and then we cut off the high q part to further reduce the size to 161×161 .

In addition to the size reduction, diffraction patterns of compactly supported objects are overwhelmingly dominated by the low q signal. To appropriately weight the higher q shape information, we divided the 2D intensities by the radial average intensity over the whole dataset before feeding them into the VAE and multiplied it back when generating the 3D intensities.

B.2. Model parameters

The depth and nodes of our encoder/decoder network depend on the quality of the sample and the spatial resolution of reconstructed 3D volumes we need. In our code we provide 3 different reconstructing volume sizes options: low volume (81^3), intermediate volume (161^3) and high volume (243^3). For the results shown in this paper, the intermediate volume model is used in which the encoder network has four 2D convolution layers with (16×81^2) , (32×27^2) , (64×9^2) and (128×9^2) nodes and (3×3) kernel size per each layer. The decoder consists of six 3D transposed convolution layers with (128×9^3) , (64×9^3) , (32×27^3) , (16×81^3) , (8×161^3) and (1×161^3) nodes and $(3 \times 3 \times 3)$ kernel size per each layer.

In Fig. 8 we show a comparison of input 2D patterns and their corresponding 2D slices of their reconstructed 3D intensities at four representative latent z -values from the model shown in Section 2.3.2. Together with Fig. 5 the first pattern is likely from half-melted contaminant particles, where the reconstruction is not reliable because the number of patterns of that shape are low and the orientation estimate by fitting against a cube is inaccurate. The other rows are well reconstructed spheres, melting cubes and cubes.




Fig. 8. Comparison between input intensities and reconstructed intensity patterns on logarithmic scale.

B.3. Devices

The network was implemented in PyTorch and run on a single node with 4 NVidia V100 GPUs with 32 GB memory each. For our 5672 2D averages, the 161^3 volume model takes less than 2 hours to calculate.

Appendix C

Direct size comparison of cubes and spheres




Fig. 9. Volume distribution of cubic particles (orange) and spherical particles (blue).

Figure 9 shows the volume comparison of classified pure cubes and pure spheres. The cube sizes were calculated directly from frames with clear streak fringes, which were assumed to have at least one set of faces parallel to the X-ray beam. The larger spread of cube sizes can be accounted for by the inclusion of frames where the faces were not exactly parallel to the beam. Nevertheless, similar to the VAE model, a clear size increase is observed from cubes to spheres.

We acknowledge European XFEL in Schenefeld, Germany, for provision of X-ray free-electron laser beamtime at Scientific Instrument SPB/SFX (Single Particles, Clusters, and Biomolecules and Serial Femtosecond Crystallography) and would like to thank the staff for their assistance. This work has been supported by the Clusters of Excellence ‘Center for Ultrafast Imaging’ (CUI,EXC 1074, ID 194651731) and ‘Advanced Imaging of Matter’ (AIM,EXC 2056, ID 390715994) of the Deutsche Forschungsgemeinschaft (DFG). This work has also been supported by the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) through the Consolidator Grant COMOTION (ERC-614507-Küpper)

and by the Helmholtz Gemeinschaft through the “Impuls-und Vernetzungsfond”. J.C.H.S. and R.A.K. acknowledge support from the National Science Foundation BioXFEL award (STC-1231306). F.R.N.C.M. acknowledges support from the Swedish Research Council, Röntgen-Ångström Cluster and Carl Tryggers Foundation for Scientific Research. P.L.X. acknowledges a fellowship from the Joachim Herz Stiftung. P.L.X. and H.N.C. acknowledge support from the Human Frontiers Science Program (RGP0010/2017).

References

- Allahgholi, A., Becker, J., Delfs, A., Dinapoli, R., Goettlicher, P., Greiffenberg, D., Henrich, B., Hirsemann, H., Kuhn, M., Klanner, R., Klyuev, A., Krueger, H., Lange, S., Laurus, T., Marras, A., Mezza, D., Mozzanica, A., Niemann, M., Poehlsen, J., Schwandt, J., Sheviakov, I., Shi, X., Smoljanin, S., Steffen, L., Sztuk-Dambietz, J., Trunk, U., Xia, Q., Zeribi, M., Zhang, J., Zimmer, M., Schmitt, B. & Graafsma, H. (2019). *Journal of Synchrotron Radiation*, **26**(1), 74–82.
URL: <https://doi.org/10.1107/s1600577518016077>
- von Ardenne, B., Mechelke, M. & Grubmüller, H. (2018). *Nature Communications*, **9**(1).
URL: <https://doi.org/10.1038/s41467-018-04830-4>
- Ayyer, K., Lan, T.-Y., Elser, V. & Loh, N. D. (2016). *Journal of Applied Crystallography*, **49**(4), 1320–1335.
URL: <https://doi.org/10.1107/S1600576716008165>
- Ayyer, K., Morgan, A. J., Aquila, A., DeMirici, H., Hogue, B. G., Kirian, R. A., Xavier, P. L., Yoon, C. H., Chapman, H. N. & Barty, A. (2019). *Optics Express*, **27**(26), 37816.
URL: <https://doi.org/10.1364/oe.27.037816>
- Ayyer, K., Xavier, P. L., Bielecki, J., Shen, Z., Daurer, B. J., Samanta, A. K., Awel, S., Bean, R., Barty, A., Bergemann, M., Ekeberg, T., Estillore, A. D., Fangohr, H., Giewekemeyer, K., Hunter, M. S., Karnevskiy, M., Kirian, R. A., Kirkwood, H., Kim, Y., Koliyadu, J., Lange, H., Letrun, R., Lübke, J., Michelat, T., Morgan, A. J., Roth, N., Sato, T., Sikorski, M., Schulz, F., Spence, J. C. H., Vagovic, P., Wöllweber, T., Worbs, L., Yefanov, O., Zhuang, Y., Maia, F. R. N. C., Horke, D. A., Küpper, J., Loh, N. D., Mancuso, A. P. & Chapman, H. N. (2021). *Optica*, **8**(1), 15–23.
URL: <http://www.osapublishing.org/optica/abstract.cfm?URI=optica-8-1-15>
- Bobkov, S. A., Teslyuk, A. B., Kurta, R. P., Gorobtsov, O. Y., Yefanov, O. M., Ilyin, V. A., Senin, R. A. & Vartanyants, I. A. (2015). *Journal of Synchrotron Radiation*, **22**(6), 1345–1352.
URL: <https://doi.org/10.1107/s1600577515017348>
- Cahn, R. W. (1986). *Nature*, **323**(6090), 668–669.
- Chen, J., Fan, X., Liu, J., Gu, C., Shi, Y., Zheng, W. & Singh, D. J. (2021). *The Journal of Physical Chemistry Letters*, **12**, 8170–8177.
- Cho, D. H., Shen, Z., Ihm, Y., Wi, D. H., Jung, C., Nam, D., Kim, S., Park, S.-Y., Kim, K. S., Sung, D., Lee, H., Shin, J.-Y., Hwang, J., Lee, S. Y., Lee, S. Y., Han, S. W., Noh, D. Y., Loh, N. D. & Song, C. (2021). *ACS Nano*, **15**(3), 4066–4076.
URL: <https://doi.org/10.1021/acsnano.0c07961>
- Decking, W., Abeghyan, S., Abramian, P. *et al.* (2020). *Nature Photonics*, **14**(6), 391–397.
URL: <https://doi.org/10.1038/s41566-020-0607-z>
- Ekeberg, T., Svenda, M., Abergel, C., Maia, F. R., Seltzer, V., Claverie, J.-M., Hantke, M., Jönsson, O., Nettelblad, C., van der Schot, G., Liang, M., DePonte, D. P., Barty, A., Seibert, M. M., Iwan, B., Andersson, I., Loh, N. D., Martin, A. V., Chapman, H., Bostedt, C., Bozek, J. D., Ferguson, K. R., Krzywinski, J., Epp, S. W., Rolles, D., Rudenko, A., Hartmann, R., Kimmel, N. & Hajdu, J. (2015). *Physical Review Letters*, **114**(9).
URL: <https://doi.org/10.1103/physrevlett.114.098102>

- Ignatenko, A., Assalauova, D., Bobkov, S. A., Gelisio, L., Teslyuk, A. B., Ilyin, V. A. & Vartanyants, I. A. (2021). *Machine Learning: Science and Technology*, **2**(2), 025014.
URL: <https://doi.org/10.1088/2632-2153/abd916>
- Kingma, D. P. & Welling, M. (2019). *Foundations and Trends® in Machine Learning*, **12**(4), 307–392.
URL: <https://doi.org/10.1561/22000000056>
- Kurta, R. P., Donatelli, J. J., Yoon, C. H., Berntsen, P., Bielecki, J., Daurer, B. J., DeMirici, H., Fromme, P., Hantke, M. F., Maia, F. R., Munke, A., Nettelblad, C., Pande, K., Reddy, H. K., Sellberg, J. A., Sierra, R. G., Svenda, M., van der Schot, G., Vartanyants, I. A., Williams, G. J., Xavier, P. L., Aquila, A., Zwart, P. H. & Mancuso, A. P. (2017). *Physical Review Letters*, **119**(15).
URL: <https://doi.org/10.1103/physrevlett.119.158102>
- Loh, N.-T. D. & Elser, V. (2009). *Physical Review E*, **80**(2).
URL: <https://doi.org/10.1103/physreve.80.026705>
- Mancuso, A. P., Aquila, A., Batchelor, L., Bean, R. J., Bielecki, J., Borchers, G., Doerner, K., Giewekemeyer, K., Graceffa, R., Kelsey, O. D., Kim, Y., Kirkwood, H. J., Legrand, A., Letrun, R., Manning, B., Morillo, L. L., Messerschmidt, M., Mills, G., Raabe, S., Reimers, N., Round, A., Sato, T., Schulz, J., Takem, C. S., Sikorski, M., Stern, S., Thute, P., Vagovič, P., Weinhausen, B. & Tschentscher, T. (2019). *Journal of Synchrotron Radiation*, **26**(3), 660–676.
URL: <https://doi.org/10.1107/s1600577519003308>
- Neutze, R., Wouts, R., van der Spoel, D., Weckert, E. & Hajdu, J. (2000). *Nature*, **406**(6797), 752–757.
URL: <https://doi.org/10.1038/35021099>
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V. *et al.* (2011). *the Journal of machine Learning research*, **12**, 2825–2830.
- Rose, B. C., Huang, D., Zhang, Z.-H., Stevenson, P., Tyryshkin, A. M., Sangtawesin, S., Srinivasan, S., Loudin, L., Markham, M. L., Edmonds, A. M., Twitchen, D. J., Lyon, S. A. & de Leon, N. P. (2018). *Science*, **361**(6397), 60–63.
URL: <https://doi.org/10.1126/science.aao0290>
- Shi, X., Levine, E., Weber, H. & Hargreaves, B. A. (2018). *Magnetic Resonance in Medicine*, **81**(4), 2247–2263.
URL: <https://doi.org/10.1002/mrm.27536>
- Shneerson, V. L., Ourmazd, A. & Saldin, D. K. (2008). *Acta Crystallographica Section A Foundations of Crystallography*, **64**(2), 303–315.
URL: <https://doi.org/10.1107/s0108767307067621>
- Singer, A. & Shkolnisky, Y. (2011). *SIAM journal on imaging sciences*, **4**(2), 543–572.
- Yefanov, O. & Vartanyants, I. (2013). *Journal of Physics B: Atomic, Molecular and Optical Physics*, **46**(16), 164013.
- Yoon, C. H., Schwander, P., Abergel, C., Andersson, I., Andreasson, J., Aquila, A., Bajt, S., Barthelmeß, M., Barty, A., Bogan, M. J. *et al.* (2011). *Optics express*, **19**(17), 16542–16549.
- Zhong, E. D., Bepler, T., Berger, B. & Davis, J. H. (2021). *Nature Methods*, **18**(2), 176–185.
URL: <https://doi.org/10.1038/s41592-020-01049-4>
- Zimmermann, J., Langbehn, B., Cucini, R., Fraia, M. D., Finetti, P., LaForge, A. C., Nishiyama, T., Ovcharenko, Y., Piseri, P., Plekan, O., Prince, K. C., Stienkemeier, F., Ueda, K., Callegari, C., Möller, T. & Rupp, D. (2019). *Physical Review E*, **99**(6).
URL: <https://doi.org/10.1103/physreve.99.063309>

Synopsis

Two unsupervised machine learning approaches are presented to characterize structural variations in nanoparticle ensembles measured using X-ray single particle imaging. The algorithms are applied to an experimental dataset where both discrete structural classes as well as continuous deformations caused by X-ray induced melting are classified.
