

## Continuous Structural Heterogeneity Detection in High Resolution Microscopy

Haghparast, S.

**DOI**

[10.4233/uuid:50fd45fd-825e-499a-848c-776f8b0e0088](https://doi.org/10.4233/uuid:50fd45fd-825e-499a-848c-776f8b0e0088)

**Publication date**

2025

**Document Version**

Final published version

**Citation (APA)**

Haghparast, S. (2025). *Continuous Structural Heterogeneity Detection in High Resolution Microscopy*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:50fd45fd-825e-499a-848c-776f8b0e0088>

**Important note**

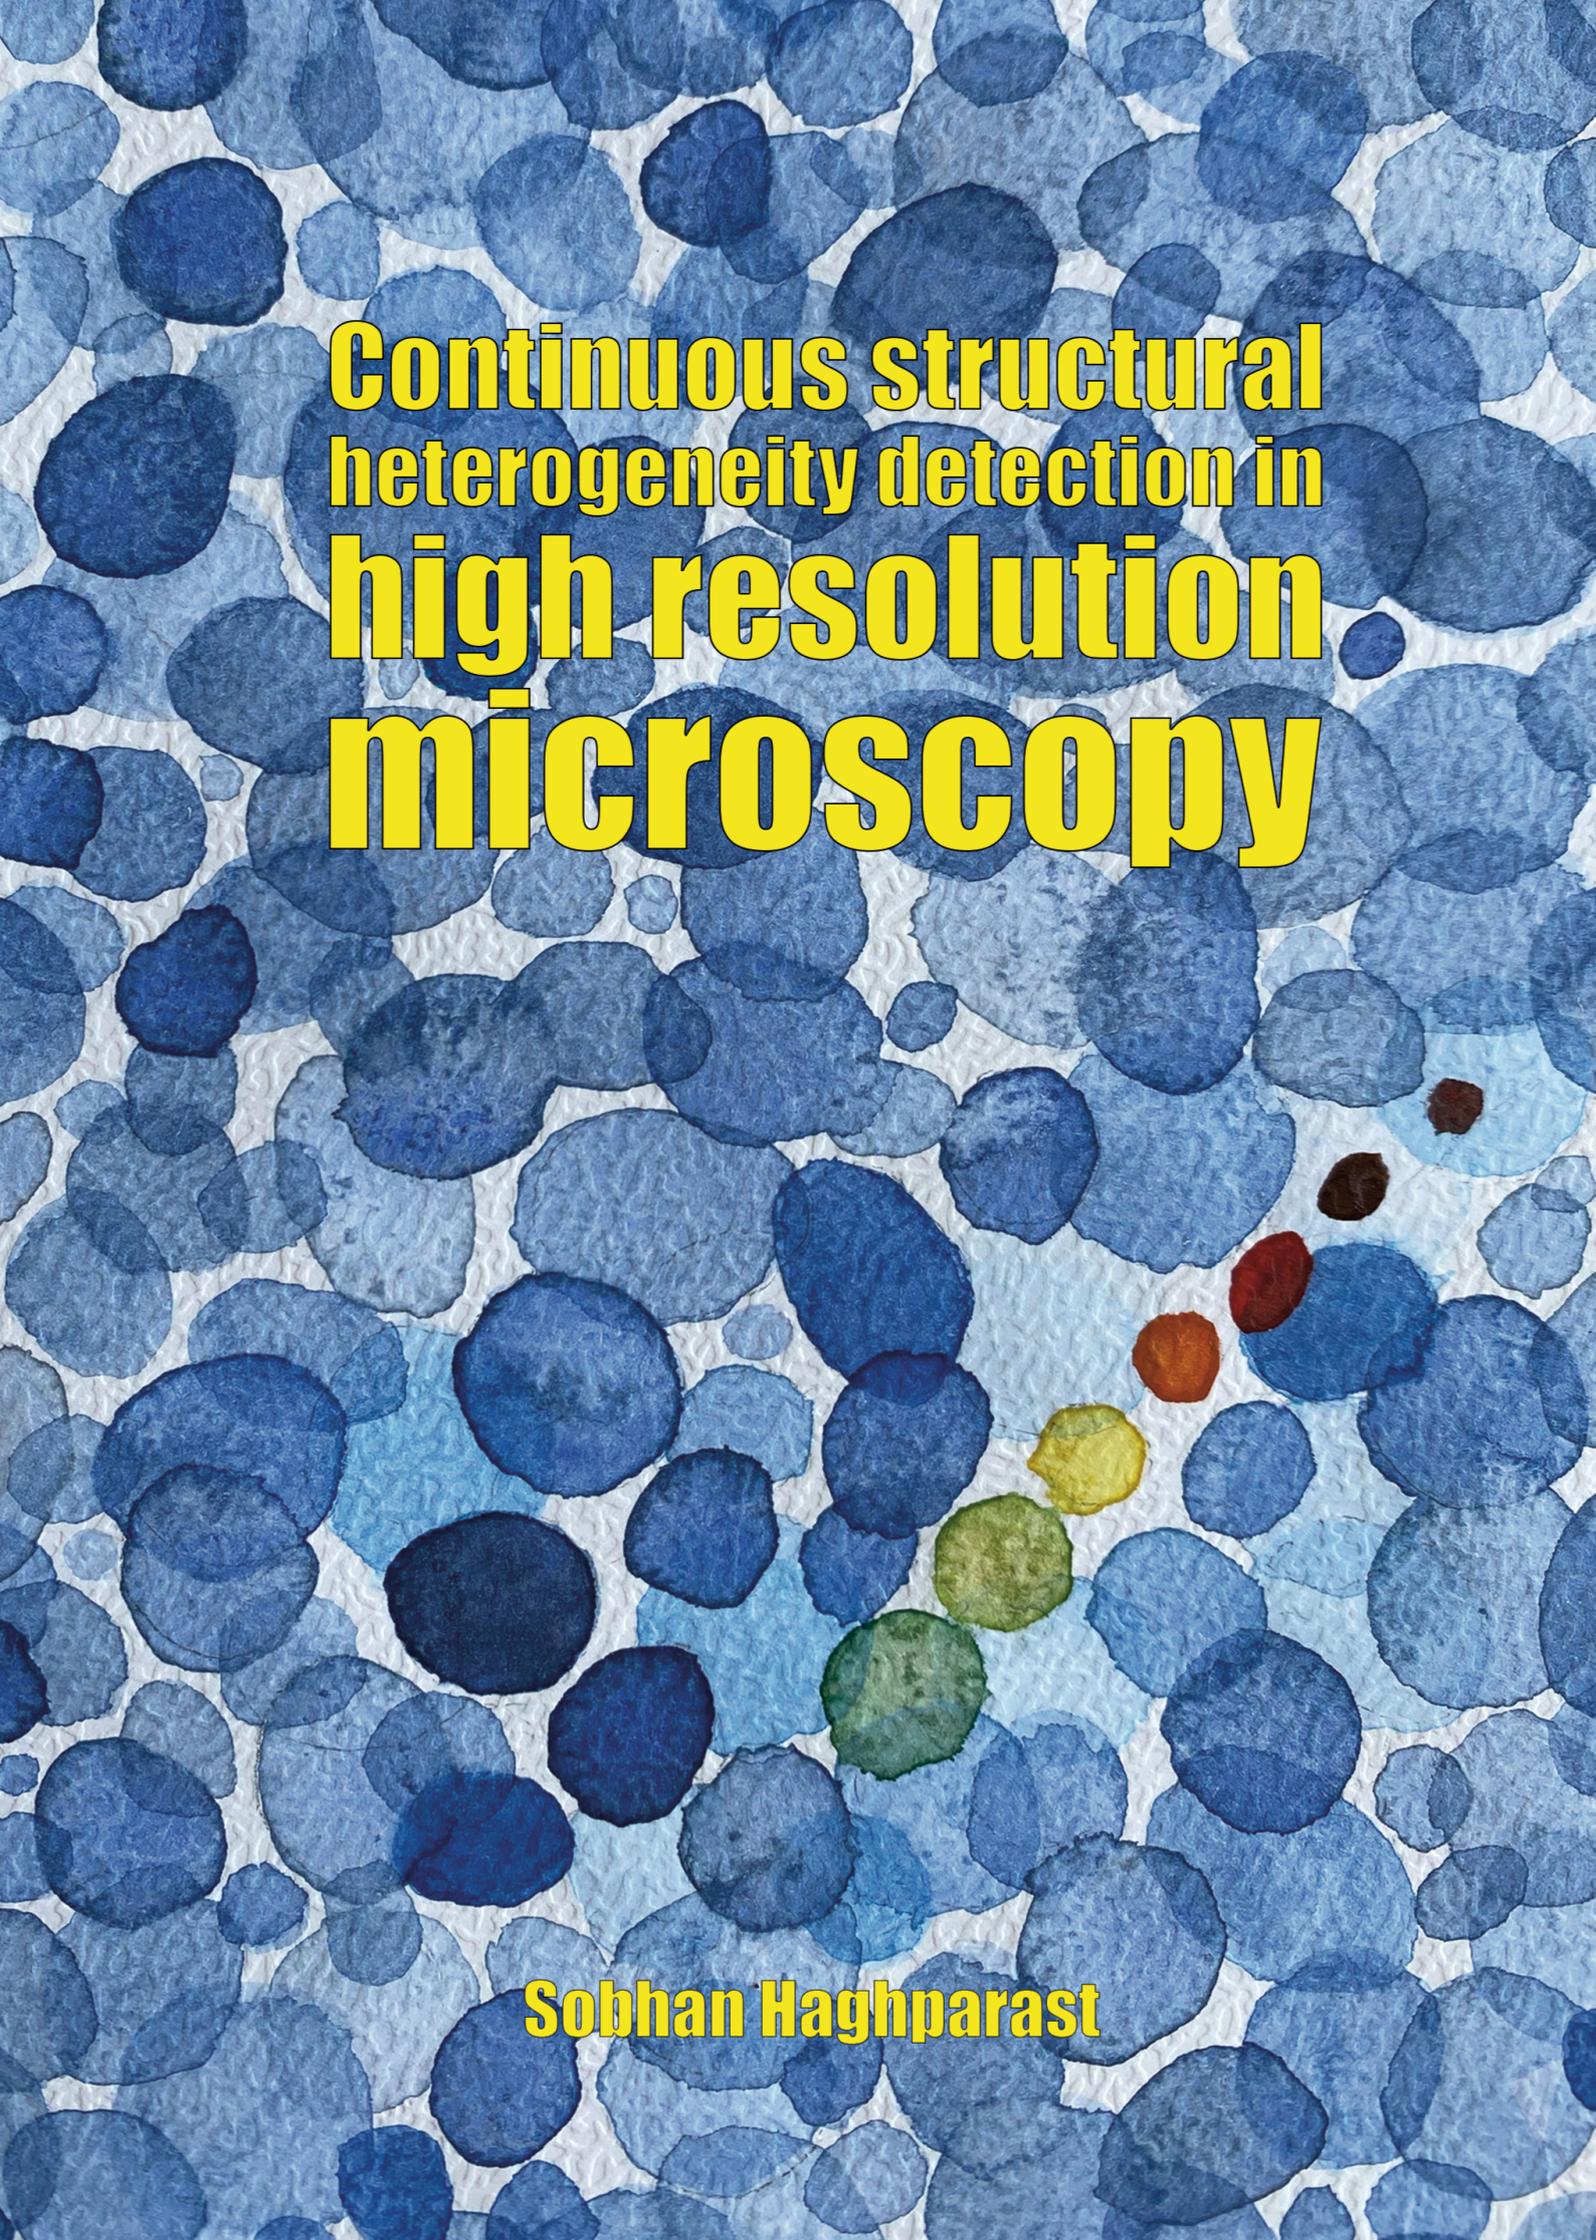
To cite this publication, please use the final published version (if applicable). Please check the document version above.

**Copyright**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

A microscopic image of cells, likely from a tissue section, showing a dense population of cells. The cells are stained in various shades of blue, with some cells appearing darker or lighter. A few cells are highlighted in different colors: a small red cell, a small orange cell, a small yellow cell, and a small green cell. The text is overlaid on the image in a bold, yellow, sans-serif font.

**Continuous structural  
heterogeneity detection in  
high resolution  
microscopy**

**Sobhan Haghparast**

**CONTINUOUS STRUCTURAL HETEROGENEITY  
DETECTION IN HIGH RESOLUTION MICROSCOPY**



# **CONTINUOUS STRUCTURAL HETEROGENEITY DETECTION IN HIGH RESOLUTION MICROSCOPY**

## **Dissertation**

for the purpose of obtaining the degree of doctor  
at Delft University of Technology  
by the authority of the Rector Magnificus, Prof. Dr.ir. T.H.J.J. van der Hagen,  
chair of the Board for Doctorates  
to be deffended publicly on Tuesday 30, September 2025 at 17:30 o'clock.

by

**Sobhan HAGHPARAST**

Master of science in Electrical Engineering (Biomedical Imaging),  
Sharif University of Technology, Tehran, Iran,  
born in Sari, Iran.

This dissertation has been approved by the promoters.

Composition of the doctoral committee:

Rector Magnificus, chairperson  
Prof. dr. B. Rieger, TNW-IMPHYS-RIETU Delft, NL *promotor*  
Prof. dr. S. Stallinga, TNW-IMPHYS-STATU Delft, NL, *promotor*

*independent members:*

Prof.dr. S. Cox KCL, UK  
Prof.dr.ir.R.L.Lagendijk EWI-INSY-CYSTU Delft, NL  
Prof.dr. K.A. Lidke UNM, USA  
Dr.F.M. Vos TNW-IMPHYS-VOSTU Delft, NL / Erasmus U., NL

*Reserve member:*

Prof.dr.ir. J.P.Hoogenboom TNW-IMPHYSTU Delft, NL

This work is part of the research program supported by the Dutch Research council (NWO), VICI grant no. 17046.



*Keywords:* Continuous Structural Heterogeneity Detection, Cryogenic Electron Microscopy (Cryo-EM), Single-Molecule Localization Microscopy (SMLM), Particle Fusion, Biological Structure Analysis

*Printed by:* Ipskamp Printing

*Front & Back:* Inspired by watercolor droplets in delft blue, resembling cells with heterogeneous structures. Designed by Zeinab Yousefnezhad.

Copyright © 2025 by Sobhan Haghparast

ISBN 978-94-6473-924-4

An electronic version of this dissertation is available at  
<http://repository.tudelft.nl/>.

# CONTENTS

<b>Summary</b>	<b>ix</b>
<b>Samenvatting</b>	<b>xi</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Fluorescence microscopy . . . . .	2
1.2 Super-resolution microscopy . . . . .	3
1.2.1 Resolution assessment in super-resolution microscopy . . . . .	3
1.3 Single molecule localization microscopy . . . . .	3
1.3.1 SMLM data processing pipeline . . . . .	5
1.4 Cryo-Electron microscopy (cryo-EM) . . . . .	7
1.4.1 Single Particle Analysis (SPA) . . . . .	8
1.5 Particle fusion in SMLM. . . . .	10
1.5.1 EM inspired particle fusion . . . . .	10
1.5.2 Template-based particle fusion model based fitting . . . . .	10
1.5.3 Template free particle fusion. . . . .	11
1.6 Motivation and outline . . . . .	11
1.6.1 Outline. . . . .	12
References . . . . .	13
<b>2 Detecting continuous structural heterogeneity in single-molecule localization microscopy data</b>	<b>19</b>
2.1 Introduction . . . . .	20
2.2 Methods . . . . .	21
2.2.1 Pairwise registration. . . . .	21
2.2.2 Multi-Dimensional Scaling (MDS) space. . . . .	22
2.2.3 Isomap. . . . .	22
2.2.4 Principal Component Analysis (PCA). . . . .	23
2.2.5 Reconstruction per bin. . . . .	23
2.2.6 Particle fusion. . . . .	23
2.2.7 Model-based shape parameter estimation. . . . .	23
2.2.8 Data acquisition of 2D NPC. . . . .	24
2.3 Results . . . . .	24
2.3.1 Continuous distribution of 3D DNA-origami tetrahedron height. . . . .	24
2.3.2 Continuous distribution of 2D NPC radius. . . . .	25
2.3.3 Simulation on impact of DOL. . . . .	27
2.3.4 Simulation on two modes of variation. . . . .	29
2.4 Discussion . . . . .	29
References . . . . .	31

---

<b>3</b>	<b>Detecting Continuous Structural Heterogeneity in SMLM Data with a Point Cloud Variational Auto-Encoder</b>	<b>35</b>
3.1	Introduction	36
3.2	Methods	37
3.2.1	Point Cloud Variational Auto-Encoder framework (PC-VAE)	37
3.2.2	Loss Function	39
3.2.3	Training Details	39
3.2.4	Binning and Visualization	41
3.2.5	Data Description.	41
3.2.6	Model based parameter estimation	41
3.3	Results	42
3.3.1	2D NPC dataset	42
3.3.2	3D tetrahedron dataset	42
3.3.3	3D NPC dataset.	46
3.4	Discussion and outlook	50
	References	51
<b>4</b>	<b>Continuous structural heterogeneity detection in Cryo-Electron Microscopy</b>	<b>55</b>
4.1	Introduction	56
4.2	Methods	57
4.2.1	Image simulation	59
4.2.2	Image processing in RELION.	59
4.2.3	Assigning orientation class.	59
4.2.4	Pairwise-distance matrix.	59
4.2.5	MDS embedding.	60
4.2.6	Polarity matching	61
4.2.7	Latent space binning and reconstruction	61
4.3	Results	62
4.4	Discussion & outlook	67
	References	71
<b>5</b>	<b>Conclusion</b>	<b>75</b>
5.1	Conclusion	76
5.1.1	Detecting continuous structural heterogeneity in single-molecule localization microscopy data.	76
5.1.2	Detecting Continuous Structural Heterogeneity in Single Molecule Localization Microscopy Data with a Point Cloud Variational Auto-Encoder	76
5.1.3	Continuous structural heterogeneity detection in Cryo-EM data	77
5.2	Outlook	78
5.2.1	Detecting continuous structural heterogeneity in single-molecule localization microscopy data.	78
5.2.2	Detecting Continuous Structural Heterogeneity in Single Molecule Localization Microscopy Data with a Point Cloud Variational Auto-Encoder	78
5.2.3	Continuous structural heterogeneity detection in Cryo-EM data.	79

---

References . . . . .	80
<b>Acknowledgements</b>	<b>83</b>
<b>Curriculum Vitæ</b>	<b>87</b>



# SUMMARY

High-resolution microscopy techniques, such as Single-Molecule Localization Microscopy (SMLM) and Cryogenic Electron Microscopy (Cryo-EM), can utilize particle fusion or averaging to reconstruct a macromolecular structure of increased signal-to-noise ratio and of potentially higher resolution. This process assumes that all fused particles are structurally equal. Structural heterogeneity, however, is often present due to biological variations and should not be ignored. In particularly continuous and subtle conformational changes present in the data lead to undesired blurring of the reconstruction. This thesis develops methods to detect continuous structural heterogeneity and to exploit it for more faithful reconstructions, enabling more accurate interpretations and insights into molecular structures.

In Chapter 2, we propose a method to detect continuous structural heterogeneity in SMLM datasets based on an all-to-all pairwise comparison of the found structures. The method is applied to both experimental and simulated data, where continuous variations such as the height of 3D DNA origami tetrahedrons and the radius of 2D Nuclear Pore Complexes (NPCs) are detected. The chapter highlights how accounting for these structural variations leads to more reliable particle fusion and reconstruction.

In Chapter 3, we propose a Point Cloud Variational Auto-Encoder (PCVAE) that operates directly on 2D and 3D localization data to detect structural heterogeneity. Unlike common neural networks that rely on pixelated images, our method utilizes raw localization coordinates. This not only reduces the required memory but also has low computational complexity and thus allows scalability to many structures. In contrast to multi-dimensional scaling approaches, where the computational complexity scales quadratically, here it remains linear with the number of particles. Our method is capable of identifying multiple modes of variation and reveals nanometer-scale changes such as radius and height variations in both simulated and experimental datasets.

In Chapter 4, we propose a method to detect continuous structural heterogeneity in Cryo-EM datasets. Recent approaches rely on machine learning models that often require large training datasets and careful tuning of hyperparameters. In contrast, our method detects underlying continuous variations in 2D projections by pairwise comparison of images within orientation classes. The approach reconstructs intermediate conformational states representing the continuous structural heterogeneity in synthetic SARS-CoV-2 spike protein data, simulated under ideal conditions. More realistic simulations, incorporating varying defocus per particle and radiation damage, do not lead to the same favourable results, still posing a challenge for future research.



# SAMENVATTING

Hoge-resolutiemicroscopietechnieken, zoals enkel molecule lokalisatiemicroscopie (SMLM) en cryogene elektronenmicroscopie (Cryo-EM), kunnen deeltjesfusie of -middeling gebruiken om een macromoleculaire structuur met grotere signaal-ruisverhouding en een potentieel hogere resolutie te reconstrueren. Dit proces veronderstelt dat alle gefuseerde deeltjes structureel gelijk zijn. Structurele heterogeniteit is echter vaak aanwezig als gevolg van biologische variaties en mag niet worden genegeerd. Vooral continue en subtiele conformatieveranderingen in de data leiden tot ongewenste vervaging van de reconstructie. Dit proefschrift ontwikkelt methoden om continue structurele heterogeniteit te detecteren en te benutten voor getrouwere reconstructies, wat nauwkeurigere interpretaties en inzichten in moleculaire structuren mogelijk maakt.

In hoofdstuk 2 stellen we een methode voor om continue structurele heterogeniteit in SMLM datasets te detecteren op basis van een paarsgewijze vergelijking van de gevonden structuren over alle mogelijke paren. De methode wordt toegepast op zowel experimentele als gesimuleerde data, waarbij continue variaties zoals de hoogte van 3D DNA-origami-tetraëders en de straal van 2D nucleaire poriecomplexen (NPC's) worden gedetecteerd. Dit hoofdstuk belicht hoe het in aanmerking nemen van deze structurele variaties leidt tot betrouwbaardere deeltjesfusie en reconstructie.

In hoofdstuk 3 stellen we een puntwolk gebaseerde variationele auto-encoder (PCVAE) voor die direct op 2D en 3D localisatiedata werkt om structurele heterogeniteit te detecteren. In tegenstelling tot gangbare neurale netwerken die afhankelijk zijn van gepixelde afbeeldingen, maakt onze methode gebruik van ruwe localisatiecoördinaten. Dit vermindert niet alleen het benodigde computergeheugen, maar heeft ook een lage rekencomplexiteit en maakt schaalbaarheid naar vele structuren mogelijk. In tegenstelling tot multidimensionale schaalmethoden, waarbij de rekencomplexiteit kwadratisch schaal, blijft deze hier lineair met het aantal deeltjes. Onze methode is in staat om meerdere variatiemodi te identificeren en toont veranderingen op nanometerschaal aan, zoals variaties in straal en hoogte, in zowel gesimuleerde als experimentele datasets.

In hoofdstuk 4 stellen we een methode voor om continue structurele heterogeniteit in Cryo-EM datasets te detecteren. Recente methoden zijn gebaseerd op kunstmatige intelligentie modellen die vaak grote trainingsdatasets en zorgvuldige afstemming van hyperparameters vereisen. Onze methode daarentegen detecteert onderliggende continue variaties in 2D-projecties door paarsgewijze vergelijking van beelden binnen oriëntatieklassen. De aanpak reconstrueert tussenliggende conformationele toestanden die de continue structurele heterogeniteit in synthetische SARS-CoV-2 spike-eiwitdata, gesimuleerd onder ideale omstandigheden, vertegenwoordigen. Realistischere simulaties, met variërende defocus per deeltje en stralingschade, leiden niet tot dezelfde gunstige resultaten, wat nog steeds een uitdaging vormt voor toekomstig onderzoek.



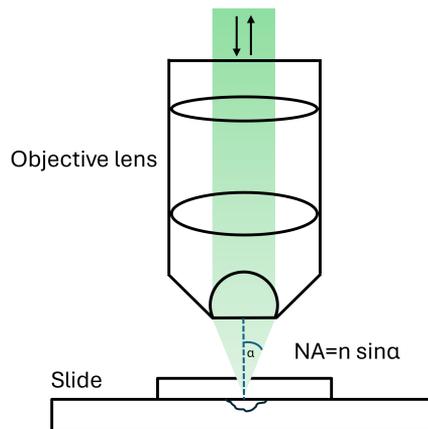
# 1

## INTRODUCTION

## 1.1. FLUORESCENCE MICROSCOPY

Fluorescence microscopy is an imaging technique for studying structures of interest in the sub-micron scale, making use of fluorescence to highlight these structures [1, 2]. Absorption of light by a fluorescent molecule causes an electron in the ground state to move to an excited state. Typically within nanoseconds, the excited electron returns to its ground state, causing emission of a photon. The emitted photon has a longer wavelength than the absorbed photon due to loss of energy to vibrations and rotations of the molecule, a phenomenon known as the Stokes shift [3]. Fluorescent molecules, fluorophores, can be selectively attached to a structure of interest by, e.g. an antibody. This labeling provides the specificity to fluorescence microscopy that makes the application of the technique so widespread in biology. Despite this major asset of fluorescence microscopy the technique also has a major drawback, namely the resolution of a conventional fluorescence microscope is limited to a few hundred nanometers. This diffraction limit to resolution was first derived in 1873 by Ernst Abbe [4] as the inverse of the spatial frequency cut-off of the optical imaging system:

$$d = \frac{\lambda}{2NA}, \quad (1.1)$$



**Figure 1.1: Diagram of excitation and emission process in epi-fluorescence microscopy.** The numerical aperture is determined by the refractive index ( $n$ ) of the immersion medium (here air  $n = 1$ ) of the objective and half of the maximum acceptance angle of light entering the objective lens.

where  $\lambda$  is the wavelength of the fluorescent light, which is typically in the visible region between 400 to 700 nm, and NA is the numerical aperture of the objective lens (about 1.4 for high immersion NA objectives) as shown in figure 1.1. This limit was a barrier to study nanoscale structures before the invention of super-resolution microscopy.

## 1.2. SUPER-RESOLUTION MICROSCOPY

Super-resolution microscopy (SRM) bypasses the diffraction limit and includes a wide range of strategies based on different physical principles [5]. The resolution of SRM microscopes can be a hundred-fold better than the resolution of conventional microscopes, and biological discoveries have been made, for instance the mapping of intricate neuronal networks [6]. It also plays a role in virology by enabling detailed imaging of virus-host [7] interactions as well as in cancer research for analyzing tumor micro-environments at the molecular level [8].

The field of super-resolution microscopy comprises of various distinct strategies. For example, Stimulated Emission Depletion (STED) microscopy [9, 10] is a widely used technique that employs both excitation and depletion beams to illuminate samples. In this method, the excitation beam activates the fluorophores, while a doughnut-shaped depletion beam deactivates them. This combination results in a significantly narrower fluorescence emission area, which improves resolution. As a result, the super-resolved images are formed from the fluorescence received from the center of the excitation focus. Photo-Activated Localization Microscopy (PALM) [11] or Fluorescence Photo-Activation Microscopy (FPALM) [12] and STochastic Optical Reconstruction Microscopy (STORM) [13] are super-resolution techniques that achieve high-resolution by capturing images of stochastically different subsets of active fluorescent molecules in successive time frames. In figure 1.2, the principles of these methods are depicted along with an example image for each technique. Expansion Microscopy (ExM) [14, 15] achieves higher resolution by embedding specimens in a polymer gel matrix and subsequently physically expanding them followed by imaging of the expanded sample with a conventional microscope, thereby improving the effective spatial resolution.

### 1.2.1. RESOLUTION ASSESSMENT IN SUPER-RESOLUTION MICROSCOPY

In the context of super-resolution microscopy, Fourier Ring Correlation (FRC) [17, 18] is commonly used to measure the image resolution. FRC quantifies the resolution of microscopy images by analyzing the similarity between two noise independent spatial frequency components of the image, and can be applied to various imaging modalities. It involves computing the correlation of concentric rings in the Fourier transform of the two image halves. FRC provides a resolution estimate based on the frequency at which the correlation drops below a threshold, indicating the spatial detail that is reliably discernible. Another method of assessing resolution is the spectral signal-to-noise ratio (SSNR) [19], which evaluates the ratio of true signal to noise across different spatial frequencies, providing a comprehensive measure of image quality and resolution.

## 1.3. SINGLE MOLECULE LOCALIZATION MICROSCOPY

Among the super-resolution techniques, single molecule localization microscopy (SMLM) is accessible to a broad community of researchers due to its ease in implementation. It reaches sub-diffraction limit resolution by making use of blinking fluorophores, i.e. by stochastically turning the emission light on and off. The family of SMLM techniques includes a plethora of acronyms such as PALM [11], STORM [13], dSTORM [20], and DNA-PAINT [21]. Figure 1.3 shows the principle of SMLM. As opposed to standard digital mi-

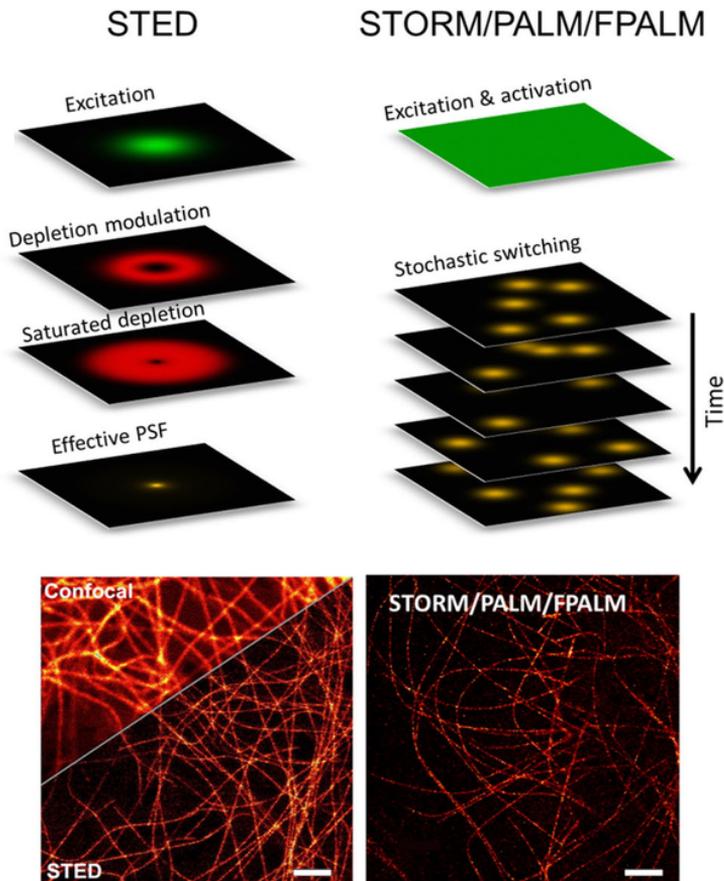


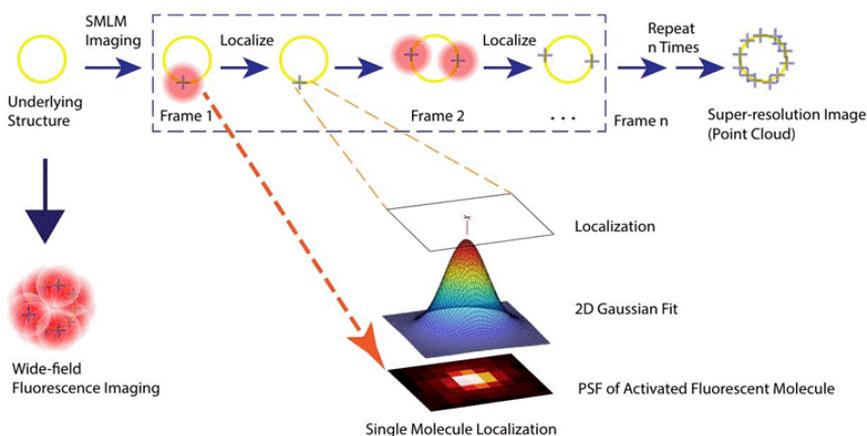
Figure 1.2: **Super-resolution microscopy.** The upper panel depicts the principles of SRM techniques. The lower panel shows confocal and super-resolution images of fluorescent protein labeled microtubules in living cells, confocal and STED microscopy of mCitrine-tubulin in a living PtK2 cell and STORM/PALM/FPALM of mEos2-tubulin in a living *Drosophila* S2 cell. The scale bar in both figures is  $2\ \mu\text{m}$ . Image from ref. [16].

scopy imaging that produces 2D pixel or 3D voxel grid data, SMLM produces 2D or 3D point cloud data, with associated uncertainties in each dimension. The acquisition of such images involves time series of frames. In each frame the emitted light from the stochastically activated fluorophores is collected using highly sensitive cameras, such as EM-CCD or sCMOS cameras, which are capable of detecting single-photon events. The emission spots are captured with a frame rate of typically 10 Hz to 100 Hz. In total the number of acquired frames can be anywhere from  $10^3$  to  $10^5$ . The total acquisition time of recording the frames can thus span from seconds to several hours. The obtained reso-

lutions are typically around 20 nm laterally and 50 nm axially [11].

### 1.3.1. SMLM DATA PROCESSING PIPELINE

SMLM raw data consist of a stack of frames recorded in the imaging experiment. In each frame, a fraction of the fluorophores is in the on-state while the other fluorophores are in the off-state. In conventional light microscopy the result of imaging is a directly interpretable (camera) image, but in SMLM this is not the case. To reconstruct a final image from SMLM data, a series of image processing steps is applied which can be divided in the following steps:



**Figure 1.3: Schematic of Single-Molecule Localization Microscopy.** Imaging the yellow circle labeled with fluorophores with wide field microscopy will result in a blurry image due to the diffraction limit. In SMLM, however, fluorophores are stochastically turned on and off resulting in sparse images. After fitting a Point Spread Function (PSF) model, often a Gaussian, to each observed spot the position of each fluorophore is estimated. These localizations are then plotted to create a super resolved SMLM image. Image from ref. [22].

### Segmentation

The process of distinguishing individual molecular signals from the background is called segmentation. This initial step in the analysis involves detecting regions that may contain signals from single molecules per time frame. This crucial first detection is particularly challenging in scenarios where the fluorescence signal is weak, or there is significant background. A key part of this detection process is defining the Region of Interest (ROI), which must be specified in terms of pixel size. The size of the ROI is critical as it determines the area over which the signal is analyzed. The back projected pixel size should be chosen based on the Nyquist sampling criterion, ensuring that the sampling rate is sufficient to reconstruct the signal without aliasing. The ROI can be as small as 7 pixels squared for standard 2D SMLM experiments. Filtering the fluorescent molecule signals

can be performed by using simple thresholding methods, or a set of advance probabilistic techniques [13, 23] and machine learning approaches [24].

### Localization

In the localization step, the aim is to accurately and precisely determine the position of the fluorescent molecule in each region of interest (ROI). In other words, the data are transformed from images to a list of coordinates (and other parameters such as background and photon count) plus uncertainty. Localization can be performed for both 2D and 3D imaging. 3D imaging techniques, such as astigmatic imaging introduces astigmatism in the optical path causing the PSF to elongate differently in the two in-plane directions depending on the axial position to capture images at different depths. In advanced models, further information such as the orientation of the fluorescent molecule and aberration can be estimated by using physical models of the PSF [25]. For such more complex imaging modalities the ROI can be as big as several tens of pixels squared in order to accommodate the complicated spot shapes that are generated on the camera. Maximum likelihood estimation (MLE) [26] and least squares (LS) fitting of a PSF model [27] to the ROIs represent the most efficient and commonly used methods for localization. Localization algorithms are inherently subject to errors due to noise, consisting of both random and systematic components as characterized by variance and bias, respectively, quantifying precision and accuracy [28]. The theoretical precision of an unbiased algorithm is limited by the Cramer-Rao Lower Bound (CRLB), which depends on the signal-to-noise ratio (SNR). The CRLB is given by:

$$\sigma_{\text{loc}} \geq \frac{\sigma_{\text{PSF}}}{\sqrt{N}}, \quad (1.2)$$

where  $\sigma_{\text{loc}}$  represents the standard deviation of the estimated coordinates,  $\sigma_{\text{PSF}}$  denotes the sigma of the Gaussian fit to the PSF, and  $N$  is the number of photons collected by the camera. Typical values of  $\sigma_{\text{PSF}}$  are around 100 nm and  $N$  ranges from  $10^2$  to  $10^4$ , resulting in predicted precision limits between 1 and 10 nm. However, this simplified formula does not account for additional factors, mostly the background, but also factors such as the non-Gaussian shape of the PSF, read and amplification noise, finite pixel size, and dipole orientation, all of which can degrade the achievable precision.

### Filtering

The quality of localizations in SMLM is often compromised by various factors such as autofluorescence signals, out-of-focus sources, and overlapping emissions from nearby emitters. Such issues lead to erroneous localizations. Filtering these out is important for enhancing the final SMLM image reconstruction. Filtering can be performed by evaluating multiple parameters; for instance, low photon count, background level, or poor quality of fit can indicate faulty localizations that need to be discarded. Additionally, evaluating the temporal correlation of localizations of the same fluorophore across consecutive frames within a given neighborhood helps identify and remove inaccurate localizations. This multi-parameter approach ensures a more robust and accurate filtering process, ultimately leading to higher quality SMLM images [29].

### Drift correction

An SMLM measurement needs a large number of image frames to localize all molecules and therefore takes considerable imaging time, up to even a few hours in some modalities like DNA-PAINT. Therefore, sample drift becomes a problem and cannot be neglected. Drift correction can be performed using fiducial markers, which are tracked continuously throughout the experiment [30] and can thus be used to correct for the moving sample based on their positions. Dai et al. [31] recently used synthetic nano-structures, e.g. grid patterns instead of beads achieving drift correction down to about 1 nm. Data analysis approaches such as (redundant) cross-correlation between blocks of consecutive frames can also be used to assess the sample drift [32]. The level of correction that was achieved was approximately 1 nm in the lateral ( $xy$ ) direction and 2 nm in the axial ( $z$ ) direction. These small values for the residual drift were achieved by using the redundant cross-correlation algorithm.

### Visualization

In the visualization step the localizations are represented in an image. There are several ways to do this. Using a scatter plot to visualize the localization coordinates is ineffective as the density of the points is not visible except for sparse cases. A better approach is to transfer the accumulated localized points to pixelated images by binning in boxes of size equal to a super-resolution pixel (chosen to be on the order of the localization uncertainty or less), possibly combined with Gaussian blurring with standard deviation the localization uncertainty [33, 34].

### Post processing and quantification

Analyzing and quantification of the SMLM data is important for finding new insights about the nanostructures and processes. Many techniques are developed to further interpret the data such as cluster analysis [35], resolution assessment [17], structural quantification [36, 37], particle fusion [38, 39], etc. Among these methods, particle fusion succeeded in pushing the resolution toward 1 nm.

## 1.4. CRYO-ELECTRON MICROSCOPY (CRYO-EM)

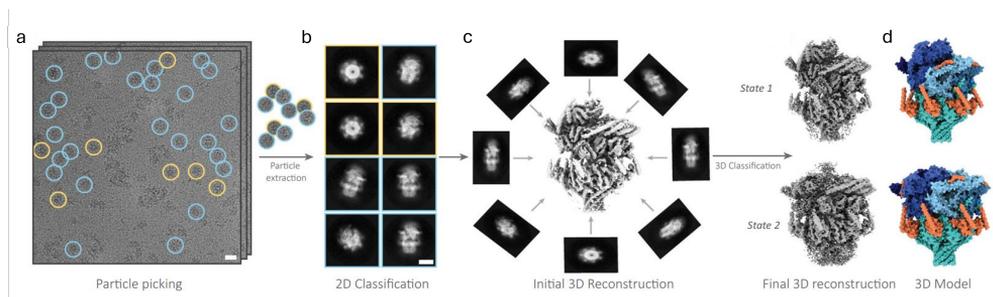
Electron Microscopes (EM) use electron beams controlled by electron optic devices that act similar to optical lenses to produce a magnified image [40]. The first Transmission Electron Microscope (TEM) [41] was introduced by Ernst Ruska in 1930. Nowadays TEM enables researchers to obtain a resolution on the Ångström level, and as such is one of the most important techniques to study (macro)molecular complexes in cells. The resolving power goes down to 50 pm in state-of-the-art machines in material science. A drawback of electron microscopes is the considerable radiation damage to the specimen caused by the highly energetic incident electrons [42]. Both amplitude contrast and phase contrast can be generated when electrons scatter against the sample. Phase contrast contains more information about biological specimens because amplitude contrast depends strongly on the atomic number, and for atoms that make up biological samples this is typically low and nearly equal (carbon, oxygen, nitrogen, etc.). Phase contrast can be transformed to intensity contrast by applying defocus (or a phase plate). The interference of the undiffracted beam with the beam that is phase shifted by the sample generates contrast in the case of

defocus. The Contrast Transfer Function (CTF) is a function that describes how the contrast of different spatial frequencies is transferred through the microscope to the image. Understanding and accounting for the CTF is needed for interpreting and processing EM images, especially in high-resolution applications.

The sample holder of an EM should be kept in vacuum as air molecules scatter the electrons. This makes many biological samples incompatible with this condition. Cryo-EM [43] is a technique in which the specimen is frozen to liquid nitrogen temperatures to cryo-immobilize (fixate) the sample. This can be achieved by either plunge freezing or high pressure freezing. A key element is that the freezing must be done very fast, so as to form amorphous and not crystalline ice, because crystalline ice will diffract the electron beam, preventing the generation of any meaningful image formation. At low temperatures, the molecular structures are excellently preserved, and thus provide a forum for ultrastructure studies. Although the TEM produces 2D images from the specimen it is possible to obtain 3D volume reconstructions using techniques such as single particle analysis (SPA) [44] or electron tomography (ET) [45].

#### 1.4.1. SINGLE PARTICLE ANALYSIS (SPA)

Single-particle analysis overcomes the limitation of the low dose required to prevent radiation damage by combining many images of copies of the same molecule in randomly different orientations. Non-biologic samples can be typically bombarded with  $500\text{--}2000\text{ e}^-/\text{\AA}^2$ , but biological samples can only tolerate a 100 times lower dose. In order to obtain a high quality reconstruction of a molecular structure the chemically identical copies must be aligned. This alignment process relies on advanced image processing methods that handle the variations in positions and orientations of the molecules, the effects of defocusing, and the extremely low signal-to-noise ratio (SNR) inherent to such data. The processing pipeline is depicted in figure 1.4.



**Figure 1.4: Steps of SPA.** (a) Particle picking or identification from the cryo-EM micrographs (b) Classification of picked particles based on their projection angle resulting in different orientation classes. (c) Particles from each orientation classes are used to calculate 3D reconstructions. (d) 3D classification and refinement of final high-resolution reconstruction to obtain density maps and atomic model. Image from ref. [46].

#### Particle picking

High-resolution cryo-EM studies, especially for asymmetric molecular structures, often

need selecting hundreds of thousands of high-quality particles from micrographs. This particle picking step is a major barrier to automating the SPA process. Historically, particle selection was done manually, a tedious and time-consuming task prone to human bias and inconsistency. To address this problem, various computational methods have been developed to assist in particle picking. These methods fall into three categories: generative [47], unsupervised [48], and discriminative approaches [49]. Generative methods identify particle candidates by measuring the similarity to a reference, often using template matching techniques. Unsupervised approaches distinguish particle-like objects from noisy background without labeled training data. Discriminative methods train a classifier on labeled datasets of positive and negative examples to detect particle images. Although these computational approaches have reduced the time and effort required for single-particle data analysis, they still fall short of a fully automated pipeline. Generative methods need an initial set of high-quality reference particles, and discriminative approaches require manually picked samples to train the classifier. Although unsupervised methods do not depend heavily on labeled data, they often do not fully exploit the intrinsic characteristics of particles, leading to their combination with template-matching [50] or classification-based methods [51] to achieve satisfactory results.

### Classification

An important part of the SPA analysis is 2D classification, which groups particles with similar 3D projections. Averaging over particles within such an orientation class improves the SNR. The classification step is challenging due to the low SNR of the individual particle projections. Typical 2D classification algorithms combine clustering with image alignment, iteratively identifying the best alignments and classification parameters. Since the early development of single-particle cryo-EM, a standard framework combining K-means clustering with multi-reference alignment (MRA) has been used. RELION [48], a widely used method, improves on this by employing a maximum-likelihood approach for MRA, comparing each image with all others in many orientations and translations. For the best classification results however, ISAC (iterative stable alignment and clustering) [52] can be used. This method validates class members through repeated stability tests, ensuring homogeneity and using modified K-means to limit class size. However, ISAC is extremely time-consuming and is typically used only for particularly challenging problems.

### 3D Reconstruction

Reconstruction of a 3D molecular structure can be achieved by combining many 2D projected images from different orientations. To this end, a variety of methods have been developed. The main challenge in 3D structure reconstruction is identifying the absolute orientation of individual 2D experimental projection images, typically indicated by projection angles. However, the low SNR, CTF of the microscope, and errors during particle picking make it challenging to determine these angles directly from the 2D projection images. The 3D reconstruction problem can be considered an optimization challenge, where the goal is to find a 3D model that aligns with the 2D projection images as closely as possible. Therefore, many current 3D reconstruction methods, such as EMAN [53] and RELION [48], employ expectation maximization to iteratively refine the estimation of projection angles and the 3D model, starting from an initial condition. Improvements in such

3D reconstruction methods are done by stochastic hill climbing [54], stochastic gradient descent [55], and particle filtering with a posterior probability density estimation [56] that achieved the resolution of about few Ångströms. In the presence of sample heterogeneity 3D reconstruction from 2D projection images is even more difficult. This problem is typically addressed by considering multiple 3D structures in the reconstruction pipeline.

## 1.5. PARTICLE FUSION IN SMLM

Methods for particle fusion in light microscopy are relatively new due to recent advances in super-resolution techniques, despite the longstanding development of SPA in cryo-EM. In the same way as for cryo-EM, where the reconstruction of the molecular complex is performed by fusion of projections obtained from chemically identical copies of the molecular complex, it is possible to leverage fusion in SMLM in cases where these similar molecular structures exist. In SMLM there are two main challenges that limit the resolution. First, the low binding affinity of fluorescent labels results in a degree of labeling that is typically on the order of 30-70%, which means that only about half of the potential binding sites is actually labeled. Second, the finite localization uncertainty, limited by the photon count of the emitters, gives rise to a spread of localizations around the binding sites. Moreover, false positive localizations can further deteriorate the quality of the data. Fusion of chemically identical structures increases the SNR, and diminishes the effect of missing labels at binding sites. This section elaborates on the main approaches for particle fusion in super-resolution light microscopy, highlighting how these methods enhance imaging quality.

### 1.5.1. EM INSPIRED PARTICLE FUSION

A logical starting point for SMLM particle fusion has been in its historical origin in SPA for cryo-EM. This requires adaptation of the SMLM data to fit the framework of cryo-EM based methods. In particular, this entails a conversion of the localization point cloud to pixelized image data, with noise and randomness of an entirely different character compared to cryo-EM data. One of the works that used EM inspired techniques to fuse SMLM data was done by Szymborska et al. [57] on Nuclear Pore Complexes (NPCs) which revealed highly symmetric arrangement of the NPC scaffold with unprecedented detail. Cryo-EM methodologies can be adapted for use in SMLM to enhance image resolution and analysis through similar principles of particle alignment, fusion, and 3D reconstruction [58]. However, these 2D/3D reconstruction attempts suffer from the mismatch between the central premise including disparities between coordinate-based (SMLM) and pixel-based image (cryo-EM) data representations.

### 1.5.2. TEMPLATE-BASED PARTICLE FUSION MODEL BASED FITTING

The structure of several molecular complexes has been revealed at a resolution of a few Ångströms with cryo-EM studies. Some of these structures are then used as a template to perform particle fusion in SMLM. This means that an accurate model is prepared in advance for the benefit of guiding the SMLM particle fusion. Löscherger et al. [37] used the eight-fold symmetric ring of the NPC as a template. In this study optimal translation and rotation were found by minimization of the sum of the distances of localizations in each segmented particle to the eight binding sites of the model. Similar approaches were

tried in the study of tegument proteins [59] and endocytic proteins [60]. Provided that the accuracy of the models can be guaranteed, it is possible to use these template driven techniques. The major benefit is the relatively low computational complexity that scales linearly with the number of particles. On the other hand, if the accuracy of the template is not well established, the fusion becomes susceptible to template bias, resulting in a wrong reconstruction.

### 1.5.3. TEMPLATE FREE PARTICLE FUSION

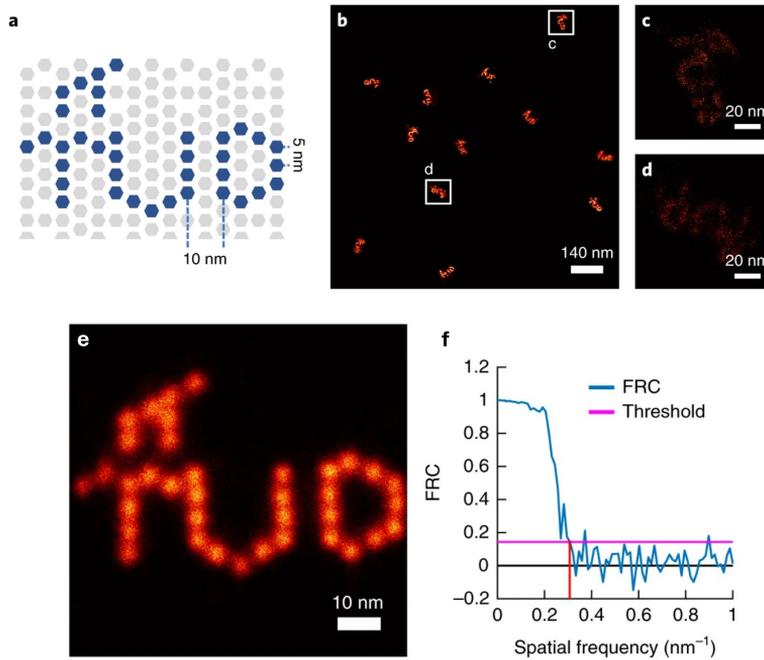
The issue of template bias is overcome by template free methods, potentially enabling discovery of unknown sub-cellular structures. Several approaches have been proposed to align and merge particle localizations, in particular pyramid registration [61], all-to-all registration based fusion [38], and fast particle fusion approach [62]. An example of prior knowledge free structure assessment is the symmetry discovery method enabled by the all-to-all method [39]. An example result from the application of this method that achieved around 3 nm resolution is illustrated in Figure 1.5. The algorithmic complexity of template free particle fusion methods is usually high. For instance, the computational complexity of the all-to-all [38, 39] approach scales quadratically with the number of particles. Recently, however, Wang et al. [62] introduced a method adapted from the so-called Joint Registration of Multiple Point Clouds (JRMPC) to speed-up the process, achieving a linear scaling with the number of particles. This computational advantage was required for a high precision study of the two-binding site sub-structure of Nup96 using thousands of NPCs.

## 1.6. MOTIVATION AND OUTLINE

The key assumption in all cryo-EM SPA and SMLM particle fusion studies is the homogeneity of the data. That is, it is assumed that all particles have a chemically identical structure (apart from the labelling). In practice, however, this assumption may not be valid. Neglecting structural heterogeneity in the particle fusion will lead to a lower SNR of the reconstruction, as overlaying dissimilar structures will smear out features. Uncovering structural heterogeneity offers potentially useful information that exists in the data.

Structural heterogeneity can be divided into two categories: discrete heterogeneity and continuous heterogeneity. Detecting discrete states of proteins/molecules can be performed by various classification approaches. In cryo-electron microscopy, Bayesian approaches [63] and deep learning methods [64, 65] are two ways of dealing with different conformation states of molecular structures. For SMLM data, Huijben et al. [36] studied discrete structural heterogeneity using a classification approach without using a model. Continuous structural heterogeneity detection, however, is a more difficult problem and cannot be addressed by common classification algorithms. Therefore, new strategies are needed to be able to detect this type of heterogeneity.

In cryo-EM it is common to use machine learning methods [66, 67] to detect continuous structural heterogeneity. Just as in cryo-EM, a learning-based method [24] for SMLM was introduced, but it requires extensive labeling for training. Another method, ECLiPSE [68], relies on statistical pattern recognition and requires segmentation, demanding high labeling accuracy and a high SNR. Methods like LocMoFit [69] fit a point cloud to a geo-



**Figure 1.5: An example of template free particle fusion in SMLM** *a)* The template used to build the DNA origami TU Delft logo with 37 docking sites. *b)* Field of view containing many logo particles. *c,d)* Two examples of single particles with different orientations and with incomplete labeling (DOL=80%). *e)* The fusion result of 383 individual particles. *f)* FRC curve showing the resolution of  $3.3 \pm 0.3$  nm. Image from ref. [38].

metric model based on prior knowledge. Despite these efforts on uncovering structural heterogeneity, an approach that can reliably shed light on continuous heterogeneity in the data remains a challenge. The continuous nature of the heterogeneity relates to one or more structural parameters that take values in a continuous range, such as the radius of the NPC. The aim of this thesis is to i) develop methods to detect continuous structural heterogeneity and ii) from that enhance the reconstruction quality by particle fusion of similar particles opposed to fusing all data.

### 1.6.1. OUTLINE

The body of the thesis consists of three content chapters. The first two chapters introduce frameworks to detect continuous structural heterogeneity in SMLM data. The first builds on a similarity measure introduced earlier in the all-to-all registration and the second leverages data-driven variational auto-encoders. In the last chapter, we introduce a method for detecting continuous structural heterogeneity in cryo-EM data. A summary of

each chapter is provided below:

In chapter 2, we present a prior-knowledge-free method for detecting continuous structural variations with localization microscopy. In experimental datasets, we show the continuous variation of the height of DNA origami tetrahedrons imaged with 3D PAINT and of the radius of Nuclear Pore Complexes imaged in 2D with STORM.

In chapter 3, we introduce a new pipeline based on a data-driven variational auto-encoder approach to detect the continuous structural heterogeneity. The aims of this approach are to address the challenges raised in chapter two, namely, to overcome the quadratic scaling of computational time with respect to the number of particles and to disentangle multiple modes of continuous variation in SMLM data simultaneously.

In chapter 4, we focus on detection of continuous structural heterogeneity in cryo-EM data for investigating variations of protein conformation. The spike protein of SARS-CoV-2 [70–72] for instance, exhibits multiple conformations that are critical for its function in host cell entry and antibody recognition.

The advanced computational techniques developed in this thesis may lead to improved access to high-value structural information from microscopy data in different high-resolution modalities.

## REFERENCES

- [1] J. R. Lakowicz, *Principles of fluorescence spectroscopy* (Springer, 1999) pp. 1–23.
- [2] B. Valeur and M. N. Berberan-Santos, *Molecular fluorescence: principles and applications* (John Wiley & Sons, 2013).
- [3] P. Abbyad, W. Childs, X. Shi, and S. G. Boxer, *Dynamic Stokes shift in green fluorescent protein variants*, [Proceedings of the National Academy of Sciences](#) **104** (2007), pp. 20189–20194.
- [4] E. Abbe, *Beiträge zur Theorie des Mikroskops und der mikroskopischen Wahrnehmung*, [Archiv für mikroskopische Anatomie](#) **9** (1873), pp. 413–468.
- [5] L. Schermelleh, R. Heintzmann, and H. Leonhardt, *A guide to super-resolution fluorescence microscopy*, [Journal of Cell Biology](#) **190** (2010), pp. 165–175.
- [6] A. Dani, B. Huang, J. Bergan, C. Dulac, and X. Zhuang, *Superresolution imaging of chemical synapses in the brain*, [Neuron](#) **68** (2010), pp. 843–856.
- [7] S. Banerjee, S. Maurya, and R. Roy, *Single-molecule fluorescence imaging: Generating insights into molecular interactions in virology*, [Journal of biosciences](#) **43** (2018), pp. 519–540.
- [8] M. Heilemann, D.-P. Herten, R. Heintzmann, C. Cremer, R. Eils, and D. Baddeley, *Super-resolution microscopy provides insights into the tumor microenvironment*, [Nature Reviews Cancer](#) **12** (2012), pp. 9–22.
- [9] S. W. Hell and J. Wichmann, *Breaking the diffraction resolution limit by stimulated emission: stimulated-emission-depletion fluorescence microscopy*, [Optics Letters](#) **19** (1994), pp. 780–782.

- [10] T. A. Klar and S. W. Hell, *Subdiffraction resolution in far-field fluorescence microscopy*, *Optics Letters* **24** (1999), pp. 954–956.
- [11] E. Betzig, G. H. Patterson, R. Sougrat, O. S. Lindwasser, S. Olenych, J. S. Bonifacino, M. W. Davidson, J. Lippincott-Schwartz, and H. F. Hess, *Imaging intracellular fluorescent proteins at nanometer resolution*, *Science* **313** (2006), pp. 1642–1645.
- [12] S. T. Hess, T. P. K. Girirajan, and M. D. Mason, *Ultra-high resolution imaging by fluorescence photoactivation localization microscopy*, *Biophysical Journal* **91** (2006), pp. 4258–4272.
- [13] M. J. Rust, M. Bates, and X. Zhuang, *Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM)*, *Nature Methods* **3** (2006), pp. 793–796.
- [14] D. Gambarotto, F. U. Zwettler, M. Le Guennec, M. Schmidt-Cernohorska, D. Fortun, S. Borgers, J. Heine, J.-G. Schloetel, M. Reuss, M. Unser, E. S. Boyden, M. Sauer, and V. Hamel, *Imaging cellular ultrastructures using expansion microscopy (U-ExM)*, *Nature Methods* **16** (2019), pp. 71–74.
- [15] P. W. Tillberg, F. Chen, K. D. Piatkevich, Y. Zhao, C. C. Yu, B. Wu, J. A. English, L. Gao, A. Martorell, H.-J. Suk, F. Yoshida, E. M. DeGennaro, D. E. Roossien, G. Gong, S. J. Sahl, D. Cai, and E. S. Boyden, *Protein-retention expansion microscopy of cells and tissues labeled using standard fluorescent proteins and antibodies*, *Nature Biotechnology* **34** (2016), pp. 987–992.
- [16] F. D. Stefani, M. H. Ulbrich, J. Ortega-Arroyo, L. Peres-Lemus, H. Cang, B. Cortés, F. Agullo-Rueda, A. Dani, S. Shoham, P. R. Selvin, P. S. Blank, and D. Fotiadis, *Fluorescent probes for nanoscopy: Four categories and multiple possibilities*, *Nature Reviews Chemistry* **4** (2012), pp. 507–520.
- [17] R. P. J. Nieuwenhuizen, K. A. Lidke, M. Bates, D. Leyton Puig, D. Grünwald, S. Stallinga, and B. Rieger, *Measuring image resolution in optical nanoscopy*, *Nature Methods* **10** (2013), pp. 557–562.
- [18] B. Rieger, I. Droste, F. Gerritsma, T. Ten Brink, and S. Stallinga, *Single image fourier ring correlation*, *Optics Express* **32** (2024), pp. 21767–21782.
- [19] M. Unser, B. L. Trus, and A. C. Steven, *A new resolution criterion based on spectral signal-to-noise ratios*, *Ultramicroscopy* **23** (1987), pp. 39–51.
- [20] M. Heilemann, S. van de Linde, M. Schüttpehl, R. Kasper, B. Seefeldt, A. Mukherjee, P. Tinnefeld, and M. Sauer, *Subdiffraction-resolution fluorescence imaging with conventional fluorescent probes*, *Angewandte Chemie International Edition* **47** (2008), pp. 6172–6176.
- [21] R. Jungmann, M. S. Avendaño, J. B. Woehrstein, M. Dai, W. M. Shih, and P. Yin, *Single-molecule kinetics and super-resolution microscopy by fluorescence imaging of transient binding on dna origami*, *Nano Letters* **10** (2010), pp. 4756–4761.

- [22] I. M. Khater, I. R. Nabi, and G. Hamarneh, *A review of super-resolution single-molecule localization microscopy cluster analysis and quantification methods*, *Patterns* **1** (2020), p. 100018.
- [23] C. S. Smith, S. Stallinga, K. A. Lidke, B. Rieger, and D. Grünwald, *Probability-based particle detection that enables threshold-free and robust in vivo single-molecule tracking*, *Molecular Biology of the Cell* **26** (2015), pp. 4057–4062.
- [24] J. S. Daniel and A. J. Garcia-Saez, *Quantitative analysis of super-resolved structures using asap*, *Nature Methods* **16** (2019), pp. 711–714.
- [25] C. N. Hulleman, R. Thorsen, E. Kim, C. Dekker, S. Stallinga, and B. Rieger, *Simultaneous orientation and 3d localization microscopy with a vortex point spread function*, *Nature Communications* **12** (2021), p. 5934.
- [26] C. S. Smith, N. Joseph, B. Rieger, and K. A. Lidke, *Fast, single-molecule localization that achieves theoretically minimum uncertainty*, *Nature Methods* **7** (2010), pp. 373–375.
- [27] A. V. Abraham, S. Ram, J. Chao, E. S. Ward, and R. J. Ober, *Quantitative study of single molecule location estimation techniques*, *Optics Express* **17** (2009), pp. 23352–23373.
- [28] M. Lelek, M. T. Gyparaki, G. Beliu, F. Schueder, J. Griffie, S. Manley, R. Jungmann, M. Sauer, M. Lakadamyali, and C. Zimmer, *Single-molecule localization microscopy*, *Nature Reviews Methods Primers* **1** (2021), p. 39.
- [29] B. Rieger, R. P. J. Nieuwenhuizen, and S. Stallinga, *Image processing and analysis for single-molecule localization microscopy: Computation for nanoscale imaging*, *IEEE Signal Processing Magazine* **32** (2014), pp. 49–57.
- [30] A. Balinovic, D. Albrecht, and U. Endesfelder, *Spectrally red-shifted fluorescent fiducial markers for optimal drift correction in localization microscopy*, *Journal of Physics D: Applied Physics* **52** (2019), p. 204002.
- [31] M. Dai, R. Jungmann, and P. Yin, *Optical imaging of individual biomolecules in densely packed clusters*, *Nature Nanotechnology* **11** (2016), pp. 798–807.
- [32] Y. Wang, J. Schnitzbauer, Z. Hu, X. Li, Y. Cheng, Z.-L. Huang, and B. Huang, *Localization events-based sample drift correction for localization microscopy with redundant cross-correlation algorithm*, *Optics Express* **22** (2014), pp. 15982–15991.
- [33] D. Baddeley, M. B. Cannell, and C. Soeller, *Visualization of localization microscopy data*, *Microscopy and Microanalysis* **16** (2010), pp. 64–72.
- [34] R. P. J. Nieuwenhuizen, S. Stallinga, and B. Rieger, *Visualization and resolution in localization microscopy*, *Cell Membrane Nanodomains* (2014), p. 409.
- [35] P. Rubin-Delanchy, G. L. Burn, J. Griffiñ, D. J. Williamson, N. A. Heard, A. P. Cope, and D. M. Owen, *Bayesian cluster identification in single-molecule localization microscopy data*, *Nature Methods* **12** (2015), pp. 1072–1076.

- [36] T. A. Huijben, H. Heydarian, A. Auer, F. Schueder, R. Jungmann, S. Stallinga, and B. Rieger, *Detecting structural heterogeneity in single-molecule localization microscopy data*, [Nature Communications](#) **12** (2021), p. 3791.
- [37] A. Löschberger, S. van de Linde, M.-C. Dabauvalle, B. Rieger, M. Heilemann, G. Krohne, and M. Sauer, *Super-resolution imaging visualizes the eightfold symmetry of gp210 proteins around the nuclear pore complex and resolves the central channel with nanometer resolution*, [Journal of Cell Science](#) **125** (2012), pp. 570–575.
- [38] H. Heydarian, F. Schueder, M. T. Strauss, B. Van Werkhoven, M. Fazel, K. A. Lidke, R. Jungmann, S. Stallinga, and B. Rieger, *Template-free 2D particle fusion in localization microscopy*, [Nature Methods](#) **15** (2018), pp. 781–784.
- [39] H. Heydarian, M. Joosten, A. Przybylski, F. Schueder, R. Jungmann, B. Van Werkhoven, J. Keller-Findeisen, *et al.*, *3D particle averaging and detection of macromolecular symmetry in localization microscopy*, [Nature Communications](#) **12** (2021), p. 2847.
- [40] L. Reimer and H. Kohl, *Transmission Electron Microscopy: Physics of Image Formation and Microanalysis* (Springer, 2008).
- [41] E. Ruska, *The early development of electron lenses and electron microscopy*, [Microscopy, microanalysis, microstructures](#) **8** (1987), pp. 97–108.
- [42] D. B. Williams and C. B. Carter, *Transmission Electron Microscopy: A Textbook for Materials Science* (Springer, 2009).
- [43] J. Dubochet, M. Adrian, J. Chang, J. Homo, J. Lepault, A. McDowell, and P. Schultz, *Cryo-electron microscopy of vitrified specimens*, [Quarterly reviews of biophysics](#) **21** (1988), pp. 129–228.
- [44] J. Frank, *Three-Dimensional electron microscopy of macromolecular assemblies: visualization of biological molecules in their native state* (Oxford University Press, 2006).
- [45] V. Lucic, F. Foerster, and W. Baumeister, *Three-dimensional imaging of native cell ultrastructure using cryo-electron tomography*, [Journal of Cell Science](#) **118** (2005), pp. 755–761.
- [46] L. Anton, A. Leonie, D. W. Cobb, and C.-M. Ho, *Structural parasitology of the malaria parasite \*plasmodium falciparum\**, [Trends in Biochemical Sciences](#) **47** (2022), pp. 149–159.
- [47] F. Guo, W. Jiang, M. Sun, and M. G. Rossmann, *Automated particle selection using a neural network classification strategy with a fuzzy segmentation and a genetic algorithm*, [Journal of Structural Biology](#) **147** (2004), pp. 237–252.
- [48] S. H. Scheres, *Relion: Implementation of a bayesian approach to cryo-EM structure determination*, [Journal of Structural Biology](#) **180** (2012), pp. 519–530.
- [49] F. Wang, H. Gong, P. Zhang, Z. Zhang, L. Chen, H. Zheng, and Y. Xing, *Deeppicker: A deep learning approach for fully automated particle picking in cryo-EM*, [Journal of Structural Biology](#) **195** (2016), pp. 325–336.

- [50] J. Rogelj, D. Sousa, J. He, D. A. Agard, and E. Nogales, *Automated particle picking for low-contrast macromolecular electron microscopy*, *Journal of Structural Biology* **153** (2006), pp. 231–243.
- [51] Y. Zhu, Q. Ouyang, D. C. Trudgian, S. H. Scheres, and D. E. Shaw, *Automated particle selection: results of a comparative study*, *Journal of Structural Biology* **145** (2004), pp. 3–14.
- [52] Z. Yang, J. Fang, J. Chittuluru, F. J. Asturias, and P. A. Penczek, *Iterative stable alignment and clustering of 2d transmission electron microscope images*, *Structure* **20** (2012), pp. 237–247.
- [53] G. Tang, L. Peng, P. R. Baldwin, D. S. Mann, W. Jiang, I. Rees, and S. J. Ludtke, *EMAN2: an extensible image processing suite for electron microscopy*, *Journal of Structural Biology* **157** (2007), pp. 38–46.
- [54] P. A. Penczek and F. J. Asturias, *Ab initio cryo-em structure determination as a validation problem*, in *Proceedings of the 2014 IEEE International Conference on Image Processing (ICIP)* (IEEE, 2014) pp. 2090–2094.
- [55] A. Punjani, J. L. Rubinstein, D. J. Fleet, and M. A. Brubaker, *cryosparc: algorithms for rapid unsupervised cryo-EM structure determination*, *Nature Methods* **14** (2017), pp. 290–296.
- [56] M. Hu, J. Yu, C. Li, X. Lei, Z. Hong Zhou, and J. Frank, *A particle-filter framework for robust cryoem 3d reconstruction*, *Nature Methods* **15** (2018), pp. 1083–1089.
- [57] A. Szymborska, A. De Marco, N. Daigle, V. C. Cordes, J. A. Briggs, and J. Ellenberg, *Nuclear pore scaffold structure analyzed by super-resolution microscopy and particle averaging*, *Science* **341** (2013), pp. 655–658.
- [58] C. Sieben, N. Banterle, K. M. Douglass, P. GÅnczy, and S. Manley, *Multicolor single-particle reconstruction of protein complexes*, *Nature Methods* **15** (2018), pp. 777–780.
- [59] R. F. Laine, A. Albecka, S. Van De Linde, E. J. Rees, C. M. Crump, and C. F. Kaminski, *Structural analysis of herpes simplex virus by optical super-resolution imaging*, *Nature Communications* **6** (2015), p. 5980.
- [60] M. Mund, J. A. van Der Beek, J. Deschamps, S. Dmitrieff, P. Hoess, J. L. Monster, A. Picco, F. Nedelec, M. Kaksonen, and J. Ries, *Systematic nanoscale analysis of endocytosis links efficient vesicle formation to patterned actin nucleation*, *Cell* **174** (2018), pp. 884–896.
- [61] J. Broeken, H. Johnson, D. S. Lidke, S. Liu, R. P. J. Nieuwenhuizen, S. Stallinga, K. A. Lidke, and B. Rieger, *Resolution improvement by 3D particle averaging in localization microscopy*, *Methods and Applications in Fluorescence* **3** (2015), p. 015003.
- [62] W. Wang, H. Heydarian, T. A. Huijben, S. Stallinga, and B. Rieger, *Joint registration of multiple point clouds for fast particle fusion in localization microscopy*, *Bioinformatics* **38** (2022), pp. 3281–3287.

- [63] S. H. Scheres, *A bayesian view on cryo-EM structure determination*, [Journal of Molecular Biology](#) **415** (2012), pp. 406–418.
- [64] B. M. Powell and J. H. Davis, *Learning structural heterogeneity from cryo-electron subtomograms with tomoDRGN*, [Nature Methods](#) (2024), pp. 1–12.
- [65] J.-G. Wu, Y. Yan, D.-X. Zhang, B.-W. Liu, Q.-B. Zheng, X.-L. Xie, S.-Q. Liu, S.-X. Ge, Z.-G. Hou, and N.-S. Xia, *Machine learning for structure determination in single-particle cryo-electron microscopy: A systematic review*, [IEEE Transactions on Neural Networks and Learning Systems](#) **33** (2021), pp. 452–472.
- [66] E. D. Zhong, T. Bepler, B. Berger, and J. H. Davis, *Cryodrgn: reconstruction of heterogeneous cryo-em structures using neural networks*, [Nature Methods](#) **18** (2021), pp. 176–185.
- [67] M. Chen and S. J. Ludtke, *Deep learning-based mixed-dimensional gaussian mixture model for characterizing variability in cryo-em*, [Nature Methods](#) **18** (2021), pp. 930–936.
- [68] S. Hugelier, H. Kim, M. T. Gyparaki, C. Bond, Q. Tang, A. N. Santiago-Ruiz, S. Porta, and M. Lakadamyali, *ECLIPSE: A versatile classification technique for structural and morphological analysis of super-resolution microscopy data*, [BioRxiv](#) (2023), pp. 2023–05.
- [69] Y.-L. Wu, P. Hoess, A. Tschanz, U. Matti, M. Mund, and J. Ries, *Maximum-likelihood model fitting for quantitative analysis of SMLM data*, [Nature Methods](#) **20** (2023), pp. 139–148.
- [70] A. C. Walls, Y.-J. Park, M. A. Tortorici, A. Wall, A. T. McGuire, and D. Velesler, *Structure, function, and antigenicity of the SARS-CoV-2 spike glycoprotein*, [Cell](#) **181** (2020), pp. 281–292.
- [71] Y. Yuan, D. Cao, Y. Zhang, J. Ma, J. Qi, Q. Wang, and G. F. Gao, *Cryo-EM structures of MERS-CoV and SARS-CoV spike glycoproteins reveal the dynamic receptor binding domains*, [Nature Communications](#) **8** (2020), p. 15092.
- [72] D. Wrapp, N. Wang, K. S. Corbett, J. A. Goldsmith, C.-L. Hsieh, O. Abiona, B. S. Graham, and J. S. McLellan, *Cryo-EM structure of the 2019-nCoV spike in the prefusion conformation*, [Science](#) **367** (2020), pp. 1260–1263.

# 2

## DETECTING CONTINUOUS STRUCTURAL HETEROGENEITY IN SINGLE-MOLECULE LOCALIZATION MICROSCOPY DATA

Fusion of multiple chemically identical complexes, so-called particles, in localization microscopy, can improve the signal-to-noise ratio and overcome under-labeling. To this end, structural homogeneity of the data must be assumed. Biological heterogeneity, however, could be present in the data originating from distinct conformational variations or (continuous) variations in particle shapes. We present a prior-knowledge-free method for detecting continuous structural variations with localization microscopy. Detecting this heterogeneity leads to more faithful fusions and reconstructions of the localization microscopy data as their heterogeneity is taken into account. In experimental datasets, we show the continuous variation of the height of DNA origami tetrahedrons imaged with 3D PAINT and of the radius of Nuclear Pore Complexes (NPCs) imaged in 2D with STORM. In simulation, we study the impact on the heterogeneity detection pipeline of Degree Of Labeling and of structural variations in the form of two independent modes.<sup>1</sup>

---

<sup>1</sup>This chapter has been published as: **Sobhan haghparast**, Bernd Rieger, Sjoerd Stallinga, *Detecting continuous structural heterogeneity in single-molecule localization microscopy data*, Scientific Reports 13.1 (2023): 19800.

## 2.1. INTRODUCTION

Single-molecule localization microscopy (SMLM) is used to image biological samples at resolutions below the diffraction limit [1, 2]. The resolution of SMLM is limited by the density of labeling and localization precision [3, 4]. Fusion of multiple SMLM image datasets of chemically identical structures (particles) that are typically multi-component protein complexes with fixed spatial relationships, can improve the resolution as it increases the signal-to-noise ratio and overcomes under-labeling [5, 6]. In SMLM the general idea is along the same line as in single particle averaging techniques (SPA) in cryo-EM [7, 8]. Alignment of these particles is commonly achieved using either model-based registration methods [9, 10] or template-free registration methods [11–14]. Both approaches assume homogeneity of the dataset: the underlying structure is assumed to be the same for all particles. Structural heterogeneity, however, could be present in the data. These structural variations could originate from e.g. biological variations [6] or from sample preparation [15]. In the field of cryo-electron microscopy (cryo-EM) [16–18] methods to detect and cluster discrete variations are commonly used [19–22]. Recently, we have proposed a method to detect discrete structural variations in SMLM data with a clustering approach [23].

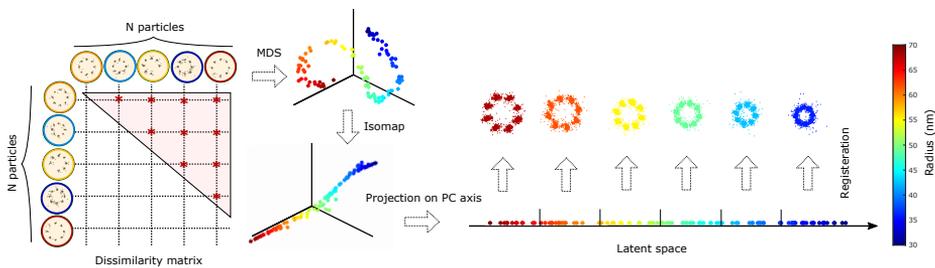
The goal of the current paper is to develop a method that can detect continuous heterogeneity in specifically SMLM datasets. This offers the potential to study naturally occurring biological variations related to e.g. dynamics or development of protein assemblies in cells. Such a method also has the advantage that blurring of the final fusion result by the underlying continuous heterogeneity can be tackled by fusing parts of the dataset that are detected to be sufficiently homogeneous. The stated task of detecting continuous heterogeneity is challenging in view of statistical variations inherent to SMLM point datasets. Relevant factors for these statistical variations are the (anisotropic) localization uncertainty, which results in scatter of localizations around the fluorophore binding sites, repeated localizations corresponding to the same fluorophore, which leads to variations of the density of localizations, and underlabeling, which gives rise to randomly different subsets of the total set of binding sites that are actually labeled with a fluorophore. These effects make the image formation for SMLM essentially different from cryo-EM, and stand in the way of direct application of methods developed for cryo-EM, such as a recent proposal to investigate continuous heterogeneity using a deep learning-based mixed-dimensional Gaussian mixture model [24]. A learning-based approach for SMLM has been introduced recently [25], but the classification tool needs manual annotation for training. Another statistical pattern recognition approach for classification (ECLiPSE [26]), requires segmentation, and therefore needs a high degree of labeling and signal-to-noise ratio. Approaches like LocMoFit (Localization Model Fit) [27] can be seen as a way of fitting a point cloud with a parametrized geometric model built on *a priori* knowledge. This is in contrast to our approach where we use a data-driven analysis by extracting information without the use of prior knowledge. Model-based approaches [28] in general are not ideal as they are susceptible to template bias and subjective model selection.

In this paper we present a model-free continuous heterogeneity detection method that works directly on localization coordinates in order to employ the full potential of the SMLM data. We apply our method to experimental data for the continuous detection of the height of DNA-origami tetrahedron structures imaged in 3D and of the radius of Nu-

clear Pore Complexes (NPCs) imaged in 2D.

## 2.2. METHODS

Our continuous heterogeneity detection (CHD) pipeline is illustrated in Fig. 2.1. It starts by computing a pairwise registration of all particles to obtain a dissimilarity matrix of the particles, similar to the approach of Huijben et al.[23]. The dissimilarity values between particles are mapped via multi-dimensional scaling[29] to a high-dimensional feature space (MDS). This typically results in a low-dimensional manifold embedded in the high-dimensional MDS space. The low-dimensionality indicates that the particles are mostly alike but vary in one or just a few features. In the next step, the Isomap algorithm[30] is used to “unroll” the data into a lower dimensional embedding. Finally, principal component analysis is used to project this representation onto the axis carrying the largest variation in a 1D latent space. This sequence of transformations preserves local ordering, and therefore, the ordering in 1D latent space carries the information on the continuous heterogeneity captured by the dissimilarity measure. Since neighboring particles in the latent space are structurally similar to each other, we can now divide them into bins and fuse the particles per bin.



**Figure 2.1: Continuous heterogeneity detection pipeline.**  $N$  particles are registered using the all-to-all registration procedure resulting in  $N(N - 1)/2$  dissimilarity values. Multi-dimensional scaling embeds the dissimilarity matrix into an abstract Cartesian space (only the first 3 dimensions are shown). Points lying on a low-dimensional manifold are unrolled using Isomap. Projection of the unrolled manifold on the main principal component axis creates a 1D latent space in which the data is ordered based on the dissimilarity captured by the dissimilarity measure.

In the following we describe the different steps of the algorithm in detail.

### 2.2.1. PAIRWISE REGISTRATION.

We use the all-to-all registration [11] to register  $N$  particles from a 2D/3D SMLM dataset. Each particle is independently registered to all other particles in the dataset using a combination of Gaussian Mixture Model (GMM) registration [31] and the Bhattacharya cost function [14]. First, particles are aligned using GMM registration with multiple initial poses, which results in a set of rotation matrices and translation vectors. The final registration parameters are those that maximize the Bhattacharya cost function over the GMM

optima for the different initializations. This procedure results in transformation parameters (rotation and translation) and the optimum cost value for each pair. The width of the Gaussian distribution (scale parameter) used in GMM registration is obtained by registering ten random groups of particles for different scales. In this scale-sweep approach, the scale parameter which gives the highest Bhattacharya cost value for all test sets is chosen as the scale value to be used for the whole dataset. It is necessary to select the proper scale value to avoid blurring of nearby binding sites or overfitting on each localization. The elements of the  $N \times N$  matrix of optimum cost values (or rather the upper triangular part of this matrix) are normalized by the number of localizations for the two corresponding particles [23]. This makes the cost function matrix less sensitive to variations in the number of localizations. Finally, all pairwise cost values (quantifying the similarity between pairs of particles) are subtracted from the maximum pairwise cost value to create a dissimilarity matrix.

### 2.2.2. MULTI-DIMENSIONAL SCALING (MDS) SPACE.

We assign coordinates to the  $N$  particles in a high-dimensional (dimension  $D = 30$ ) space, by interpreting the values of the dissimilarity matrix as the Euclidean distance in this high-dimensional space, the so-called MDS space. This is done by iteratively updating the coordinates in MDS space to minimize the stress function:

$$S = \sum_{i,j} (d_{ij} - \|x_i - x_j\|)^2, \quad (2.1)$$

where  $d_{ij}$  is the pairwise dissimilarity between particle  $i$  and  $j$ , and  $x_i$  is the MDS position vector of particle  $i$  ( $i = 1, 2, \dots, N$ ). We found that a dimensionality  $D > 15$  resulted in a value for the stress function smaller than  $10^{-4}$ , which was sufficient for all our applications. As a rule of thumb, the number of dimensions is set to be  $D = 30$ [23].

### 2.2.3. ISOMAP.

We have observed that the point cloud in MDS space can and often is distributed across a lower-dimensional manifold. In addition, the ordering of the points on this manifold correlates with the ordering from the dissimilarity metric. In order to take advantage of this in case it occurs, we use a global geometric framework for nonlinear dimensionality reduction, the so-called Isomap algorithm [30]. The Isomap algorithm is used to unroll the low-dimensional manifold embedded in the high-dimensional MDS space by flattening the curved manifold, preferably into a non-curved shape. This in turn enables easier detection of the dominant directions of variation in the dissimilarity measure across the data in MDS space. Isomap unrolls the manifold while keeping the number of dimensions the same as the original MDS space using the following steps: 1) Clustering the particles[32] to find the  $k$  nearest neighbors based on the Euclidean distance for all  $N$  particles in MDS space. 2) Connecting the  $k$  neighborhoods for each point to construct a proximity graph. 3) Computing the shortest pairwise distance for all pairs of points in the graph. This results in a new  $N \times N$  matrix in which the elements represent the geodesic distance. 4) Embedding the geodesic distance matrix into MDS space by minimizing the stress function, keeping the number of dimensions the same. The parameter  $k$ , which corresponds to the number of neighbors used in Isomap, should be empirically chosen based on the distribution of the particles in MDS space (typically between 4 to 12). We empirically found  $k = 4$

works well. If  $k$  is too high, unrolling cannot be performed properly as neighbors can then be found not just along the manifold, but also via a "shortcut" through empty space. In cases where the successive differences between explained variance on the first three axes are higher than half of the average variance explained, Isomap can reveal the low dimensional latent space. Otherwise, it can not further reduce the number of dimensions due to an isotropic distribution in MDS space.

#### 2.2.4. PRINCIPAL COMPONENT ANALYSIS (PCA).

After unrolling the data in the high dimensional MDS space, we used PCA to identify the direction of the largest variation. Here the axis with maximum variance explained corresponding to the highest eigenvalue was selected as the one embedding the maximum mode of variation in the data measured by the Bhattacharya cost function. Subsequently, by projecting all particles on its main principal component axis, a 1D latent space is created, which corresponds to some of the continuous heterogeneity that exists in the dataset.

#### 2.2.5. RECONSTRUCTION PER BIN.

To visualize the continuous heterogeneity revealed in the latent space, particles are divided into bins, and a single reconstruction for each bin is made. The way in which the particles are distributed over the bins is adapted to the distribution of the particles in the latent space. For a (near) uniform distribution of particles, the bin width is set such that all bins cover the dynamic range of values in latent space. For a more bell-shaped distribution of particles, the bin width is set to be equal to half of the fitted standard deviation of the distribution. The total number of bins is typically chosen in the range of 5-10, but can in principle be chosen arbitrarily, provided there are more than approximately 10 particles per bin. The resulting reconstructions per bin are expected to be more faithful since the particles in each bin are structurally close to each other.

#### 2.2.6. PARTICLE FUSION.

The superparticle reconstructions per bin are made with a template-free method developed previously by us [13], based on earlier work by Evangelides and Horaud [33], the so-called Joint Registration of Multiple Point Clouds (JRMPC).

#### 2.2.7. MODEL-BASED SHAPE PARAMETER ESTIMATION.

We identify the coordinate axis in 1D latent space indicating the heterogeneity parameter. In an experimental 2D NPC dataset, we compare the latent space coordinate to the estimated radius of the ring structure. To estimate the radius, we center all particles by subtracting the mean of all localizations in a particle from the localizations, as in Heydari et al. [11]. Subsequently, the localization coordinates are transformed into polar coordinates. We take the mean of the radial coordinate as an estimation of NPC radius. To estimate the precision of the model we calculate the FWHM of the radius histogram and divide it by the square root of the number of localization events to find the standard error of the mean. In an experimental 3D DNA-origami tetrahedron dataset, we compare the latent space coordinate to the estimated height of the tetrahedron structure. To this end, we manually align the particles in each bin along the  $z$ -axis and project all localizations

on the  $z$ -axis, giving a histogram of  $z$  coordinates of all localizations in the particle. This histogram has two peaks, one corresponding to the three binding sites in the base plate of the tetrahedron, and the other corresponding to the tip of the tetrahedron. By fitting a mixture of two Gaussian distributions to this histogram, we find the height as the difference between the mean of the two fitted Gaussian distributions. In a simulation study on elliptically shaped 2D NPCs we estimate the ellipticity of the ring structure as in Huijben et al.[23] by finding the center of the 8 blobs in each particle using  $k$  means clustering, followed by fitting an ellipse to the centers of the 8 blobs.

### 2.2.8. DATA ACQUISITION OF 2D NPC.

Experimental NPC data was acquired with the following protocol. U2OS Nup96-SNAP cells (Cell Lines Services, from Jan Ellenberg, EMBL) were seeded on collagen-coated 8-well chambers slide (1.5NA, LabTek II #155409) with 3  $\text{\AA}$ tg/mL aphidicolin (Millipore #178273-1MG) and incubated overnight at 37°C. Cells were then pre-fixed with 2% Paraformaldehyde (PFA, Electron Microscopy Sciences #1570-S) for 30 sec, permeabilized with 0.1% Digitonin (RPI #43065-0.1 ) for 30 min and additionally fixed with 2% PFA for 10 min. PFA was quenched with 20 mM Tris pH 8.0 (G-Biosciences #R002) for 5min. Samples were then blocked with 10% Fetal Bovine Serum (FBS, Hyclone #SV30014.03) for 30 min and then stained with 500 nM AlexaFluor647-BG for SNAP-tag (NewEngland Biolabs #S9136S) with 1  $\mu$ M of Dithiothreitol (DTT, American Bioanalytical #AB00490) in 10% FBS for 1 h at room temperature. Samples were prepared for STORM imaging with 100 mM Cysteamine Hydrochloride (MEA, Sigma #M6500-25G), 1% GLOX oxygen scavenger buffer (40  $\mu$ g/mL catalase (Sigma #C40-100MG) and 500  $\mu$ g/mL glucose oxidase (Sigma #G2133-50KU)) and a second buffer (10 mM NaCl, 50 mM Tris pH 8 and 10% Glucose). Images were taken using an Oxford Nanoimager (ONi) STORM microscope with oil-immersion objective (100X, 1.4NA) and laser power density of 4  $kW/cm^2$ . For imaging acquisition, a 641 nm laser was used to excite AF647-BG and a 405 nm laser was used to enhance blinking. Samples were pre-bleached with 190 mW of 641 nm laser (1500 frames), then acquisition was done with 120 mW of the 641 nm laser (5000 frames) and enhanced blinking with 0.2 mW of 405 nm and 120 mW of 641 nm (5,000 frames).

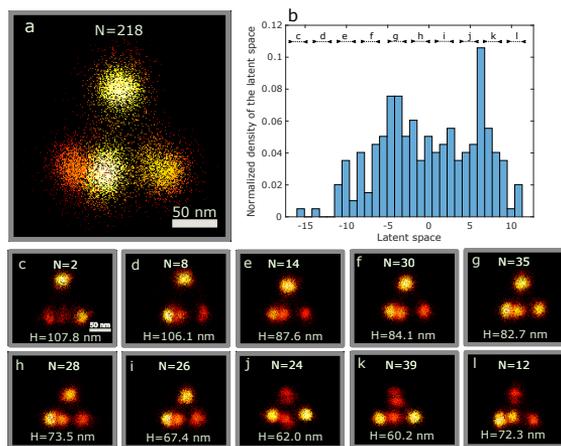
## 2.3. RESULTS

We applied the proposed CHD analysis pipeline to two experimental datasets, a 3D DNA origami tetrahedron dataset, and a 2D NPC dataset. We also investigated limitations of the method by two simulation studies, on the impact of the Degree Of Labeling (DOL) and on the sensitivity to two independent modes of continuous variation.

### 2.3.1. CONTINUOUS DISTRIBUTION OF 3D DNA-ORIGAMI TETRAHEDRON HEIGHT.

We applied our CHD pipeline to a three-dimensional DNA-origami tetrahedron data set, imaged with DNA-PAINT [23, 34]. This data set consists of 218 particles with an edge length of around 100 nm and a height of around 90 nm. There is a variation in tetrahedron height, which was analyzed previously by Huijben et al.[23] with their clustering approach. We have applied our CHD method to this data set as well and the key results are

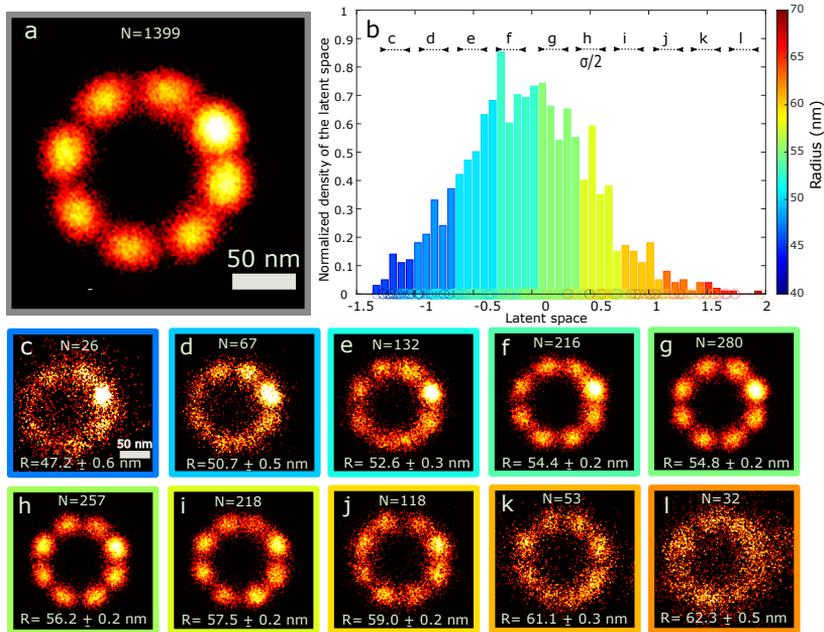
shown in Fig. 2.2. The reconstruction of the whole data set (Fig. 2.2a) was obtained using fast particle fusion based on JRMP [13]. Fig. 2.2b shows the distribution of particles over the latent space coordinate. There is a clear correlation of the found latent space coordinate with a continuous height variation of about 45 nm in the entire data set, providing additional support for the validity of our method. The variance explained in the first PC axis is 23% (clearly larger than in the next directions, which give rise to 14, 9, 5, and 4% variance explained). We divided the latent space into 10 bins with equal length in latent space as the particles are almost uniformly distributed in the latent space. Fig. 2.2c-h shows the reconstructions for each bin. These reconstructions have less elongated blobs than the overall reconstruction of Fig. 2.2a. A 3D reconstruction of each bin can be seen in Supplementary video 1.



**Figure 2.2: Continuous distribution of DNA origami tetrahedron height.** *a)* All particles are registered using fast particle fusion based on the homogeneity assumption. The fusion of structurally different particles causes blurry and elongated blobs. *b)* The distribution of the particles in latent space. *c-l)* Fusion of particles in 10 bins sorted along the latent space coordinate, showing a clear variation in the height of tetrahedrons. Scale bar in *c* applies to *d-l* as well. Colormap represents the localization density.

### 2.3.2. CONTINUOUS DISTRIBUTION OF 2D NPC RADIUS.

The NPC dataset consists of 1,339 particles imaged with 2D STORM (see Methods section). Each particle was picked (cropped) manually from a single SMLM image. The shape of lower-dimensional projections of the distribution of points in MDS space turns out to be closer to an ellipsoid structure as opposed to a more flat shape. As a result, the Isomap step does not provide additional value in this case, and we directly applied PCA to the distribution in MDS space. We attribute the more diffuse topology of the manifold in MDS space to other modes of variation in addition to the dominant one (NPC radius). These confounding modes of variation could be related to other modes of structural variation, and SMLM specific statistical variations such as Degree Of Labeling (DOL).



**Figure 2.3: Continuous distribution of 2D NPC radius.** *a)* Fusion of 1,399 NPCs imaged with STORM using fast particle fusion with homogeneity assumption. *b)* The normalized histogram of the latent space is fitted by a Gaussian distribution, and the size of each bin is selected as half of the fitted standard deviation. Images *c-l* correspond to registered particles inside each division using particle fusion[13]. Scale bar in *c* Applies to *d-l*.

By estimating the size of each particle as explained in the Model-based shape parameter estimation in Method section, we observed a continuous variation in the radius of the scaffold that seems drawn from a Gaussian distribution with a mean of around 55 nm, and a standard deviation of around 4 nm. This continuous heterogeneity causes blurriness in the reconstruction of all particles in the dataset (Fig. 2.3a) that assumes homogeneity of the underlying data. Fig. 2.3b shows the results of analyzing this dataset with our CHD algorithm. There is a very clear correlation between the latent space coordinate and the independently assessed radius, validating our method for discovery of continuous heterogeneity without any prior assumptions. Although the first principal axis shows a very good correlation with radius, the variance explained in this axis is only 4%. This may be related to the more diffuse topology of the MDS manifold compared to the 3D DNA-origami tetrahedron case. We also inspected the second or higher principal axis but did not observe a correlation to another structural feature.

Fig. 2.3c-l shows reconstructions of particles in 10 bins, defined according to the observed bell-shaped distribution of the latent space coordinate (see Methods section). The

results indicate that the registration of particles in the middle bins (Fig. 2.3f-i) leads to better reconstructions compared to the first (Fig. 2.3c-e) and last bins (Fig. 2.3j-l). As a quantitative measure, we calculated the spectral signal to noise[35] (SSNR) curves for each bin. To that end we divided each bin into two halves and registered each part independently using the fast particle fusion approach. After alignment of the two halves, we applied a random rotation in view of the 8-fold rotational symmetry of the structure, thereby avoiding hotspots. Fig. 2.4 shows that for the middle latent space bins (4-7), the SSNR curves are higher than for the first (1-3) and last (8-10) bins, in agreement with the visual quality of the reconstructions per bin in (Fig. 2.3c-l). All SSNR curves level off to a noise plateau for spatial frequencies higher than approximately  $0.15 \text{ nm}^{-1}$ , indicating that the smallest features in the reconstructions are about 6 nm. The reason for the relatively poor reconstruc-

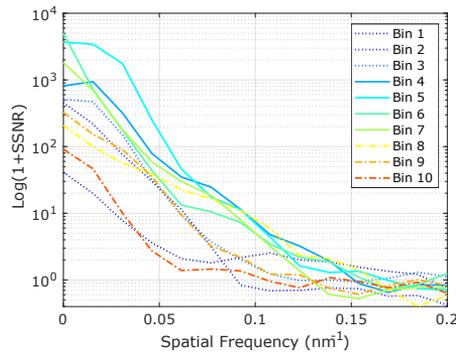


Figure 2.4: **Spectral signal to noise curves for each division.** This figure illustrates that bins 4-7 have higher SSNR than bins 1-3 and 8-10, which confirms the better reconstruction quality of the middle bins.

tion of the extreme bins is not just the lower number of similar particles that contribute to the reconstruction, but also that these bins are contaminated with outlier particles. This suggests an additional value of our proposed method, namely outlier particle detection. Fig. 2.5 shows examples of particles that can be designated as outlier particles compared to randomly selected valid particles. As threshold for the definition of outlier, we take the extreme 1% of the distribution of particles in latent space. Visually, these outlier particles are relatively remote from the expected ring shaped point clouds.

### 2.3.3. SIMULATION ON IMPACT OF DOL.

A low Degree Of Labeling (DOL) of binding sites is a common problem in SMLM. It can be expected that DOL also has an impact on the ability to detect continuous heterogeneity in a dataset, as the randomness of which binding sites are labeled (and which not) affects the imaged structure of each individual particle. We investigated the impact of DOL by a simulation study of a 2D NPC structure. We applied our CHD pipeline to simulated data sets with five DOL values (30, 50, 70, 90, and 100%), and for particle radii drawn from uniform distributions with four ranges. The simulation is performed as Huijben et al. [23] with 250 particles in each dataset. Fig. 2.6 shows the variance explained on the first PC axis

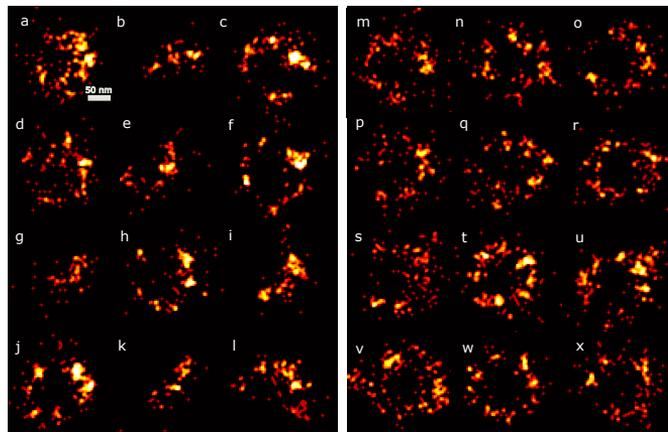


Figure 2.5: **Outliers detected by the CHD pipeline.** Images **a-l** correspond to particles that are in low-density regions of the latent space. As reference, images **m-x** are randomly selected from the whole data set.

(the latent space coordinate) as a function of DOL for the four distributions. As expected the variance explained decreases with decreasing DOL. If we take 50% as a minimum value for this performance criterion, we can conclude that the DOL should be at least in the range 50-70%.

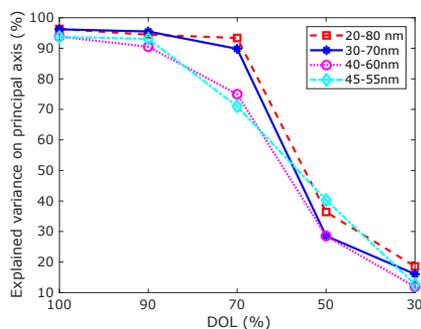


Figure 2.6: **Performance of continuous heterogeneity detection with DOL.** Each curve consists of 5 data points, pertaining to datasets with five DOL values (30, 50, 70, 90, 100%). The radius of the particles is drawn from a uniform distribution with ranges indicated in the legend.

### 2.3.4. SIMULATION ON TWO MODES OF VARIATION.

So far, we have only considered cases in which there was a single, dominant mode of variation. To evaluate the ability of the CHD method to detect two modes of variation simultaneously, we made a simulation study of 2D NPC particles with both a variation in radius (uniform distribution ranging from 30 nm to 70 nm) and a variation in ellipticity of the ring (uniform distribution ranging from 0.6 to 1.0). The outcome was compared to simulations with only a variation in radius or in ellipticity. Fig. 2.7 shows the similarity matrices for these different simulations. The matrices for single-mode variation simulations (Fig. 2.7a-b) were ordered such that the particles are sorted based on the ellipticity or radius of the particles. As a consequence, by moving from left to right in each row in Fig. 2.7a-b, the similarity value decreases since each particle is compared with a less similar structure in terms of ellipticity or radius. This applies to every row in the matrix and creates a diagonal band of higher similarity values. The width and the average hue of this diagonal band in the matrix are related to the sensitivity of the similarity measure to the ellipticity or radius. If we apply the same particle ordering procedure for either the ellipticity or the radius for two-mode variation simulation (Fig. 2.7c-d), it turns out that the width and hue of the diagonal band is limited or seemingly absent. This indicates that the single similarity metric we use might have difficulties to provide sufficient information to detect multiple modes of variation.

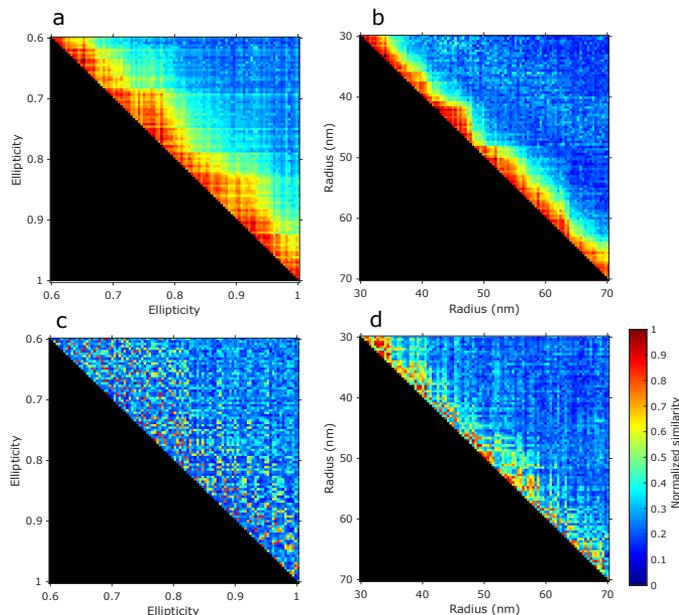
In the simulated dataset with two modes of variation it appeared nevertheless possible to detect both modes. Fig. 2.8 shows the distribution of particles in the 2D latent space (first and the second principal axes), where the color code shows the radius and ellipticity ground truth, respectively. Clearly, the variation in radius corresponds with the first principal axis, while the variation in ellipticity corresponds with the second principal axis. While these initial simulation results are encouraging, we have not been able to detect multiple modes of variation in both experimental datasets we studied. Several confounding factors, in particular a stronger underlabeling, could be limiting in detecting multiple variation modes in experiment.

## 2.4. DISCUSSION

In summary, we have developed a model-free continuous structural heterogeneity tool to sort particles based on a dissimilarity measure. We successfully detected continuous structural variation in different localization microscopy datasets, such as DNA-origami tetrahedrons and NPCs, which led to more faithful fusions.

The method should be applicable to any SMLM dataset that consists of particles that share a similar structure but vary in conformational state. In cryo-EM studying structure variation is applied to e.g. variation of position or motion of side groups. Similar variations could potentially be visible in our approach albeit at a lower resolution due to the imaging modality.

It is not clear how many samples are required per bin to detect structural heterogeneity, let alone how this number of samples depends on localisation precision and degree of labelling. Already for normal averaging, with the assumption of structural homogeneity, it is unclear how the FRC resolution of the reconstruction depends on the number of particles. This is very different from cryo-EM SPA where each extra particle reduces the noise and the improvement follows the expected  $1/\sqrt{N}$  scaling). As best practice, we inspect the

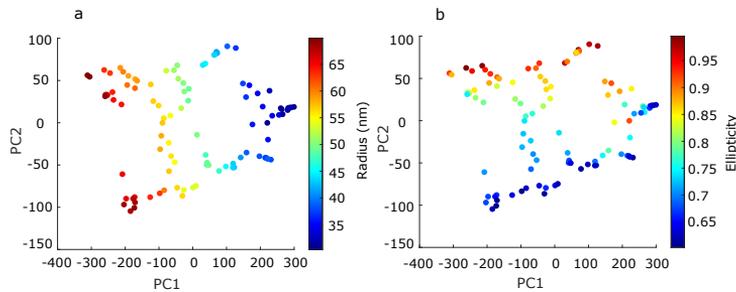


**Figure 2.7: Sensitivity analysis of the cost function to two modes of variation.** Plots illustrating the normalized pairwise cost value (similarity measure) for a variation in radius and/or ellipticity of simulated 2D NPC. **a)** Cost values of a dataset with uniform distribution over ellipticity ratio ranging from 0.6 to 1. Particles are sorted based on their ellipticity from top to down and left to right in the matrix. The values in each column show the sensitivity of a particle with specific ellipticity to the other particles. As a result, the color separation over the diagonal band refers to the sensitivity of the method to the variations in ellipticity. **b)** Cost values of a dataset with uniform distribution over the radius ranging from 30 nm to 70 nm. **c,d)** Cost values of a dataset with two modes of variations. In **c** particles are sorted based on their ellipticity, while in **d** sorting is based on radius. A comparison of the diagonal bands in **c** and **d** shows that the method is more sensitive to variations in radius than in ellipticity.

result and then judge if the numbers have been sufficient.

We found several limitations of the proposed template-free continuous heterogeneity detection method. Firstly, it is not *a priori* clear if the Isomap unrolling step is of use or not. Secondly, picking up modes of variation like variations in DOL, that have no clear geometrical interpretation such as size parameters, turn out in simulation to be too challenging to detect. Thirdly, the detection of multiple continuous modes of shape variation in SMLM data remains unsolved, despite initially hopeful simulation results. This may be due to the poorer quality of experimental data and variation modes that are entangled differently than foreseen in simulation. Finally, the method is based on a single dissimilarity metric that may be expected to have different sensitivities to different modes of variation.

An alternative to the proposed method may be to fit an *a priori* model to the data and



**Figure 2.8: 2D Latent space for two modes of variation.** *a)* Distribution of particles in the 2D latent space color coded respectively with the ground truth radius. *b)* Distribution of particles in the 2D latent space color coded respectively with ellipticity.

then sort the data by the distribution of model parameters obtained by the fit. A major drawback of such an approach, however, is that the outcome of the analysis is prone to biases induced by the assumed model. A more fruitful next step may be to consider multiple metrics that quantify specific features and/or particle similarity for sorting the particles. This can be performed either by a semi-template-free approach with multiple features that are designed to be sensitive to specific modes of variation or by generating sets of more abstract features using deep neural networks and auto-encoders[36].

## REFERENCES

- [1] T. Klein, S. Proppert, and M. Sauer, *Eight years of single-molecule localization microscopy*, *Histochemistry and Cell Biology* **141** (2014), pp. 561–575.
- [2] S. W. Hell, *Microscopy and its focal switch*, *Nature Methods* **6** (2009), pp. 24–32.
- [3] M. Bates, B. Huang, and X. Zhuang, *Super-resolution microscopy by nanoscale localization of photo-switchable fluorescent probes*, *Current Opin. in Chemical Biology* **12** (2008), pp. 505–514.
- [4] R. P. Nieuwenhuizen, K. A. Lidke, M. Bates, D. L. Puig, D. Grünwald, S. Stallinga, and B. Rieger, *Measuring image resolution in optical nanoscopy*, *Nature Methods* **10** (2013), pp. 557–562.
- [5] A. Löschberger, S. van de Linde, M.-C. Dabauvalle, B. Rieger, M. Heilemann, G. Krohne, and M. Sauer, *Super-resolution imaging visualizes the eightfold symmetry of gp210 proteins around the nuclear pore complex and resolves the central channel with nanometer resolution*, *Journal of Cell Science* **125** (2012), pp. 570–575.
- [6] A. Szymborska, A. De Marco, N. Daigle, V. C. Cordes, J. A. Briggs, and J. Ellenberg, *Nuclear pore scaffold structure analyzed by super-resolution microscopy and particle averaging*, *Science* **341** (2013), pp. 655–658.

- [7] M. Kudryashev, D. Castaño-Dañez, and H. Stahlberg, *Limiting factors in single particle cryo electron tomography*. Computational and structural biotechnology journal **1** (2012).
- [8] G. Tang, L. Peng, P. R. Baldwin, D. S. Mann, W. Jiang, I. Rees, and S. J. Ludtke, *Eman2: An extensible image processing suite for electron microscopy*, Journal of Structural Biology **157** (2007), pp. 38–46.
- [9] C. Sieben, N. Banterle, K. M. Douglass, P. Gönczy, and S. Manley, *Multicolor single-particle reconstruction of protein complexes*, Nature Methods **15** (2018), pp. 777–780.
- [10] R. D. Gray, D. Albrecht, C. Beerli, G. H. Cohen, R. Henriques, and J. Mercer, *Nanoscale polarization of the vaccinia virus entry fusion complex drives efficient fusion*, Nat. Microbiol. **4** (2019), p. 1636–1644.
- [11] H. Heydarian, F. Schueder, M. T. Strauss, B. Van Werkhoven, M. Fazel, K. A. Lidke, R. Jungmann, S. Stallinga, and B. Rieger, *Template-free 2D particle fusion in localization microscopy*, Nature Methods **15** (2018), pp. 781–784.
- [12] H. Heydarian, M. Joosten, A. Przybylski, F. Schueder, R. Jungmann, B. v. Werkhoven, J. Keller-Findeisen, J. Ries, S. Stallinga, M. Bates, *et al.*, *3D particle averaging and detection of macromolecular symmetry in localization microscopy*, Nature Communications **12** (2021), pp. 1–9.
- [13] W. Wang, H. Heydarian, T. A. Huijben, S. Stallinga, and B. Rieger, *Joint registration of multiple point clouds for fast particle fusion in localization microscopy*, Bioinformatics **38** (2022), pp. 3281–3287.
- [14] J. Broeken, H. Johnson, D. S. Lidke, S. Liu, R. P. Nieuwenhuizen, S. Stallinga, K. A. Lidke, and B. Rieger, *Resolution improvement by 3D particle averaging in localization microscopy*, Meth. Appl. Fluor. **3** (2015), p. 014003.
- [15] A. Jimenez, K. Friedl, and C. Leterrier, *About samples, giving examples: optimized single molecule localization microscopy*, Methods **174** (2020), pp. 100–114.
- [16] R. Danev, H. Yanagisawa, and M. Kikkawa, *Cryo-electron microscopy methodology: current aspects and future directions*, Trends in Biochemical Sciences **44** (2019), pp. 837–848.
- [17] G. Tang, L. Peng, P. R. Baldwin, D. S. Mann, W. Jiang, I. Rees, and S. J. Ludtke, *EMAN2: an extensible image processing suite for electron microscopy*, Journal of Structural Biology **157** (2007), pp. 38–46.
- [18] J. Frank, *Single-particle reconstruction of biological macromolecules in electron microscopy—30 years*, Quarterly Reviews of Biophysics **42** (2009), pp. 139–158.
- [19] S. H. Scheres, M. Valle, R. Nuñez, C. O. Sorzano, R. Marabini, G. T. Herman, and J.-M. Carazo, *Maximum-likelihood multi-reference refinement for electron microscopy images*, Journal of Molecular Biology **348** (2005), pp. 139–149.

- [20] S. H. Scheres, H. Gao, M. Valle, G. T. Herman, P. P. B. Eggermont, J. Frank, and J.-M. Carazo, *Disentangling conformational states of macromolecules in 3D-EM through likelihood optimization*. *Nature Methods* **4** (2007), pp. 27–29.
- [21] I. S. Gabashvili, R. K. Agrawal, R. Grassucci, and J. Frank, *Structure and structural variations of the escherichia coli 30 s ribosomal subunit as revealed by three-dimensional cryo-electron microscopy*, *Journal of Molecular Biology* **286** (1999), pp. 1285–1291.
- [22] D.-H. Chen, J.-L. Song, D. T. Chuang, W. Chiu, and S. J. Ludtke, *An expanded conformation of single-ring groel-groes complex encapsulates an 86 kDa substrate*, *Structure* **14** (2006), pp. 1711–1722.
- [23] T. A. Huijben, H. Heydarian, A. Auer, F. Schueder, R. Jungmann, S. Stallinga, and B. Rieger, *Detecting structural heterogeneity in single-molecule localization microscopy data*, *Nature Communications* **12** (2021), pp. 1–8.
- [24] M. Chen and S. J. Ludtke, *Deep learning-based mixed-dimensional gaussian mixture model for characterizing variability in cryo-em*, *Nature Methods* **18** (2021), pp. 930–936.
- [25] J. S. H. Danial and A. J. Garcia-Saez, *Quantitative analysis of super-resolved structures using asap*, *Nature Methods* **16** (2019), pp. 711–714.
- [26] S. Hugelier, H. Kim, M. T. Gyparaki, C. Bond, Q. Tang, A. N. Santiago-Ruiz, S. Porta, and M. Lakadamyali, *Eclipse: A versatile classification technique for structural and morphological analysis of super-resolution microscopy data*, *BioRxiv* (2023).
- [27] Y.-L. Wu, P. Hoess, A. Tschanz, U. Matti, M. Mund, and J. Ries, *Maximum-likelihood model fitting for quantitative analysis of smlm data*, *Nature Methods* **20** (2023), pp. 139–148.
- [28] R. Henderson, *Avoiding the pitfalls of single particle cryo-electron microscopy: Einstein from noise*, *Proceedings of the National Academy of Sciences* **110** (2013), pp. 18037–18041.
- [29] A. Mead, *Review of the development of multidimensional scaling methods*, *Journal of the Royal Statistical Society: Series D* **41** (1992), pp. 27–39.
- [30] J. B. Tenenbaum, V. d. Silva, and J. C. Langford, *A global geometric framework for nonlinear dimensionality reduction*, *Science* **290** (2000), pp. 2319–2323.
- [31] B. Jian and B. C. Vemuri, *Robust point set registration using Gaussian mixture models*, *IEEE Trans. Pat. Ana. Mach. Int.* **33** (2010), pp. 1633–1645.
- [32] A. K. Jain, M. N. Murty, and P. J. Flynn, *Data clustering: a review*, *ACM Computing Surveys (CSUR)* **31** (1999), pp. 264–323.
- [33] G. D. Evangelidis and H. Radu, *Joint alignment of multiple point sets with batch and incremental expectation-maximization*, *IEEE Trans. Pat. Ana. Mach. Int.* **40** (2017), pp. 1397–1410.

- [34] R. Iinuma, Y. Ke, R. Jungmann, T. Schlichthaerle, J. Woehrstein, and P. Yin, *Polyhedra self-assembled from DNA tripods and characterized with 3D DNA-PAINT*, *Science* **344** (2014), pp. 65–69.
- [35] M. Unser, B. L. Trus, and A. C. Steven, *A new resolution criterion based on spectral signal-to-noise ratios*, *Ultramicroscopy* **23** (1987), pp. 39–51.
- [36] M. Sewak, K. Sanjay, and R. Hemant, *An overview of deep learning architecture of deep neural networks and autoencoders*, *Journal of Computational and Theoretical Nanoscience* **17** (2020), pp. 182–188.

## DATA AND CODE AVAILABILITY

Data and codes are publicly available. Updated versions of the software can be downloaded from <https://gitlab.tudelft.nl/imphys/ci/chd>. The single molecule localization data is accessible via 4TU.research repository at [https://data.4tu.nl/private\\_datasets/ML40deqg5qxT0aLstirZjiuZg82GLELr09FSK-qzB0s](https://data.4tu.nl/private_datasets/ML40deqg5qxT0aLstirZjiuZg82GLELr09FSK-qzB0s).

## ACKNOWLEDGEMENTS

We thank Dr. Melikian Gregory and Dariana Torres Rivera for generously sharing the 2D NPC data. This work has been supported by the Dutch Research Council (NWO), VICI grant no. 17046 for B.R. and W.W.

## AUTHOR CONTRIBUTIONS STATEMENT

S.H developed the method, wrote the code, performed the simulations, and analyzed data. B.R and S.S initiated and directed the research. The paper was written by all authors.

# 3

## DETECTING CONTINUOUS STRUCTURAL HETEROGENEITY IN SMLM DATA WITH A POINT CLOUD VARIATIONAL AUTO-ENCODER

The low degree of labeling and limited photon count of fluorescent emitters in single molecule localization microscopy results in poor quality images of macro-molecular complexes. Particle fusion provides a single reconstruction with high signal-to-noise ratio by combining many single molecule localization microscopy images of the same structure. The underlying assumption of homogeneity is not always valid, heterogeneity can arise due to geometrical shape variations or distinct conformational states. We introduce a Point Cloud Variational Auto-Encoder that works directly on 2D and 3D localization data to detect multiple modes of variation in such datasets. The computing time is on the order of a few minutes, enabled by the linear scaling with dataset size, and fast network training in just four epochs. The use of lists of localization data instead of pixelated images leads to just minor differences in computational burden between 2D and 3D cases. With the proposed method, we detected radius variation in 2D Nuclear Pore Complex data, height variations in 3D DNA origami tetrahedron data, and both radius and height variations in 3D Nuclear Pore Complex data. In all cases, the detected variations were on the few nanometer scale.

1

---

<sup>1</sup>Sobhan haghparast, Yi Zhang, Qian Tao, Sjoerd Stallinga and Bernd Rieger *Detecting Continuous Structural Heterogeneity in Single Molecule Localization Microscopy Data with a Point Cloud Variational Auto-Encoder.*

### 3.1. INTRODUCTION

Single Molecule Localization Microscopy (SMLM) is a widely applied super resolution microscopy technique that enables below diffraction limit imaging [1–4]. The resolution of SMLM is mainly limited by the localization precision and the degree of labeling to values typically between 10-50 nm [5]. If chemically identical copies of a macromolecular structure exist in the image, the resolution can be increased by methods to fuse these so-called particles to an overall reconstruction of the targeted macromolecular structure [6–9]. These particle fusion methods combine the localizations of many particles to increase the signal-to-noise in the final reconstruction. Such methods are based on the assumption of homogeneity of the underlying data, that is, it is assumed that all imaged particles are structurally fully identical. This homogeneity assumption, however, does not necessarily hold, not even for chemically identical structures. Heterogeneity in the set of to-be fused particles can originate from biological changes [10] or sample preparation [11] and can be continuous or in distinct classes [12]. In Cryogenic Electron Microscopy (Cryo-EM) [13, 14] heterogeneity detection of distinct classes is commonly applied [15–17]. SMLM data are point sets (coordinates of all localization events) that thus differ from pixelated images of Cryo-EM. Moreover, the complex photophysics of the molecular on-off switching leads to SMLM data that often come with repeated localizations of the same emitter such that the number of localizations is not a linear representation of the actual density of fluorophores [18, 19]. These differences in acquisition and data stand in the way of direct application of Cryo-EM methods for particle fusion and for detecting heterogeneity in the underlying dataset, such as the use of learning-based mixed-dimensional Gaussian mixture models [20].

Detecting distinct heterogeneity in SMLM data can be performed using template-free clustering approaches such as done by Huijben et. al [12]. Such methods, however, cannot be used to identify continuous variations such as geometric shape parameters of the imaged macromolecular structure. Recently, a template-free continuous structural heterogeneity detection method has been proposed by us [21]. A key limitation of the method is its computational complexity, which scales as  $N^2$ , with  $N$  the number of particles. In practice, the dataset size is therefore limited to a few hundred particles, which can be analyzed in a few hours. In addition, the ability to detect multiple modes of variation and non-geometrical modes of variation, such as the Degree Of Labeling (DOL) is limited. Alternative approaches for detecting continuous structural heterogeneity include methods with varying requirements and limitations. For example, a proposed learning-based method [22] is heavily dependent on manual annotation during the training phase. This process is not only time-consuming, but is also prone to human error, and has poor scalability for large datasets. Similarly, the statistical pattern recognition method ECLIPSE [23] requires precise segmentation and a high signal-to-noise ratio to operate optimally, which limits its effectiveness in scenarios with low DOL or complex environments. LocMoFit [24], a model-based fitting technique, utilizes a parametric geometric model based on prior knowledge, constraining its adaptability to cases where this prior knowledge is available and reliable.

Here, we propose to overcome these limitations by applying a Variational Auto-Encoder (VAE) to extract generative features from the SMLM dataset, and thereby identifying continuous modes of variation. The VAE [25] is a probabilistic generative model designed to

map the input data, the SMLM point dataset, to a lower-dimensional latent space while preserving the essential statistical properties of the data distribution. Unlike traditional autoencoders that directly learn a deterministic encoding-decoding mapping, VAEs learn a probabilistic mapping by encoding inputs into distributions over latent variables, creating a continuous and smooth latent space where similar data points are mapped to nearby regions. This continuous nature of the latent representation enables interpolation between different data points (and the generation of new samples). In the context of SMLM data, the continuous latent space makes VAEs particularly suitable for characterizing and detecting heterogeneous spatial distributions, as the continuous latent space naturally captures the gradual variations in particle clustering patterns. Application of the VAE neural network architecture directly to the point datasets of SMLM, rather than to pixelated datasets, provides a major step to lower memory requirements and processing speed. Moreover, it provides access to multiple modes of variation via the different latent space dimensions, and is fully template free, minimizing the need for human interaction and avoiding user bias.

In view of the specific application to point clouds, we will refer to the proposed method as Point Cloud Variational Auto-Encoder (PC-VAE) in the following. We evaluated our PC-VAE network on multiple experimental datasets with different numbers of particles and structures, namely: Nuclear Pore Complex (NPC) data imaged in 2D and 3D with Stochastic Optical Reconstruction Microscopy (STORM) [26], and DNA origami tetrahedrons imaged in 3D with DNA Points Accumulation for Imaging in Nanoscale Topography (DNA-PAINT) [27].

## 3.2. METHODS

### 3.2.1. POINT CLOUD VARIATIONAL AUTO-ENCODER FRAMEWORK (PC-VAE)

The purpose of applying VAE to SMLM data is to learn a continuous latent representation of the point cloud that captures the spatial heterogeneity inherent in the particles. The training process leverages a mini-batch-based stochastic gradient descent (SGD) strategy: In each iteration, a mini-batch of  $M$  particles is sampled, and the model computes the loss per particle, which combines reconstruction accuracy and latent space regularization. The loss, averaged across the  $M$  particles in the mini-batch, is used to update the model parameters via backpropagation. The optimization process is conducted in multiple epochs, ensuring that the entire dataset is traversed multiple times. As a result, the latent space induced by the trained encoder can effectively capture the spatial diversity and heterogeneity across all particles in the dataset. After training, the model encodes each particle into its respective latent representation, effectively capturing its spatial features for downstream tasks of identifying the structural heterogeneity.

To illustrate the model's functionality, we take a single particle with  $N$  localizations as an example. Such a particle with  $N$  localizations is denoted as a matrix  $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N] \in \mathbb{R}^{N \times p}$ , with  $p = 2, 3$  the spatial dimension of the dataset. The VAE learns a latent representation  $\mathbf{z}$  in a lower-dimensional space  $\mathbb{R}^d$  ( $d \ll N$ ) through an encoder network  $q_\phi(\mathbf{z}|\mathbf{X})$  and a decoder network  $p_\theta(\mathbf{X}|\mathbf{z})$ , where  $\phi$  and  $\theta$  are the learnable parameters of the encoder and decoder. The encoder maps the input  $\mathbf{X}$  to a multivariate Gaussian distribution in the latent space, which can be described by its mean vector  $\mu$  and standard deviation

$\sigma$ , from which we use the reparameterization trick [25] to obtain the latent representation of  $\mathbf{X}$ . The reparameterization trick is a technique for making the sampling process differentiable, which is essential for training the VAE using backpropagation:

$$\begin{aligned}\mu, \sigma &= f_\phi(\mathbf{X}) \\ \mathbf{z} &= \mu + \sigma \odot \epsilon,\end{aligned}\tag{3.1}$$

where  $f_\phi$  is the forward operation of the encoder neural network,  $\odot$  denotes the element-wise product, and  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ . With the reparameterization trick, gradients can be computed with respect to the encoder parameters  $\phi$ , enabling end-to-end training of the VAE using backpropagation. The decoder network  $p_\theta(\mathbf{X}|\mathbf{z})$  aims to reconstruct the cloud of localization points from its latent sample representation. Given a latent vector  $\mathbf{z} \in \mathbb{R}^d$ , the decoder learns to generate a matrix of spatial coordinates  $\hat{\mathbf{X}} = [\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_N] \in \mathbb{R}^{N \times p}$  that follows the original localizations.

To effectively encode and reconstruct the point cloud, we need to design a neural network architecture that can handle the properties of point clouds, which can be considered as sets. The key property we need for processing point cloud data is permutation invariance, as the ordering of points in a point cloud should not affect its representation. Formally, for any permutation matrix  $\mathbf{P} \in \{0, 1\}^{N \times N}$ , it satisfies:

$$f_\phi(\mathbf{P}\mathbf{X}) = f_\phi(\mathbf{X}).\tag{3.2}$$

This property is crucial for SMLM data since localizations are inherently orderless and any meaningful feature extraction should be invariant to their input ordering. To achieve this permutation invariance, we adopt a PointNet-based architecture [28] for our encoder network, which processes each point cloud independently and then aggregates information through symmetric operations. Specifically, in our implementation, we employ three layers of shared-weight multilayer perceptrons (MLPs) with output dimensions of 64, 128, and 1024 to transform each  $p$ -dimensional localization (coordinates) into a higher-dimensional feature space. This transformation enables the network to capture more complex spatial relationships within the point cloud. Each MLP layer is followed by batch normalization (BN) and rectifier linear unit (ReLU) activation functions. The use of identical MLP weights when processing each localization ensures that the output remains unchanged regardless of how the input localizations are ordered. This architectural design directly enforces permutation invariance in point cloud processing. We further apply mean pooling as the final symmetric aggregation function to get the latent representation vector. This differs from the original PointNet architecture [28], which uses max pooling for classification and segmentation tasks. Our choice of mean pooling is motivated by the generative nature of our task. While max pooling selectively preserves the most prominent features useful for discrimination tasks, mean pooling retains channel-wise average spatial information for robust latent representation and reconstruction.

For the VAE, the dimensionality of the compressed representation needs manual selection. This often involves a balance between capturing essential variability in the data while maintaining computational efficiency and interpretability. Although there is no universally optimal number of latent dimensions, the choice is guided by empirical testing. The criterion of dimensionality selection is to ensure that the model can effectively cap-

ture key generative features without overfitting to noise or irrelevant details. Through iterative evaluation, it is possible to identify a latent space size that balances these factors, enabling the model to capture the relevant modes of variation while avoiding redundancy of latent dimensions. For our architecture, the number of latent dimension is empirically selected to  $d = 8$  based on this criterion.

Our proposed decoder network progressively reconstructs the coordinates through a series of expanding transformations. The network first expands the latent vector through fully connected MLPs (256, 512, and 1024 neurons) with hyperbolic tangent (Tanh) activations and batch normalization. We choose the Tanh activations in the decoder since its bounded output range (-1 to 1) naturally suits the generation of normalized point coordinates during reconstruction. A skip connection is used at the 1024 neuron layer for direct information flow. The skip connection allows information to bypass intermediate layers and flow directly from an earlier layer to a later one for stable backpropagation during training [29]. The network then applies two 1024 channel localized feature transformations through 1D convolutions, followed by a 4-headed self-attention [30] that captures global dependencies between different regions of the point cloud. Finally, the network outputs the spatial coordinates of the reconstructed particles. This architecture enables the decoder to reconstruct both the overall structure and the local density variations characteristic of the SMLM data. An overview of the pipeline is shown in Figure 3.1.

### 3.2.2. LOSS FUNCTION

. The loss function of our proposed method consists of two components that balance reconstruction quality and latent space regularization:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{recon}} + \beta \mathcal{L}_{\text{KLD}}, \quad (3.3)$$

where  $\beta \in \mathbb{R}^+$  is the balancing coefficient. To avoid dependence on the ordering of localizations, we measure the quality of the reconstruction using the symmetric Chamfer distance (CD) [31]. The Chamfer distance computes the average minimum distance between points in two sets:

$$\mathcal{L}_{\text{recon}} = \frac{1}{N} \left( \sum_{\mathbf{x} \in \mathbf{X}} \min_{\hat{\mathbf{x}} \in \hat{\mathbf{X}}} |\mathbf{x} - \hat{\mathbf{x}}|_2^2 + \sum_{\hat{\mathbf{x}} \in \hat{\mathbf{X}}} \min_{\mathbf{x} \in \mathbf{X}} |\mathbf{x} - \hat{\mathbf{x}}|_2^2 \right). \quad (3.4)$$

We follow the ordinary use of the Kullback-Leibler divergence (KLD) to regularize the latent space to follow the  $d$ -dimensional Gaussian distribution which has a closed form:

$$\mathcal{L}_{\text{KLD}} = D_{\text{KL}}(q_{\phi}(\mathbf{z}|\mathbf{X}) || \mathcal{N}(0, \mathbf{I})) = \frac{1}{2} \sum_{i=1}^d (\mu_i^2 + \sigma_i^2 - \log(\sigma_i^2) - 1). \quad (3.5)$$

The loss function is computed for each particle in a sampled mini-batch and then averaged over all particles in the mini-batch.

### 3.2.3. TRAINING DETAILS

. We implemented all our network-related codes with PyTorch 2.0 [32] on a workstation with an NVIDIA RTX A40 GPU and an Intel Xeon (3.4 GHz, 4 cores) CPU. We conducted

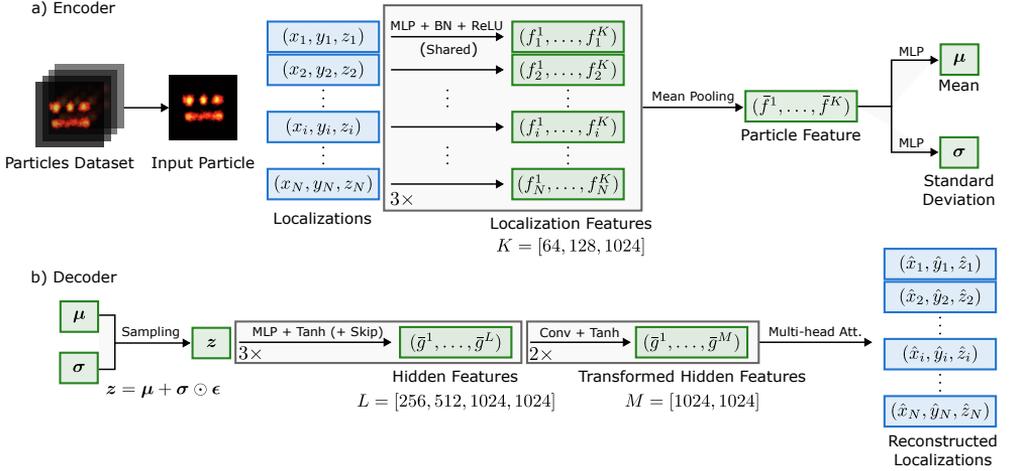


Figure 3.1: An overview of the proposed PC-VAE framework. For the sake of simplicity, we show the processing steps for a single particle only. The loss function is computed by averaging over a mini-batch of  $M$  sampled particles. In the encoder, each input localization point is transformed independently through three parallel MLP blocks to higher dimensions. These MLP blocks are applied to each localization individually, meaning that every localization undergoes the same transformation regardless of its position in the sequence. The resulting features are pooled to generate a global particle representation, which is then passed through two separate multi-layer MLPs to produce Gaussian latent variables. The encoder is designed to be invariant to the order of the input points. In the decoder, samples drawn from the Gaussian distribution are transformed to match the shape of the input localizations. In the figure, Skip denotes the skip connection, and Multi-head Att. denotes the multi-head attention mechanism. The reconstruction loss function also ensures that the order of the reconstructed points is irrelevant, maintaining the framework's robustness to input permutations.

in-house ablation experiments via grid search to determine the hyperparameters for the best model performance and generalization capabilities in our settings. For the results presented here, we used a mini-batch scheme with a batch size of  $M = 8$ . Next, we use the Adam [33] optimizer with an initial learning rate of  $10^{-4}$  and a weight decay of  $10^{-5}$ . Although VAE is known for its ability to capture the representation of the underlying data, it is also known to be difficult to train with a fixed value of  $\beta$ [34]. Thus, we used a Sigmoid warm-up of the KLD in the first 5000 optimization steps. The stochastic nature of the VAE training, facilitated by the reparameterization trick, ensures that the learning process effectively captures the underlying data distribution.

#### 3.2.4. BINNING AND VISUALIZATION

. To visualize the modes of variations, we ordered the input particles based on each latent unit. This results in 8 different orderings of the particles. The ordered particles are then binned in 20 bins of equal particle number. Next, the particles in each bin are registered together, resulting in a so-called super particle per bin. Then we register the 20 super particles with respect to each other and generate reconstructions per bin for each latent unit. For registration of each point cloud we employed the fast particle fusion approach of Wang et al.[9]. Each super particle then represents a specific state or shape of the particles, ordered according to the bin number.

#### 3.2.5. DATA DESCRIPTION.

Our proposed algorithm works on datasets consisting of point clouds ( $x, y$ , and depending on the data also,  $z$  coordinates) derived from SMLM experiments. By segmenting particles from the field of view, a dataset is formed. These datasets typically have a few hundred to a few thousand particles where each particle has typically on the order of hundred localizations. The dimensionality can be 2D or 3D depending on the image acquisition. Here we use 2D and 3D datasets [21, 35] of Nuclear Pore Complexes (NPCS), and a DNA origami tetrahedron dataset[21], with in total 1,399, 3,810 and 218 particles, respectively.

#### 3.2.6. MODEL BASED PARAMETER ESTIMATION

. In order to verify our PC-VAE approach we need to compare the latent space ordering with independent estimates for the different parameters such as height and radius. As these parameters are not known for experimental data we employ a model based approach described here. For the experimental 2D NPC dataset, we estimate the radius as follows. All particles are centered by subtracting the mean localization from all localizations of a particle, as described in Heydarian et al. [36]. The localization coordinates are then transformed into polar coordinates, and the mean of the radial coordinate is taken as the estimated NPC radius.

For the experimental 3D DNA-origami tetrahedron dataset, we compare the latent space coordinate to the estimated height of the tetrahedron structure. Particles are projected onto the  $z$ -axis to generate a histogram of  $z$ -coordinates. This histogram exhibits two peaks, corresponding to the three binding sites at the base plate of the tetrahedron

and the top of the tetrahedron, respectively. A mixture of two Gaussian distributions is fitted to this histogram, and the height is calculated as the difference between the means of the two fitted Gaussian distributions.

For the 3D NPC dataset, we detect both the radius and height of the structure. To estimate the radius, the particles are centered as in the 2D case, and the median of the radial coordinate in polar space is used. To determine the height, all localizations are projected onto the  $z$ -axis. By analyzing the  $z$ -coordinate histogram, height is estimated as the difference between the median of projected points above and below the central plane defined by the mean of all particles.

To evaluate the precision of each parameter estimate, we calculate the full width at half maximum (FWHM) of the respective histogram and divide it by the square root of the number of localizations to determine the standard error of the mean.

### 3.3. RESULTS

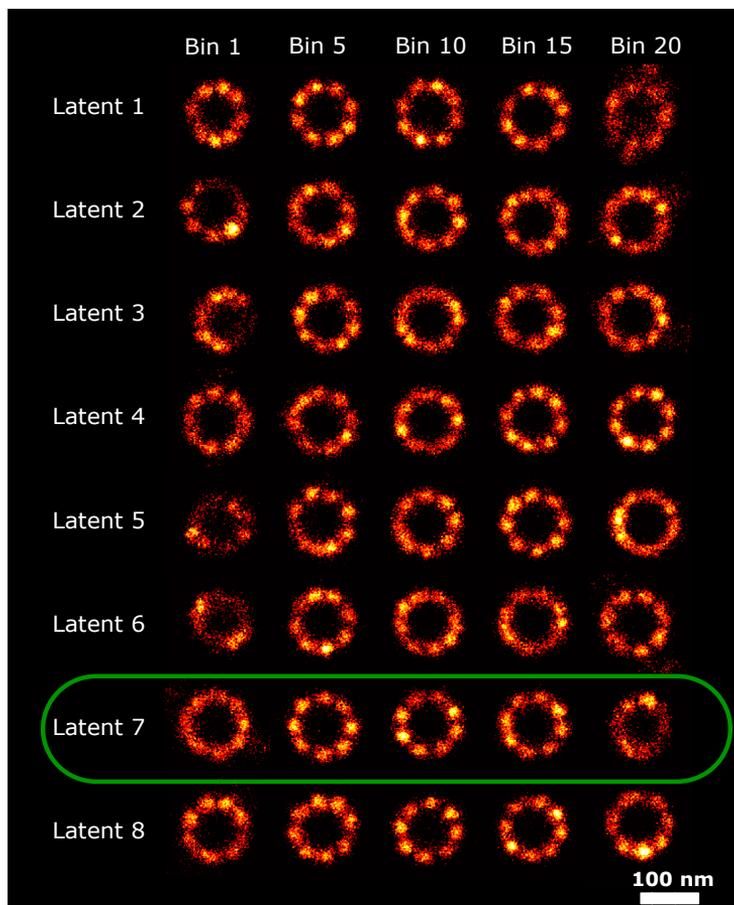
We applied our proposed PC-VAE method on the above mentioned three experimental datasets and measured different modes of variation. Validation of these template-free modes of variation can only be performed by eye or by an independent model. We therefore show and analyze the super particles obtained from registered particles in each bin for all latent dimensions.

#### 3.3.1. 2D NPC DATASET

. This dataset consists of 1,399 particles, with a continuous distribution of ring radii with a mean around 55 nm, and on average 155 localizations per particle. Since each latent space is organized based on the generative features of the data, particles within each bin are considered homogeneous with respect to the generative feature to which the latent space is most sensitive. In Figure 3.2 we show super particles for bin number 1, 5, 10, 15 and 20, respectively, in all 8 latent dimensions. This figure shows that latent space number 7 has the clearest correlation with the radius. Other latent dimensions also carry information about the radius as will be explained in the discussion section. A variation of around 25 nm in radius is detected by the PC-VAE (see Figure 3.3), similar to our earlier continuous heterogeneity detection (CHD) method[21], applied to the same data, but resolves this range into twice as many bins. This is achieved because the bins are more homogeneous, allowing for a finer detection compared to the multi dimensional scaling (MDS) based approach.

#### 3.3.2. 3D TETRAHEDRON DATASET

. This dataset consists of 218 particles with 4 binding site per structure. The average number of localizations per particle is 5,188, giving over a 1,000 localizations per binding site, which is typical for PAINT imaging. Figure 3.4 shows super particles per bin of just 11 particles (20 bins in total, 20th bin has 9 particles) ordered along the 8 latent dimensions. We manually selected latent dimension 2 for analyzing the relation of this latent dimension to tetrahedron height (see Figure 3.5). The PC-VAE framework is able to detect continuous



*Figure 3.2: Reconstructions ordered along all latent dimensions for the 2D-NPC dataset. Super-particles reconstructed by registration of particles in 20 bins for all 8 latent space dimensions. For the purpose of concise illustration we only show bins 1, 5, 10, 15 and 20. The reconstructions for latent dimension 7 shows continuous heterogeneity in the radius of the NPC ring. The scale bar applies to all images.*

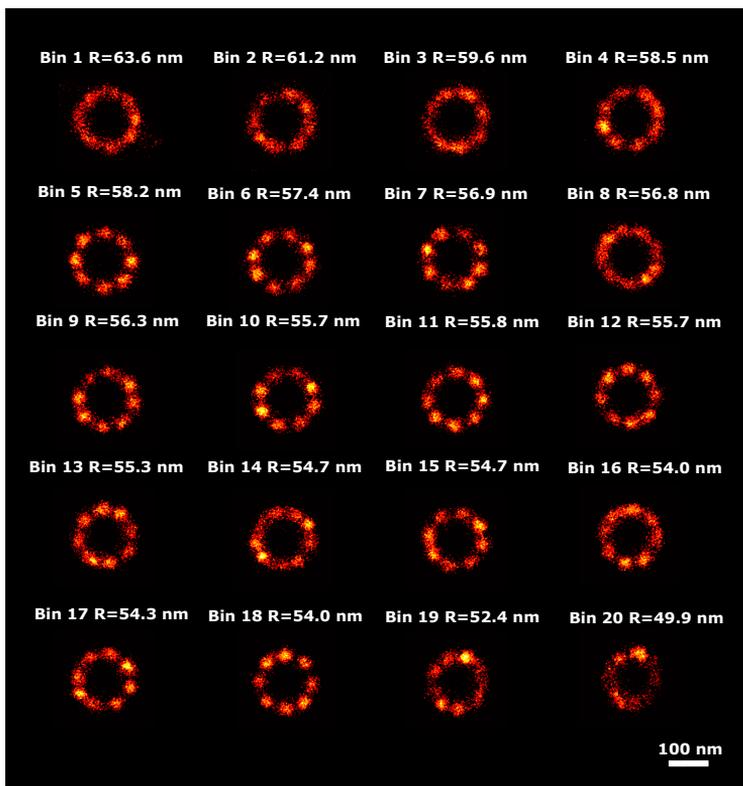
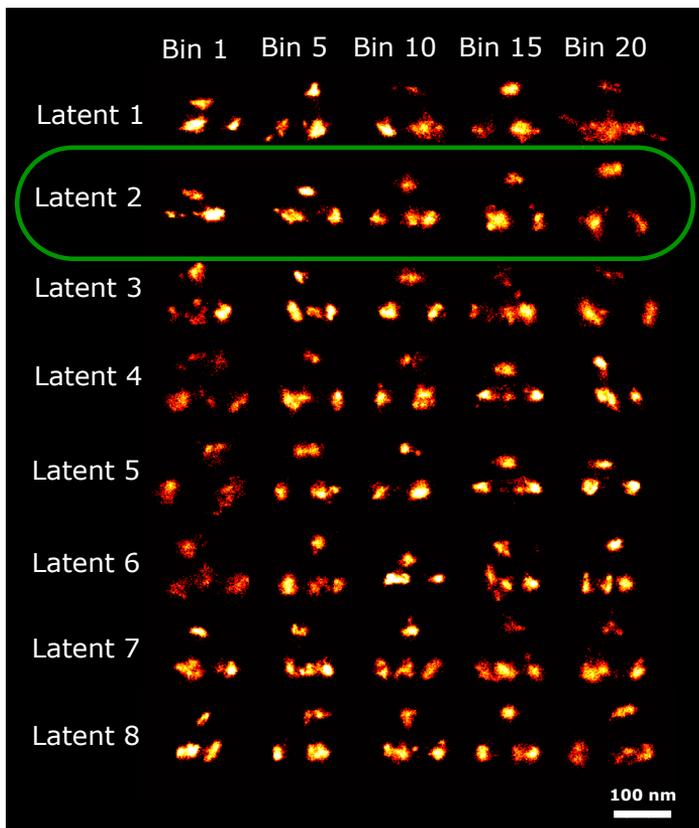


Figure 3.3: **Detected radius variation in the 2D-NPC dataset.** Super particles for all 20 bins ordered along latent dimension 7, with estimated radius, ranging from 63.6 nm to 49.9 nm. The uncertainty of the radius estimation is 0.2 nm for all bins. The scale bar applies to all images.

changes in height in 20 bins ranging from 42.4 nm to 107.3 nm. The variation of height in this dataset was initially reported in three distinct classes (45 nm, 65 nm, 95 nm) by Huijben et al.[12]. Later, the height variation in this dataset was analyzed by us with CHD[21] in ten bins from 60 nm to 107 nm. Similar to the 2D-NPC dataset, the resolving power along the chosen latent dimension for the structural parameter is better for PC-VAE than for CHD.



*Figure 3.4: Reconstructions ordered along all latent dimensions of the 3D tetrahedron dataset. Super-particles reconstructed by registration of particles in 20 bins for all 8 latent space dimensions. For the purpose of concise illustration we only show bins 1, 5, 10, 15 and 20. The reconstructions for latent dimension 2 shows continuous heterogeneity in the height of the tetrahedrons. The scale bar applies to all images.*

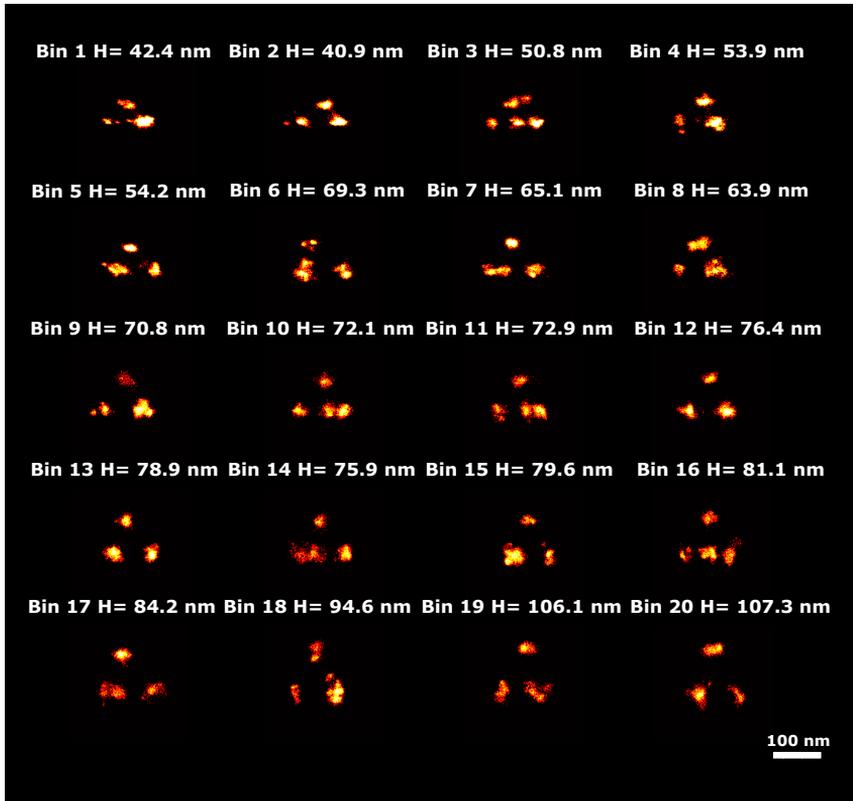


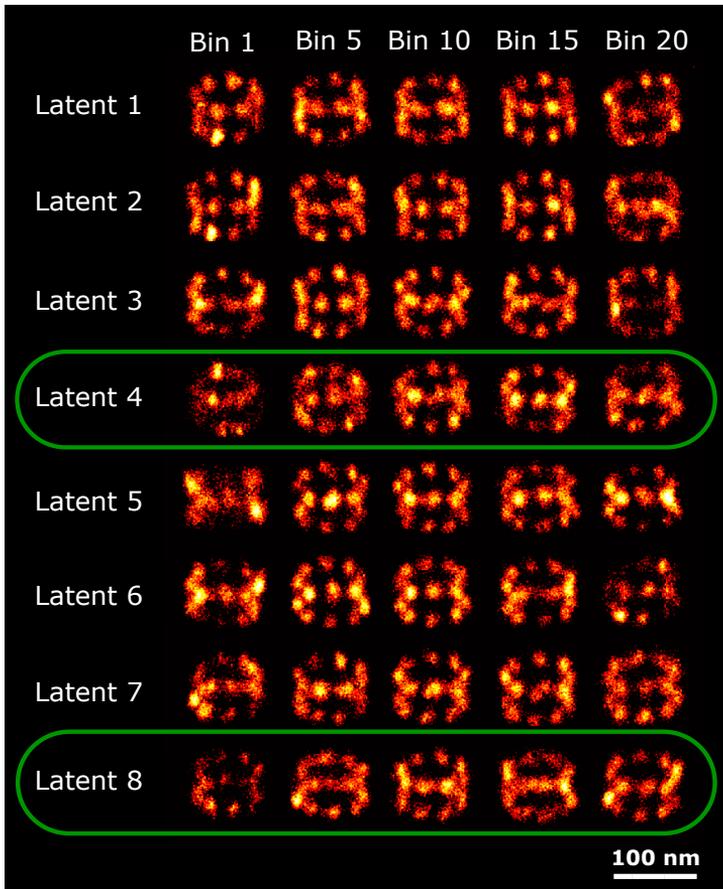
Figure 3.5: **Detected height variation in the 3D tetrahedron dataset.** Super particles for all 20 bins ordered along latent dimension 2, with estimated heights, ranging from 42.4 nm to 107.3 nm. The average model estimation precision is around 0.1 nm for all bins. The scale bar applies to all images.

### 3.3.3. 3D NPC DATASET.

This dataset consists of 3,810 cylindrically shaped 3D NPC structures imaged with 3D STORM. Each structure has 32 binding sites in two rings and the average number of localizations per particle is around 80. Recently, this dataset was analyzed by Wang et al. [35] and compared against an electron microscopy derived representation. This study revealed that the distance between the two rings of the NPC is approximately 48.5 nm, with each ring having a radius of around 55 nm. While variations in radius are expected, they have not been directly measured with light microscopy. Additionally, a 5 nm height difference has been observed between in situ NPCs in HeLa and HEK cells, which could correspond to distinct functional states, such as dilation of the inner ring [37, 38]. Our recent CHD approach[21] is not well-suited for such large datasets, in view of the scaling of the computational complexity with the square of the number of particles. Furthermore,

the CHD method struggles with datasets that have a low Degree Of Labeling (DOL), particularly when dealing with complex datasets that include multiple modes of variation. As a result, it failed to identify continuous structural heterogeneity in this dataset.

Figure 3.6 shows the reconstructed super particles for all latent dimensions, indicating both radius and height variations in latent dimensions 4 and 8. In Figure 3.7 and Figure 3.8 we show the reconstructions for all 20 bins ordered along those latent dimensions, along with the estimated radius and height. This results in detected continuous heterogeneity in the radius ranging from 51.9 nm to 56.7 nm and a height variation ranging from 43.7 nm to 51.7 nm.



*Figure 3.6: Reconstructions ordered along all latent dimensions of the 3D-NPC dataset. Super-particles reconstructed by registration of particles in 20 bins for all 8 latent space dimensions. For the purpose of concise illustration we only show bins 1, 5, 10, 15 and 20. The reconstructions for latent dimension 4 shows continuous heterogeneity in the radius of the NPCs, latent dimension 8 in the height of the NPCs. The scale bar applies to all images.*



Figure 3.7: **Detected height variation in the 3D NPC dataset.** Super particles for all 20 bins ordered along latent dimension 8, with estimated heights, ranging from from 43.7 nm to 51.7 nm. The average model estimation precision is around 0.3 nm for all bins. The scale bar applies to all images.

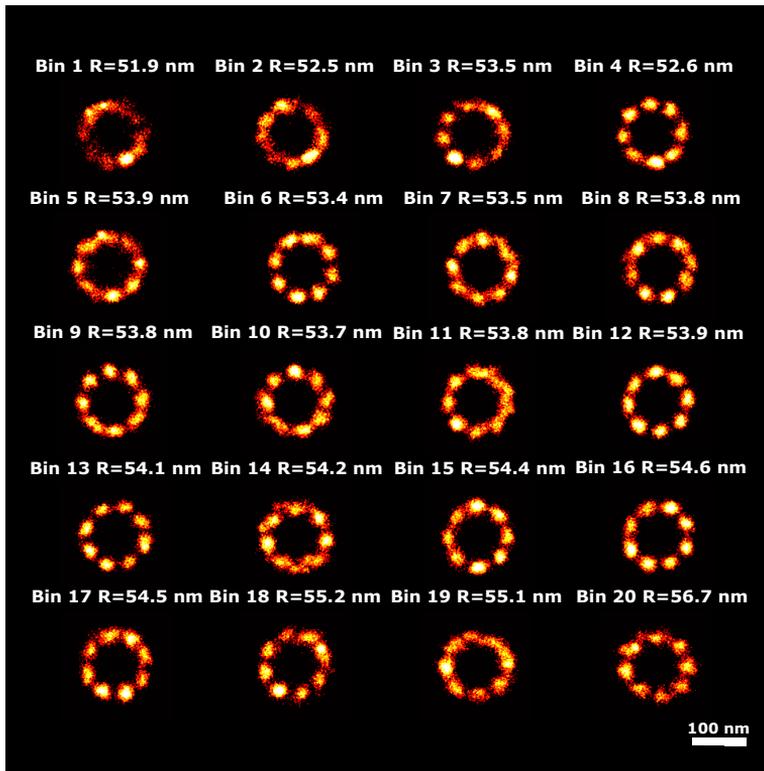


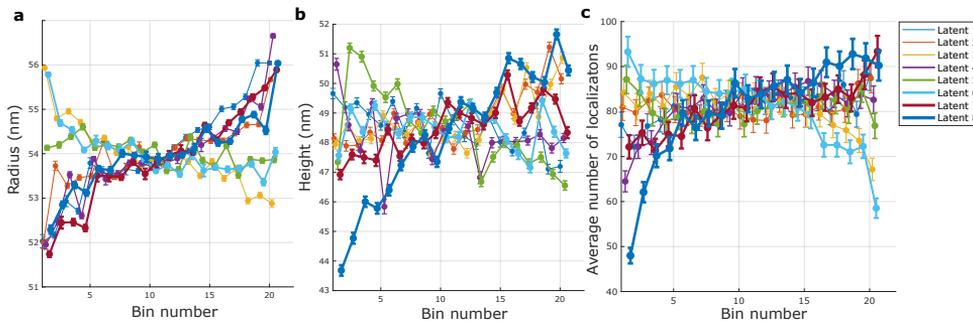
Figure 3.8: **Detected radius variation in the 3D NPC dataset.** Super particles for all 20 bins ordered along latent dimension 4, with estimated radii, ranging from from 51.9 nm to 56.7 nm. The average model estimation precision is around 0.2 nm for all bins. The scale bar applies to all images.

### 3.4. DISCUSSION AND OUTLOOK

VAEs are commonly used for (lossy) compression of data while remaining the most important data structural components. This happens mainly by detecting the generative features or the most important information to reproduce the data. By mapping the input data into a latent space, VAEs aim to capture the underlying structure in a compressed representation. In this study, we have developed a VAE based on point clouds to capture the low-dimensional representations of SMLM data, with the goal of identifying the inherent continuous structural variations. Our PC-VAE method demonstrates major improvements in computational efficiency compared to the MDS-based CHD approach [21]. First of all, the computational time scales linearly with the number of particles, as opposed to the quadratic scaling of computing the pairwise distance metric in the CHD approach. Second, the PC-VAE method works directly on lists of localizations, which does not lead to significant differences in computational time for 2D and 3D cases. Third, the learning phase of the algorithm converges very quickly, only 4 epochs turned out to be necessary for the cases that were considered. The resulting computing times are very short indeed. The 2D NPC dataset with 1,399 particles takes 24 seconds per epoch, the 3D tetrahedron dataset with 218 particles takes 13 seconds per epoch, and the 3D NPC dataset with 3,810 particles takes 279 seconds per epoch. This results in typical computing times on the order of a few minutes. Even with this effort, the CHD method was not able to detect height and radius variations because of the lack of sensitivity of the cost function to complex structures such as the NPC in 3D. The dramatic reduction in processing time of PC-VAE proves feasibility of scaling to larger datasets without sacrificing performance.

Regarding the disentanglement of different modes of variation, it turns out that the majority of the latent dimensions is correlated with some mode of variation that is present in the data. This makes isolating distinct factors of variation challenging, as the latent space does not naturally separate them into independent components. In Figure 3.9 the correlation between each latent space and 3 different modes of variation (radius, height, number of localizations) is plotted for the 3D NPC dataset. Here we used the number of localizations as a rough estimate of the DOL. The first plot indicates that, while many of the latent dimensions are correlated with the change in radius, some exhibit an increasing trend, while others show a decreasing behavior. In particular, in latent dimension number 4, the radius variation is clearly observed. Similarly, the second plot shows increasing and decreasing dependence of height on the latent dimension coordinate for different latent dimensions, and latent dimension 8 showing the most clear trend in the height for each bin. The third plot shows a clear trend with the number of localizations especially for latent dimension 8. The impact of a low and high degree of labeling can also be observed in figure 3.7. Approaches such as the total correlation variational auto encoder (TC-VAE) [39], which aims to reduce dependencies between latent variables, could be explored to improve the disentanglement of different modes of variation, allowing for more interpretable and controllable latent representations.

The designed PC-VAE can also be adapted to incorporate localization uncertainty, by modifying the Chamfer cost function to account for uncertainties, and so designing a new model that integrates these variations. However, the design and potential convergence of such a model would also need to be carefully considered, as the introduction of uncertainties could affect the stability and efficiency of training. By accounting for potential



**Figure 3.9: Detected heterogeneity parameters along all latent dimensions for the 3D NPC dataset.** This figure illustrates the correlation between each latent dimension and three modes of variations: (a) radius, (b) height, and (c) number of localizations. The legend applies to all panels.

variations in the data, this approach could offer a more reliable representation by balancing the competing demands of reconstruction accuracy and latent space heterogeneity. Finally, we foresee that PC-VAE can be adapted for a variety of other applications within SMLM, such as image registration and particle tracking, as it is a machine learning approach specifically tailored for SMLM data.

## REFERENCES

- [1] L. Schermelleh, A. Ferrand, T. Huser, C. Eggeling, M. Sauer, O. Biehlmaier, and G. Drummen, *Super-resolution microscopy demystified*, *Nature Cell Biology* **21** (2019), pp. 72–84.
- [2] T. Klein, S. Proppert, and M. Sauer, *Eight years of single-molecule localization microscopy*, *Histochemistry and Cell Biology* **141** (2014), pp. 561–575.
- [3] S. Hell, *Microscopy and its focal switch*, *Nature Methods* **6** (2009), pp. 24–32.
- [4] G. Vicidomini, P. Bianchini, and A. Diaspro, *STED super-resolved microscopy*, *Nature Methods* **15** (2018), pp. 173–182.
- [5] R. P. Nieuwenhuizen, K. A. Lidke, M. Bates, D. Leyton Puig, D. Grünwald, S. Stallinga, and B. Rieger, *Measuring image resolution in optical nanoscopy*, *Nature methods* **10** (2013), pp. 557–562.
- [6] A. Löschberger, S. Van de Linde, M. Dabauvalle, B. Rieger, M. Heilemann, G. Krohne, and M. Sauer, *Super-resolution imaging visualizes the eightfold symmetry of gp210 proteins around the Nuclear Pore Complex and resolves the central channel with nanometer resolution*, *Journal of Cell Science* **125** (2012), pp. 570–575.
- [7] H. Heydarian, F. Schueder, M. Strauss, B. Van Werkhoven, M. Fazel, K. Lidke, R. Jungmann, S. Stallinga, and B. Rieger, *Template-free 2D particle fusion in localization microscopy*, *Nature methods* **15** (2018), pp. 781–784.

- [8] H. Heydarian, M. Joosten, A. Przybylski, F. Schueder, R. Jungmann, B. Van Werkhoven, J. Keller-Findeisen, J. Ries, S. Stallinga, M. Bates, and B. Rieger, *3D particle averaging and detection of macromolecular symmetry in localization microscopy*, *Nature Communications* **12** (2021), p. 2847.
- [9] W. Wang, H. Heydarian, T. Huijben, S. Stallinga, and B. Rieger, *Joint registration of multiple point clouds for fast particle fusion in localization microscopy*, *Bioinformatics* **38** (2022), pp. 3281–3287.
- [10] A. Szymborska, A. De Marco, N. Daigle, V. C. Cordes, J. A. Briggs, and J. Ellenberg, *Nuclear Pore scaffold structure analyzed by super-resolution microscopy and particle averaging*, *Science* **341** (2013), pp. 655–658.
- [11] A. Jimenez, K. Friedl, and C. Leterrier, *About samples, giving examples: optimized single molecule localization microscopy*, *Methods* **174** (2020), pp. 100–114.
- [12] T. Huijben, H. Heydarian, A. Auer, F. Schueder, R. Jungmann, S. Stallinga, and B. Rieger, *Detecting structural heterogeneity in single-molecule localization microscopy data*, *Nature Communications* **12** (2021), p. 3791.
- [13] L. Earl, V. Falconieri, J. Milne, and S. Subramaniam, *Cryo-EM: beyond the microscope*, *Current Opinion in Structural Biology* **46** (2017), pp. 71–78.
- [14] J. Dubochet, M. Adrian, J. Chang, J. Homo, J. Lepault, A. McDowell, and P. Schultz, *Cryo-electron microscopy of vitrified specimens*, *Quarterly reviews of biophysics* **21** (1988), pp. 129–228.
- [15] S. Scheres, H. Gao, M. Valle, G. Herman, P. Eggermont, J. Frank, and J. Carazo, *Disentangling conformational states of macromolecules in 3D-EM through likelihood optimization*, *Nature Methods* **4** (2007), pp. 27–29.
- [16] D.-H. Chen, J.-L. Song, D. Chuang, W. Chiu, and S. Ludtke, *An expanded conformation of single-ring GroEL-GroES complex encapsulates an 86 kDa substrate*, *Structure* **14** (2006), pp. 1711–1722.
- [17] I. Gabashvili, R. Agrawal, R. Grassucci, and J. Frank, *Structure and structural variations of the Escherichia coli 30S ribosomal subunit as revealed by three-dimensional cryo-electron microscopy*, *Journal of Molecular Biology* **286** (1999), pp. 1285–1291.
- [18] F. Baumgart, A. M. Arnold, K. Leskovaar, K. Staszek, M. FÄlser, J. Weghuber, and G. J. SchÄijtz, *Varying label density allows artifact-free analysis of membrane-protein nanoclusters*, *Nature Methods* **13** (2016), pp. 661–664.
- [19] R. P. Nieuwenhuizen, M. Bates, A. Szymborska, K. A. Lidke, B. Rieger, and S. Stallinga, *Quantitative localization microscopy: effects of photophysics and labeling stoichiometry*, *PLOS ONE* **10** (2015), p. e0127989.
- [20] M. Chen and S. J. Ludtke, *Deep learning-based mixed-dimensional Gaussian mixture model for characterizing variability in Cryo-EM*, *Nature Methods* **18** (2021), pp. 930–936.

- [21] S. Haghparast, S. Stallinga, and B. Rieger, *Detecting continuous structural heterogeneity in single-molecule localization microscopy data*, *Scientific Reports* **13** (2023), p. 19800.
- [22] J. Danial and A. Garcia-Saez, *Quantitative analysis of super-resolved structures using ASAP*, *Nature Methods* **16** (2019), pp. 711–714.
- [23] S. Hugelier, H. Kim, M. Gyparaki, C. Bond, Q. Tang, A. Santiago-Ruiz, S. Porta, and M. Lakadamyali, *Eclipse: a versatile classification technique for structural and morphological analysis of super-resolution microscopy data*, *BioRxiv* (2023), pp. 2023–05.
- [24] Y. Wu, P. Hoess, A. Tschanz, U. Matti, M. Mund, and J. Ries, *Maximum-likelihood model fitting for quantitative analysis of SMLM data*, *Nature Methods* **20** (2023), pp. 139–148.
- [25] D. P. Kingma, *Auto-encoding variational bayes*, arXiv preprint arXiv:1312.6114 (2013).
- [26] M. J. Rust, M. Bates, and X. Zhuang, *Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM)*, *Nature Methods* **3** (2006), pp. 793–796.
- [27] J. Schnitzbauer, M. T. Strauss, T. Schlichthaerle, F. Schueder, and R. Jungmann, *Super-resolution microscopy with DNA-PAINT*, *Nature Protocols* **12** (2017), pp. 1198–1228.
- [28] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, *Pointnet: Deep learning on point sets for 3D classification and segmentation*, in *Proceedings of the IEEE conference on computer vision and pattern recognition* (2017) pp. 652–660.
- [29] K. He, X. Zhang, S. Ren, and J. Sun, *Deep residual learning for image recognition*, in *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016) pp. 770–778.
- [30] A. Vaswani, *Attention is all you need*, *Advances in Neural Information Processing Systems* (2017).
- [31] D. P. Huttenlocher, G. A. Klanderman, and W. J. Rucklidge, *Comparing images using the Hausdorff distance*, *IEEE Transactions on pattern analysis and machine intelligence* **15** (1993), pp. 850–863.
- [32] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, and A. Desmaison, *Pytorch: An imperative style, high-performance deep learning library*, *Advances in Neural Information Processing Systems* **32** (2019).
- [33] D. P. Kingma, *Adam: A method for stochastic optimization*, arXiv preprint arXiv:1412.6980 (2014).
- [34] C. P. Burgess, I. Higgins, A. Pal, L. Matthey, N. Watters, G. Desjardins, and A. Lerchner, *Understanding disentangling in  $\beta$ -VAE*, arXiv preprint arXiv:1804.03599 (2018).

- [35] W. Wang, A. Jakobi, Y. Wu, J. Ries, S. Stallinga, and B. Rieger, *Particle fusion of super-resolution data reveals the unit structure of Nup96 in Nuclear Pore Complex*, *Scientific Reports* **13** (2023), p. 13327.
- [36] H. Heydarian, F. Schueder, M. T. Strauss, B. Van Werkhoven, M. Fazel, K. A. Lidke, R. Jungmann, S. Stallinga, and B. Rieger, *Template-free 2D particle fusion in localization microscopy*, *Nature Methods* **15** (2018), pp. 781–784.
- [37] S. Petrovic, G. Mobbs, C. Bley, S. Nie, A. Patke, and A. Hoelz, *Structure and function of the Nuclear Pore Complex*, *Cold Spring Harbor Perspectives in Biology* **14** (2022), p. a041264.
- [38] S. Mosalaganti, A. Obarska-Kosinska, M. Siggel, R. Taniguchi, B. TuroÅŁovÃ¡, C. Zimmerli, K. Buczak, F. Schmidt, E. Margiotta, M. Mackmull, and W. Hagen, *AI-based structure prediction empowers integrative structural analysis of human Nuclear Pores*, *Science* **376** (2022), p. eabm9506.
- [39] R. Chen, X. Li, R. Grosse, and D. Duvenaud, *Isolating sources of disentanglement in Variational Autoencoders*, *Advances in Neural Information Processing Systems* **31** (2018).

## DATA AND CODE AVAILABILITY

Data and code are publicly available. Updated versions of the software can be downloaded from <https://github.com/Sobhanhaghpast/PCVAE.git>. The single-molecule localization data is accessible via the 4TU Research Data repository at [https://data.4tu.nl/private\\_datasets/dbgcLOAnBsWzZUJjoVbhcXgirKJmP9ycA3FlmfFcJj0](https://data.4tu.nl/private_datasets/dbgcLOAnBsWzZUJjoVbhcXgirKJmP9ycA3FlmfFcJj0).

## ACKNOWLEDGMENTS

This work has been supported by the Dutch Research Council (NWO), VICI grant no. 17046.

## AUTHOR CONTRIBUTIONS STATEMENT

S.H and Y.Z developed the method, wrote the code, performed the simulations, and analyzed data. B.R, S.S and Q.T. directed the research. The paper was written by all authors.

# 4

## CONTINUOUS STRUCTURAL HETEROGENEITY DETECTION IN CRYO-ELECTRON MICROSCOPY

1

---

<sup>1</sup>**Sobhan haghparast, Maarten joosten**, Arjen J. Jakobi, Bernd Rieger, Sjoerd Stallinga, *Continuous structural heterogeneity detection in Cryo-electron microscopy.*

## 4.1. INTRODUCTION

Cryogenic-sample Electron Microscopy (cryo-EM) has become a fast-growing method for the determination of the three-dimensional (3D) structure of biological macromolecules at resolutions that make atomic interpretation possible, i.e. 3 Ångströms or better [1, 2]. The sample may consist of purified flash-frozen molecules in a thin layer of amorphous ice in random orientations. The 3D reconstruction of a particle's electrostatic potential depends on the averaging of many hundreds to hundreds of thousands of individual 2D projection images of copies of the macromolecule of interest. This method is called Single-Particle Averaging (SPA). If the relative orientations of all particles are known, the 3D electrostatic potential map of the structure can be assembled from the 2D projection images according to the Fourier slice theorem. Typically, these relative orientations are not known in advance and need to be estimated. Most popular software packages for SPA treat this inverse problem as Bayesian and employ maximum-likelihood algorithms to estimate the relative orientations of the particles and the electrostatic potential[3–6].

Averaging over particles inherently means averaging over different conformational or compositional states of the molecule of interest. Many proteins are conformationally heterogeneous, with conformational changes often being related to the protein's function. A single 3D volume only depicts the ensemble average of all conformational states of particles in the data, and is called the *consensus map*. Not much is known about the effect of rapid cooling on the conformational landscape accessible to proteins, but recent work suggests that the cooling rate in Cryo-EM is fast enough to preserve most of the native conformation landscape[7].

Maximum-likelihood approaches to 3D reconstruction are able to reconstruct different conformational states of a protein by also estimating a 3D class assignment along with relative particle orientations[8]. This makes the algorithm computationally more costly and limited to a small number of discrete 3D classes. Conformational heterogeneity can reduce the resolution of 3D electrostatic potential as it makes perfect alignment between particles using only affine transformations impossible. Particularly in flexible regions of the protein, this causes blurring of the reconstructed volume. For this reason the dataset is often filtered to remove as much of the heterogeneity as possible, potentially discarding a large amount of useful information.

A more recent development is formed by a series of methods that instead map conformational heterogeneity onto a continuous manifold[9], often using machine learning[10]. In Frank and Ourmazd's manifold embedding method[9], aligned projection particles in a single projection direction are mapped to a low-dimensional manifold using Principal Component Analysis (PCA) of the particle images. Zhong et al. proposed the use of a neural network[10] to accomplish a similar mapping to a low-dimensional latent space that learns to encode for the conformational state of each particle. This latent space can then be sampled and used to generate the 3D reconstruction of the protein from a secondary decoder network. Almost all variations of this idea follow the framework of a variational auto-encoder. Examples of programs implementing a manifold learning method include CryoDRGN[11], 3DFlex[12] and DynaMight[13].

Similarly to Frank and Ourmazd we propose to create an embedding of 2D projection particles onto a low-dimensional manifold. Our approach compares projection images in a single projection direction (or in practice a small cone around a single projection direc-

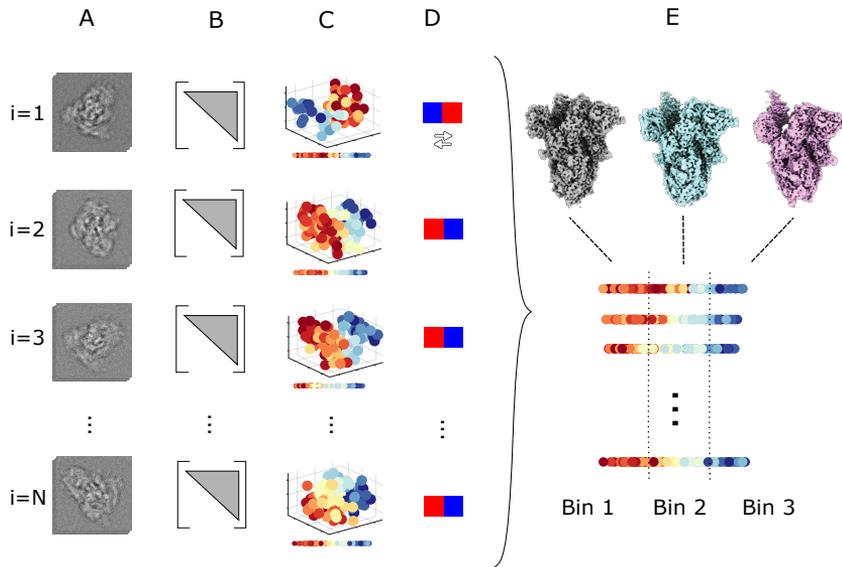
tion) by means of a simple cost function. This results in  $N(N - 1)/2$  similarity measurements per projection direction, which can then be embedded in a low-dimensional manifold using Multi-Dimensional Scaling (MDS) as was previously done for Single-Molecule Localization Microscopy (SMLM) data [14]. An important prerequisite for this approach is that the largest source of dissimilarity between particles stems from conformational or compositional differences rather than differences in orientation, defocus or other non-structural factors. Subsequently, these MDS embeddings are integrated across all orientations to identify the continuous states of the imaged structure in 3D. Using this framework, we can also generate 3D reconstructions that capture the continuous heterogeneity of structural states.

Particle fusion under the homogeneity assumption is a common practice in SMLM and different approaches have been developed for this purpose[15–18]. Methods for continuous structural heterogeneity detection, however, are but recently developed[14, 19]. The currently proposed method has been applied in SMLM to detect continuous structural heterogeneity in localization data, but it has not yet been applied to Cryo-EM data to the best of our knowledge. A fundamental difference between SMLM and Cryo-EM data is that in SMLM comparisons between particles are made based on the localisation point clouds, whereas for Cryo-EM data this needs to be done on pixelised images. As a result, (dis)similarity measures such as the Bhattacharya distance[20] cannot be used directly. Furthermore, subtle and complicating factors for the approach in SMLM such as low degree of labeling and limited photon count of emitters do not appear in Cryo-EM, where rather factors like defocus, unknown viewing direction and radiation damage complicate the analysis. For these reasons direct application of SMLM methods, even for a single orientation class, is not possible.

In the remainder of this chapter, we will first discuss the proposed continuous heterogeneity detection algorithm and then apply the algorithm to two synthetic datasets of the SARS-CoV-2 spike glycoprotein. As a first dataset we study a toy problem with much better signal-to-noise ratio than one should typically expect for experimental images and with no variation in defocus. The second dataset will be simulated with more realistic imaging parameters, specifically incorporating radiation damage.

## 4.2. METHODS

The proposed pipeline presented in Figure 4.1 begins the segmentation of different orientation classes and the alignment of particles in each orientation class. Each orientation class is analyzed separately, where pairwise image comparisons are used to build dissimilarity matrices. These matrices are embedded into a high-dimensional space using Multi-Dimensional Scaling (MDS), and the first principal component of each embedding is extracted to represent a 1D latent space that is hypothesized to capture the continuous heterogeneity. Since the polarity of the latent space dimension is in principle arbitrarily defined an automated polarity matching algorithm is applied. Next, each latent space is binned into intervals, and 3D reconstructions are performed per bin to visualize structural transitions across the dataset.



*Figure 4.1: Schematic of the proposed pipeline. First, particles are aligned and separated into subsets based on their relative orientation estimated during consensus reconstruction. **A)** Shows example orientation subset averages. **B)** Shows the all-to-all pairwise comparison matrix, which is used for MDS embedding. **C)** Shows scatter plots of the first 3 dimensions of MDS space, where particles are color-coded from blue to red to represent progression along the ground-truth motion. **D)** Shows the final combined 1D latent space's polarity which shows if the heterogeneity is from high to low or vice-versa. **E)** Depicts the binned polarity matched latent space along with 3D reconstructions per bin.*

### 4.2.1. IMAGE SIMULATION

Synthetic datasets with inherent conformational heterogeneity are generated using the Cryo-EM simulation software Roodmus[21]. Conformational heterogeneity is created by performing a steered molecular dynamics simulation in openMM version 8.0.0 [22] that interpolates between the closed (pdb id 6xm5) and partially open (pdb id 6xm4) states of the SARS-CoV-2 glycoprotein. Harmonic restraints centred at the target state are used to drive the conformational change with a force constant of 1 kJ/mol/nm<sup>2</sup>.

For the first dataset the steered MD simulation ran for 300 ps with a time step of  $dt = 2$  fs and a frame is saved every 1 ps. The trajectory is downsampled to 20 equispaced conformations. A total of 500 Cryo-EM micrographs with 300 particles each are simulated from these 20 conformations using Roodmus. The imaging parameters used are as follows: 1 Å pixel size, 45 e/Å<sup>2</sup> fluence, 300 kV acceleration potential, -0.5 μm defocus, and no radiation damage. The same defocus value is used for all micrographs.

For the second dataset the steered MD simulation ran for 3000 ps to create a smoother trajectory with a time step of  $dt = 2$  fs and a frame is saved every 5 ps. A total of 900 Cryo-EM micrographs with 300 particles each are simulated from the MD trajectory with 600 conformations. The imaging parameters are 1 Å pixel size, 45 e/Å<sup>2</sup> fluence, 300 kV acceleration potential, defocus -2.0 μm with a standard deviation of 0.5 μm. Radiation damage is also enabled with a radiation sensitivity coefficient of 0.022 Å<sup>2</sup>/e and as a result micrographs are generated as a 40-frame movie, which about matches the experimental conditions for the frame rate of the camera.

### 4.2.2. IMAGE PROCESSING IN RELION

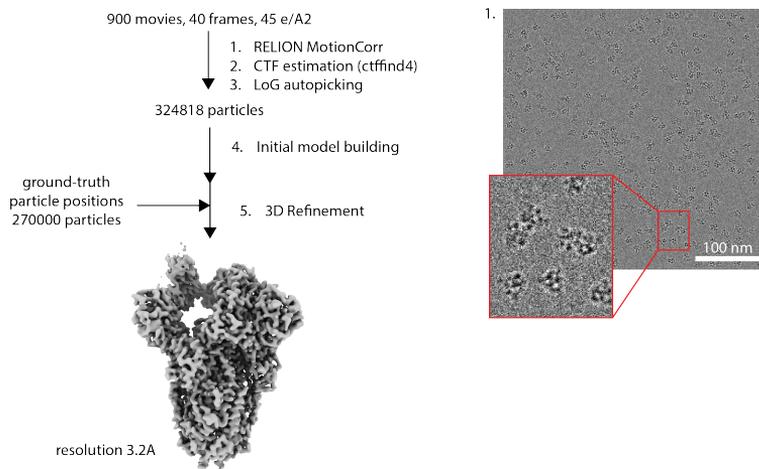
All image processing was done in RELION version 5.0-beta-3-commit-12cf15 [23]. Micrographs are imported, or multi-frame movies are imported and motion-corrected using RELION's implementation of the MotionCorr algorithm. CTF estimation is done using ctfind 4.1.14 [24]. The full processing workflow for the dataset with radiation damage is described in Figure 4.2.

### 4.2.3. ASSIGNING ORIENTATION CLASS

Particles are first aligned and CTF-corrected using the 'relion\_ctf\_toolbox' utility in RELION. Relative particle orientations in RELION are stored as Euler angles in the ZYZ intrinsic axis convention with angles labeled rlnAngleRot, rlnAngleTilt and rlnAnglePsi. Particles with similar Rot and Tilt angles are projected along the same axis and only differ by an in-plane rotation around the projection axis by the Psi angle. As a result, such projection images can be compared directly. Tilt and Rotation angles are separated into 36 evenly spaced bins and particles are assigned to a single bin based on their relative orientation.

### 4.2.4. PAIRWISE-DISTANCE MATRIX

The continuous structural heterogeneity of each 2D orientation class is independently detected using pairwise comparisons between all pairs of particles. To compare the 2D images, we started with cropping. Particles are first aligned and centered using RELION software. By averaging and visualization of the image stacks per orientation, a cropping box is manually selected to ensure that, for all orientation classes, only the background is removed. This operation reduces the pairwise comparison area from 256Å×256 to 175Å×175



*Figure 4.2: Overview of image processing in RELION. Synthetic data created with Roodmus can be processed in the same manner as experimental data. Motion correction and CTF estimation were performed, followed by particle picking and immediate Initial Model building. To reduce noise in the dataset coming from potential bad particle picks, ground-truth positions of particles were used for 3D refinement.*

pixels while the signal is preserved. For the comparison of two images, we selected the Mean Square Error (MSE) as the dissimilarity measure. We evaluated several dissimilarity measures, including peak signal-to-noise ratio (PSNR) and structural similarity index measure (SSIM). When analyzing the normalized histograms of pairwise similarity values, MSE showed a wider distribution, indicating better sensitivity to variations. This characteristic along with low computational cost made MSE the most suitable measure for our analysis. Since MSE is sensitive to pixel intensities, the presence of misalignment commonly occurring during image registration can cause MSE to pick up these misalignments and can potentially dominate the continuous heterogeneity. To address this, we minimize the MSE between each image pair by applying a number (125) of affine transformations to one of the images. The typical error made by registration software is around 2 degrees and 2 pixel shifts (pixel size 1 Å). Therefore, we opted to shift images in both the  $x$  and  $y$  directions by -2, -1, 0, 1, 2 pixels and rotate images by -2, -1, 0, 1, 2 degrees for a total of 125 transformations. In principle, more transformations could be considered, but we expect that a further increase in number of shifts and rotations would not significantly improve the result and would rather make the algorithm too computationally expensive. This process results in an  $N \times N$  upper triangular matrix containing relative information about the  $N$  images within a single orientation class.

#### 4.2.5. MDS EMBEDDING

The found dissimilarity values are then embedded into an abstract high-dimensional Cartesian space using MDS [25]. For this purpose, we select the number of dimensions as a

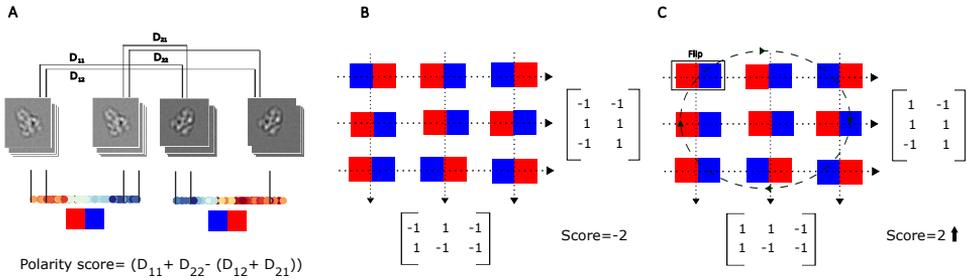
rule of thumb to 30, a choice also adopted in prior studies [18, 26]. Each particle in the orientation class is assigned one specific point in the MDS space, which is placed based on its dissimilarity with respect to all other particles. In the presence of continuous heterogeneity, these points of particles are ordered accordingly, and the (high-dimensional) manifold formed by these points carries information about the continuous structural heterogeneity. In our simulation study, where the dominant mode of variation is continuous heterogeneity, applying PCA to the MDS space reveals that the variance explained by the first principal component (PC) axis is the highest. It is expected that the first PC axis is sensitive to continuous structural heterogeneity. Next, we project the high-dimensional manifold onto the first PC axis, resulting in a 1D latent space for each orientation class, with expected ordering based on the continuous structural heterogeneity. This procedure is repeated for all orientation classes, resulting in a 1D latent space for each class.

#### 4.2.6. POLARITY MATCHING

Since the MDS space is abstract and shaped directly from the dissimilarity matrix information, the 1D latent space for different orientation classes has in principle arbitrary polarities. It is therefore necessary to match the polarities and obtain an ordered heterogeneity along the continuous motion parameter. For a 36-by-36 grid of orientation classes, manual polarity matching would not be an option. Therefore, we employed an automatic approach consisting of four steps: First, we order all particles in each orientation class based on the 1D latent space. Second, for each orientation class, we create two particle stacks that fall into the first and last quantiles of the respective latent space, resulting in two stacks for each orientation class. Third, we form two score matrices comparing neighboring classes in the horizontal and vertical directions of the orientation class grid. Scores are calculated by pairwise comparison of the images in two particle stacks. This is performed by computing the average MSE values of the two particle stacks between adjacent neighbors as explained in Figure 4.3. This process determines whether two orientation classes have matching orientations. Furthermore, we take the average MSE value as a confidence criterion, indicating how reliable the assignment of relative polarity for adjacent orientation classes is. Finally, we maximize the number of matched polarities by iteratively flipping the polarity for the different orientation classes in the grid. In each iteration starting from the top left to the bottom right, the latent space of an orientation class is flipped, and the total matching score is obtained by summing all pairwise polarity matching scores both vertically and horizontally. If the score increases, the flip is retained. If the score decreases, the flip is not accepted and the latent space is returned to its original polarity. Through successive iterations, the matching score improves, eventually aligning the latent space polarities.

#### 4.2.7. LATENT SPACE BINNING AND RECONSTRUCTION

Each polarity-matched latent space is divided into three bins, where the 2D images in each bin are considered to be homogeneous and represent a specific state of the protein. This results in three sets of 36-by-36 orientation class grids. Next, for each grid of orientation classes, we perform a 3D reconstruction using the RELION[23] software. Reconstruction results for 10 bins have also been made, and appear in the appendix.



**Figure 4.3: Comparing particle stack subsets between neighboring orientation classes for polarity matching.** A) Each latent space is arranged based on the conformation of the orientation class. However, since the MDS space is an abstract Cartesian space formed iteratively, the polarity of the latent space is in principle arbitrary. To assess polarity matching, the average mean squared error is computed for all pairs of the four subsets. A positive value indicates matching polarity, while a negative value suggests a mismatch. The absolute value serves as a confidence measure. B) The top-left panel shows the initial calculation of the binary matching metric for neighboring classes in the vertical and horizontal direction of the orientation class grid. C) The latent space is flipped, and the matching score is calculated. If the score improves, the flip is retained; otherwise, it is reverted. Iterations continue to align the polarities.

### 4.3. RESULTS

To demonstrate the feasibility of the continuous heterogeneity detection method we first applied it to a high-quality synthetic dataset of the SARS-CoV-2 spike glycoprotein. In this dataset we limited the conformational heterogeneity to 20 distinct states sampled from a 200 ps steered molecular dynamics trajectory. We simulated 500 micrographs containing a total of 150,000 particles. Further details on the simulation are presented in the methods section.

The first step in the algorithm, as described earlier, is to assign all particles to orientation classes. In Figure 4.5A we show the consensus 3D reconstruction of all 150,000 particles with relative orientations estimated by RELION. The global resolution of the map is 2.42 Å, although the measurement of resolution can be compromised by the CTF oscillations if all particles have the same defocus. We also determined the B-factor to be 29.8 Å<sup>2</sup> by means of a ResLog plot, which is comparable to previous simulated datasets with conformational heterogeneity [21] Figure 4.5B shows the orientation distribution of the particles as a histogram on a sphere as well as projected to a 2D grid. Seemingly the estimated orientation distribution does not match the ground-truth uniform distribution of orientations. In order to test the difference between orientation distributions we transformed both the ground-truth and estimated Euler angles to rotation matrices and computed the Wasserstein distance on an element wise basis between their distributions. This distance turned out to be 0.003, from which it can be concluded that the distributions are in fact equivalent. We conclude that the peak in the estimated orientations is caused by the fact

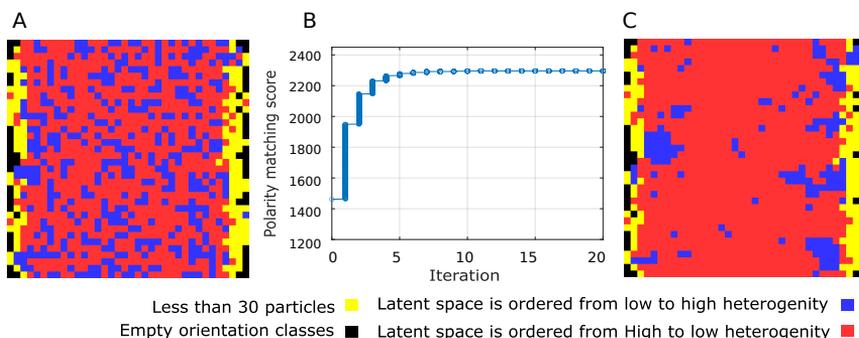
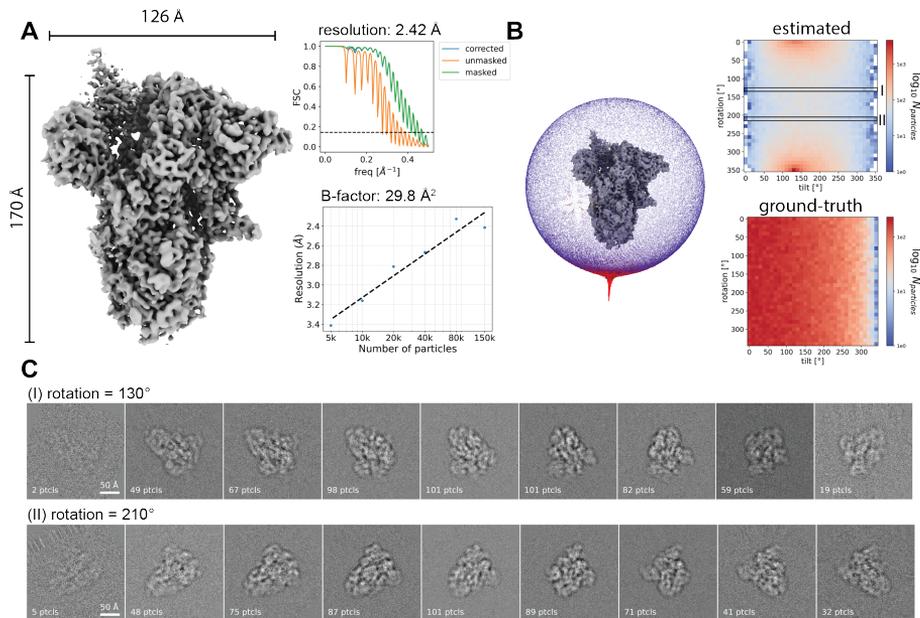


Figure 4.4: Polarity matching result. **A)** Initial polarity of the 1D latent space for first simulated dataset based on PC1. **B)** Iterative update of polarity matching score. **C)** Final polarity matching result. **D)** Initial polarity of the 1D latent space for second simulated dataset based on PC4. **E)** Iterative update of polarity matching score of the second simulated dataset. **F)** Final polarity matching result of second simulated dataset.

that there is no one-to-one mapping of orientations in  $SO(3)$  to Euler angles. In fact the visualization of Figure 4.5B is not appropriate. The non-uniqueness of the mapping between the Euler-angles and orientation in  $SO(3)$  is the phenomenon of *gimbal lock*, where the second Euler angle is close to 0, and the first and third rotations are around the same axis, and cannot be well separated. During simulation, a uniform distribution over  $SO(3)$  is created by uniformly sampling a point on a sphere to act as rotation vector, and then sampling a rotation angle uniformly between 0 and  $2\pi$ . We subsequently split the particles into 36 orientation classes based on their Euler angles. Example orientation class averages of bins highlighted in Figure 4.5B are shown in Figure 4.5C.

A few examples of the first three dimensions of this MDS space for different orientation classes are shown in Figure 4.6. Each point in the scatter plots represents a single projection image, coloured to reflect the ground-truth conformational state of the protein, indexed between 0 and 19. Below each scatter plot the one-dimensional projection of the MDS space on the first Principal Component (PC) is shown. In nearly all example orientations an organisation of the particles can be observed that corresponds to the conformational change of the particle in chronological order. It can also be seen that the polarity of these 1D projections is arbitrary. The next step in the algorithm is therefore to align the polarity of each projected MDS space in order to create a unified representation of the heterogeneity in the dataset from all orientation classes. In Figure 4.4(A-C) the polarity matching result is shown. For the first simulated dataset polarity matching resulted in 93% of all particles to be in a latent space with matched polarity.

After polarity alignment, the 1-dimensional MDS representation can be used for clustering, or for splitting the particles in the dataset across discrete conformational classes. Figure 4.7A shows splitting the total number of particles into three conformational subsets based on their latent representation. The total number of particles in each subset was kept the same. Density maps reconstructed from each subset show the different conformational states of the protein. Figure 4.7B shows the distribution of the aligned MDS



**Figure 4.5:** **A)** Consensus reconstruction of the conformationally heterogeneous dataset. FSC[27] plot shows a global resolution of 2.42 Å. From the ResLog plot we calculated a B-factor of 29.8 Å<sup>2</sup>. **B)** Orientation distribution shown as a histogram on a sphere with larger peaks corresponding to more particles in that orientation. The same distribution is also shown in a projected heatmap. Each pixel in the heatmap corresponds to different orientation class. **C)** Example particle averages taken along a single row of orientation classes.

embedding versus the ground-truth index of the conformational state of the molecule. It shows that splitting the MDS space into three subsets also partitions the ground-truth conformations, with especially class 1 and 3 corresponding to mostly non-overlapping conformations, while class 2 includes more particles from the entire range of conformational states. In Figure 4.7C we show 2D averages of particles in the same orientation class after splitting into the three conformational subsets. In directions where the Receptor Binding Domain (RBD) is clearly visible the different positions of this domain can be seen even in 2D. Figure 4.7D shows a correlation heatmap comparing the 3D reconstructions of the three conformational subsets to the 20 ground-truth conformational states of the particle. Class 1 correlates more strongly with the closed state of the protein, class 3 with the open state and class 2 lies in between.

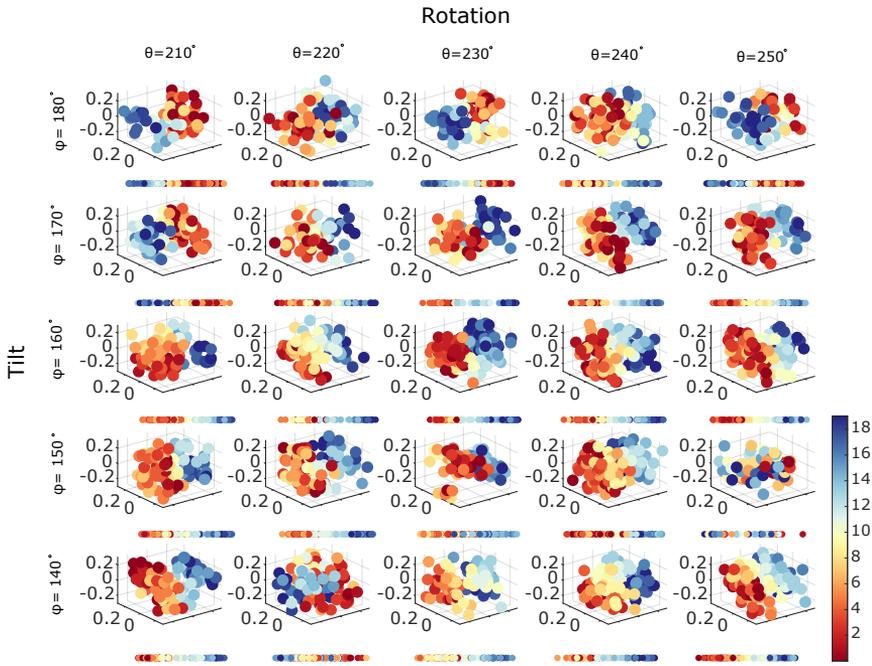
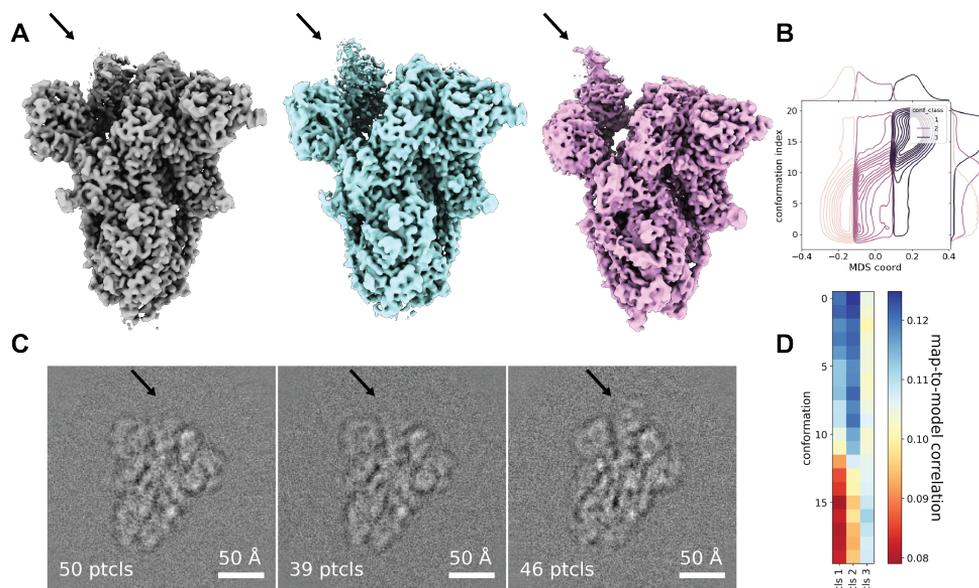
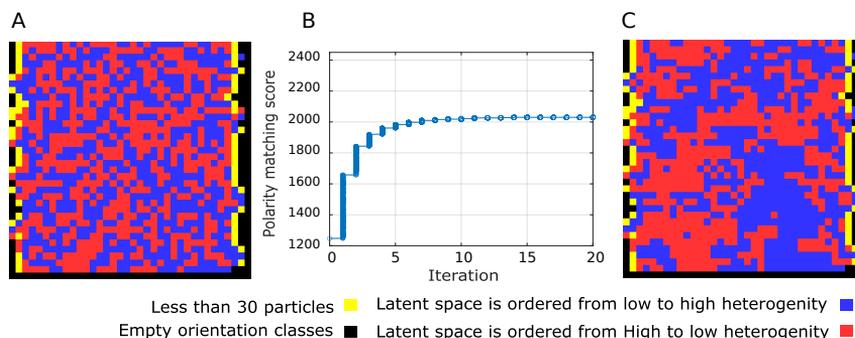


Figure 4.6: **Multi-dimensional scaling space for each orientation class for the first dataset.** This plot illustrates the MDS space for orientation classes from 21 to 25 for rotation and 14 to 18 for tilt orientation classes, where each tilt or rotation orientation class spans 10 degrees. The 1D latent spaces for all classes are formed by projecting the MDS data onto the first Principal Component (PC) axis and are plotted below each MDS space.



**Figure 4.7:** **A** Resulting density maps from each conformation class of particles obtained by binning the polarity matched MDS spaces. Each class results in a distinct conformation of the protein, mainly distinguished by the position of the RBD. **B** Distribution of the polarity matched MDS space versus the ground-truth conformation index. Different conformation classes based on MDS coordinates are shown in different colours and show that the organisation follows the ground-truth conformations. **C** Example of averaged particles from one orientation class, separated into conformation classes based on the MDS embedding. The arrow points to the RBD and shows the difference in conformations is detectable in 2D projection. **D** Heatmap of the real-space map-to-model correlation between each of the three conformation classes and the 20 atomic models that represent the ground-truth conformations of the protein.

While the method worked well on the initial simulated dataset, it did not perform well on the lower quality dataset in which radiation damage and variation in defocus was taken into account which is shown in Figure 4.10. We found that in this case the first PC axis of the MDS space was no longer informative of the conformational difference between particles. Instead, the fourth PC seemed to reflect these conformational differences better. This indicates that the method of pairwise comparison we opted for may be too sensitive to other forms of heterogeneity between particles, such as lower signal to noise ratio resulting from defocus or misalignment of particles in each orientation class. In Figure 4.9 the dimensions 3 to 5 of the MDS space are plotted along with the fourth PC axis as latent space. This projection of MDS space has a more spherical structure compared to the previous dataset. As a result, the dominant mode of variation represented by MDS space is not the sought-after conformational heterogeneity. The detection of the PC axis that carries the most information about the structural heterogeneity can therefore not be simply



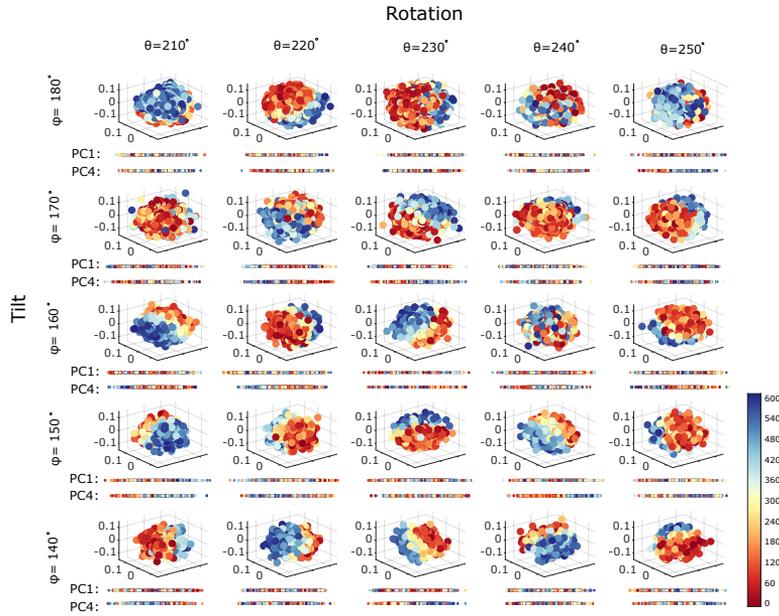
*Figure 4.8: Polarity matching result for second dataset. A) Initial polarity of the 1D latent space for second simulated dataset based on PC4. B) Iterative update of polarity matching score of the second simulated dataset. C) Final polarity matching result of second simulated dataset.*

selected as PC1 and requires further visualizations and analysis. As shown in figure 4.9 the 1-dimensional latent space according to PC4 has a better separation of the red/blue color indicative of the ground truth heterogeneity parameter than the one according to PC1. We also noticed that the lower signal-to-noise ratio of the images affects the polarity matching procedure too. This results in lower decision values for comparing the polarity of the latent spaces of adjacent orientation classes as shown in Figure 4.8(A-C) resulting in 53% polarity matching score for all particles.

#### 4.4. DISCUSSION & OUTLOOK

While we found that our proposed method worked well on the first synthetic dataset with simplified imaging parameters, introducing real-world complexities like variable defocus, radiation damage (which reduces the strength of high-resolution information in the image and hampers alignment), and an increased number of conformational states to resolve, revealed where the method begins to struggle. In the second synthetic dataset where we included all of the above mentioned effects, we could not see that the first Principal Component of the MDS space was representative of the conformational heterogeneity in the images. One of the possible reasons for the method's incapability to detect continuous structural heterogeneity in this second simulated dataset is that taking the MSE as the pairwise comparison metric between 2D images comes with an undesired pixel level sensitivity.

In the presence of two sources of blurring in the second dataset (namely the effects of radiation damage and errors in CTF-correction), MSE may not be able to capture the structural differences between particles effectively. This results in an inaccurate formation of MDS space and further analysis is affected as a consequence of this effect. The proposed method does not depend on the use of MSE as a cost function, and may give better results with alternatives that can be fitted into the pipeline without modification. There are several (dis)similarity measures that could potentially address this issue such as



**Figure 4.9: Multi-dimensional scaling space for each orientation class in the second simulated dataset.** This plot illustrates dimensions 3,4 and 5 of MDS space for orientation classes from 21 to 25 for both rotation and 14 to 18 for tilt orientation classes, where each tilt or rotation class spans 10 degrees. The 1-dimensional latent spaces for each class are formed by projecting the MDS data into the first and fourth principal component (PC) axes and are plotted below each MDS space indicating that ordering particles based on PC4 matches better with the structural heterogeneity than PC1 does.

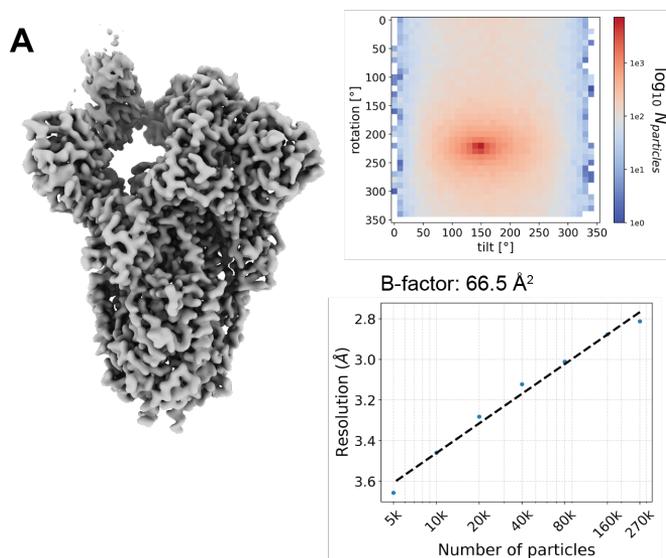


Figure 4.10: **A)** Consensus reconstruction of the second dataset with different defocus and radiation damage. The global resolution of the density map was  $3.2 \text{ \AA}$ . **B)** particle orientation distribution. Like before this shows a single peak due to the degeneracy of Euler angles. **C)** ResLog plot with estimated B-factor. Compared to the previous dataset the B-factor has roughly doubled due to the increase in number of ground-truth conformations and the application of Parakeet's beam damage model.

performing the pairwise comparison in Fourier domain with Fourier Ring Correlation[28], or more complex similarity measures such as modified SSIM, Earthmover/Wasserstein distance[29], etc. Furthermore, it could be possible to use features of the images and perform the pairwise comparison in a suitably defined feature domain. In case of any changes in similarity measure there are several factors which should be taken into account, namely, the computational complexity, and the robustness in low signal-to-noise ratio conditions.

Even if the cost function is sufficiently sensitive to small changes in the conformation of the particles, errors in pre-processing steps can also cause potential vulnerabilities to the analysis. If the segmentation into orientation classes is noisy as a result of misalignment of particles during 3D reconstruction, the quality of the latent MDS space and subsequent reconstructions can drop significantly, suggesting a sensitivity to upstream errors. Similar problems are faced by machine learning approaches, which may be solved by refining the relative orientation of each image iteratively during the pipeline.

To avoid some of the above mentioned problems there are approaches that could improve the method's performance. Currently, the use of a 1D latent space captures only the dominant mode of variation, which may be insufficient for systems with multiple modes of variation, or where the variation is overshadowed by other sources of heterogeneity. Using higher dimensional embeddings from MDS such as 2D or 3D latent spaces could preserve relationships among particles at the cost of making the latent space again less interpretable to a human observer. Furthermore, using more than one latent dimension for reconstruction could potentially help with accurate binning. Additionally, rather than treating each particle as a single unit, breaking particle images into smaller subsections (e.g., 4, 16, or more tiles) and analyzing these patches individually could help localize structural variability by down-weighting image patches that contain less variation. This patch-wise comparison could improve the method's ability to detect small, localized movements that would otherwise be obscured in global image comparisons. It would also be possible to constrain the pairwise comparison between particles with a mask around an area of interest, such as the receptor binding domain that undergoes the conformational change in the example case of the SARS-CoV-2 spike protein. This approach is already common practice in methods such as multibody refinement[30].

However, one of the key advantages of the proposed method is its scalability for large number of particles. Because each orientation class is processed independently, the pipeline is parallelizable, allowing for fast execution with proper implementation. While the disadvantage of this method is that it relies on each orientation class independently. Here, deep learning based methods have the advantage of being able to consider all orientations of the particle at the same time while building a latent space representation. Another downside of the current implementation of the method is that the analysis is centered on MSE which is on the basis of the pixel level and not directly on structural features.

In conclusion, we developed a method for capturing continuous heterogeneity in Cryo-EM SPA data by creating a 1D embedding of particles representing different conformational states. We applied this method to two simulated datasets and demonstrated its effectiveness in reconstructing different conformations of the SARS-CoV-2 spike glycoprotein. Future steps could involve refining the pairwise comparison step for dealing with radiation damage and different defocus of particles.

## REFERENCES

- [1] E. Nogales, *The development of cryo-EM into a mainstream structural biology technique*, [Nature Methods](#) **13** (2016), pp. 24–27, publisher: Nature Publishing Group.
- [2] T. M. de Oliveira, L. van Beek, F. Shilliday, J. E. Debreczeni, and C. Phillips, *Cryo-EM: The Resolution Revolution and Drug Discovery*, [SLAS DISCOVERY: Advancing the Science of Drug Discovery](#) **26** (2021), pp. 17–31, publisher: SAGE Publications Inc STM.
- [3] F. J. Sigworth, P. C. Doerschuk, J.-M. Carazo, and S. H. W. Scheres, *Chapter Ten - An Introduction to Maximum-Likelihood Methods in Cryo-EM*, in [Methods in Enzymology](#), Cryo-EM, Part B: 3-D Reconstruction, Vol. 482, edited by G. J. Jensen (Academic Press, 2010) pp. 263–294.
- [4] A. Punjani, J. L. Rubinstein, D. J. Fleet, and M. A. Brubaker, *cryoSPARC: algorithms for rapid unsupervised cryo-EM structure determination*, [Nature Methods](#) **14** (2017), pp. 290–296, publisher: Nature Publishing Group.
- [5] S. H. Scheres, *Maximum-likelihood methods in cryo-EM. Part II: application to experimental data*, [Methods in enzymology](#) **482** (2010), pp. 295–320.
- [6] S. H. W. Scheres, *RELION: Implementation of a Bayesian approach to cryo-EM structure determination*, [Journal of Structural Biology](#) **180** (2012), pp. 519–530.
- [7] L. V. Bock, M. Igaev, and H. GrubmÄijller, *Single-particle Cryo-EM and molecular dynamics simulations: A perfect match*, [Current Opinion in Structural Biology](#) **86** (2024), p. 102825.
- [8] S. H. W. Scheres, *Chapter Six - Processing of Structurally Heterogeneous Cryo-EM Data in RELION*, in [Methods in Enzymology](#), The Resolution Revolution: Recent Advances In cryoEM, Vol. 579, edited by R. A. Crowther (Academic Press, 2016) pp. 125–157.
- [9] J. Frank and A. Ourmazd, *Continuous changes in structure mapped by manifold embedding of single-particle data in cryo-EM*, [Methods Single Particle Cryo-EM](#), from sample to reconstruction, **100** (2016), pp. 61–67.
- [10] E. D. Zhong, T. Bepler, J. H. Davis, and B. Berger, *Reconstructing continuous distributions of 3D protein structure from cryo-EM images*, (2020), arXiv:1909.05215 [q-bio].
- [11] E. D. Zhong, T. Bepler, B. Berger, and J. H. Davis, *CryoDRGN: reconstruction of heterogeneous cryo-EM structures using neural networks*, [Nature Methods](#) **18** (2021), pp. 176–185, publisher: Nature Publishing Group.
- [12] A. Punjani and D. J. Fleet, *3DFlex: determining structure and motion of flexible proteins from cryo-EM*, [Nature Methods](#) **20** (2023), pp. 860–870, publisher: Nature Publishing Group.

- [13] J. Schwab, D. Kimanius, A. Burt, T. Dendooven, and S. H. W. Scheres, *DynaMight: estimating molecular motions with improved reconstruction from cryo-EM images*, *Nature Methods* **21** (2024), pp. 1855–1862, publisher: Nature Publishing Group.
- [14] S. Haghparast, S. Stallinga, and B. Rieger, *Detecting continuous structural heterogeneity in single-molecule localization microscopy data*, *Scientific Reports* **13** (2023), p. 19800.
- [15] A. Löscherger, S. van de Linde, M.-C. Dabauvalle, B. Rieger, M. Heilemann, G. Krohne, and M. Sauer, *Super-resolution imaging visualizes the eightfold symmetry of gp210 proteins around the nuclear pore complex and resolves the central channel with nanometer resolution*, *Journal of Cell Science* **125** (2012), pp. 570–575.
- [16] H. Heydarian, F. Schueder, M. T. Strauss, B. Van Werkhoven, M. Fazel, K. A. Lidke, R. Jungmann, S. Stallinga, and B. Rieger, *Template-free 2D particle fusion in localization microscopy*, *Nature Methods* **15** (2018), pp. 781–784.
- [17] H. Heydarian, M. Joosten, A. Przybylski, F. Schueder, R. Jungmann, B. v. Werkhoven, J. Keller-Findeisen, J. Ries, S. Stallinga, M. Bates, *et al.*, *3D particle averaging and detection of macromolecular symmetry in localization microscopy*, *Nature Communications* **12** (2021), pp. 1–9.
- [18] W. Wang, H. Heydarian, T. A. Huijben, S. Stallinga, and B. Rieger, *Joint registration of multiple point clouds for fast particle fusion in localization microscopy*, *Bioinformatics* **38** (2022), pp. 3281–3287.
- [19] J. Danial and A. Garcia-Saez, *Quantitative analysis of super-resolved structures using ASAP*, *Nature Methods* **16** (2019), pp. 711–714.
- [20] F. J. Aherne, N. A. Thacker, and P. I. Rockett, *The bhattacharyya metric as an absolute similarity measure for frequency coded data*, *Kybernetika* **34** (1998), pp. 363–368.
- [21] M. Joosten, J. Greer, J. Parkhurst, T. Burnley, and A. J. Jakobi, *Roodmus: a toolkit for benchmarking heterogeneous electron cryo-microscopy reconstructions*, *IUCr* **11** (2024), pp. 951–965.
- [22] P. Eastman, J. Swails, J. D. Chodera, R. T. McGibbon, Y. Zhao, K. A. Beauchamp, L.-P. Wang, A. C. Simmonett, M. P. Harrigan, C. D. Stern, R. P. Wiewiora, B. R. Brooks, and V. S. Pande, *OpenMM 7: Rapid development of high performance algorithms for molecular dynamics*, *PLOS Computational Biology* **13** (2017), p. e1005659.
- [23] J. Zivanov, T. Nakane, B. O. Forsberg, D. Kimanius, W. J. H. Hagen, E. Lindahl, and S. H. W. Scheres, *New tools for automated high-resolution cryo-EM structure determination in RELION-3*, *eLife* **7** (2018), p. e42166, using RELION 5.0 beta, based on this foundational work.
- [24] A. Rohou and N. Grigorieff, *CTFFIND4: Fast and accurate defocus estimation from electron micrographs*, *Journal of Structural Biology* **192** (2015), pp. 216–221.

- [25] A. Mead, *Review of the development of multidimensional scaling methods*, *Journal of the Royal Statistical Society: Series D* **41** (1992), pp. 27–39.
- [26] T. A. Huijben, H. Heydarian, A. Auer, F. Schueder, R. Jungmann, S. Stallinga, and B. Rieger, *Detecting structural heterogeneity in single-molecule localization microscopy data*, *Nature Communications* **12** (2021), pp. 1–8.
- [27] M. Van Heel and M. Schatz, *Fourier shell correlation threshold criteria*, *Journal of Structural Biology* **151** (2005), pp. 250–262.
- [28] R. P. Nieuwenhuizen, K. A. Lidke, M. Bates, D. L. Puig, D. Grünwald, S. Stallinga, and B. Rieger, *Measuring image resolution in optical nanoscopy*, *Nature methods* **10** (2013), pp. 557–562.
- [29] Y. Rubner, C. Tomasi, and L. J. Guibas, *The earth mover’s distance as a metric for image retrieval*, *International Journal of Computer Vision* **40** (2000), pp. 99–121.
- [30] T. Nakane, D. Kimanius, E. Lindahl, and S. H. W. Scheres, *Characterisation of molecular motions in cryo-em single-particle data by multi-body refinement in relion*, *eLife* **7** (2018), p. e36861.

## DATA AND CODE AVAILABILITY

Data and code are publicly available at <https://github.com/Sobhanhaghparast/CHD-Cryo-EM>.

## ACKNOWLEDGEMENTS

This work has been supported by the Dutch Research Council (NWO), VICI grant no. 17046.



# 5

## CONCLUSION

## 5.1. CONCLUSION

### 5.1.1. DETECTING CONTINUOUS STRUCTURAL HETEROGENEITY IN SINGLE-MOLECULE LOCALIZATION MICROSCOPY DATA

In Chapter 2, we introduced a template-free continuous structural heterogeneity detection method for single-molecule localization microscopy (SMLM) data. Traditional particle fusion in SMLM relies on the assumption of structural homogeneity, which ignores the potentially useful information on variations within the particle dataset, and which may lead to unreliable particle averaging results. While recent efforts have focused on classifying structural heterogeneity into discrete groups [1], these approaches fail to capture continuous modes of variation, such as expansion, contraction, or other gradual conformational changes. Our proposed method overcomes these limitations by minimizing human intervention, eliminating the need for high-quality data, avoiding template bias and working directly on 2D/3D localization datasets compared to existing continuous heterogeneity detection methods [2, 3].

Our approach is based on processing point clouds, beginning with the parallel computation of a dissimilarity measure for all particle pairs in the dataset. The dissimilarity matrix is then embedded into a high-dimensional abstract Cartesian space using multidimensional scaling (MDS) [4]. Subsequently, dimensionality reduction is applied, yielding a 1D latent space that orders all particles according to a continuous mode of heterogeneity. From this latent space, subsets of structurally similar particles are identified and grouped for reconstruction, producing multiple reconstructions that enhance signal-to-noise ratio (SNR) while preserving the homogeneity assumption within each group.

We successfully detected continuous structural variations in various localization microscopy datasets, such as height variations in 3D DNA-origami tetrahedrons and radius variations in 2D Nuclear Pore Complex (NPCs) datasets. These refinements resulted in more faithful particle fusions. Our findings demonstrate that even with fewer particles, more homogeneous particle stacks can achieve high-resolution reconstructions while accounting for structural heterogeneity. introduced method scales quadratically with the number of particles, making it impractical for large datasets. Additionally, its sensitivity is limited by the choice of dissimilarity measure namely Bhattacharya cost function [5], which fails to capture complex structural variations, especially in 3D experimental datasets like the NPC including 16 to 32 binding sites.

### 5.1.2. DETECTING CONTINUOUS STRUCTURAL HETEROGENEITY IN SINGLE-MOLECULE LOCALIZATION MICROSCOPY DATA WITH A POINT CLOUD VARIATIONAL AUTO-ENCODER

To address these challenges, we developed the Point Cloud Variational Auto-Encoder (PC-VAE), a deep learning framework tailored for SMLM data. PC-VAE allows template-free, continuous heterogeneity detection directly from 2D and 3D localization coordinates, avoiding the need for voxelization and reducing memory usage and computational complexity.

PC-VAE encodes SMLM point clouds into a latent space, where particles are ordered along continuous modes of variation. Unlike the previous MDS-based approach, which require computationally expensive pairwise comparisons, scaling quadratically with the number of particles in the dataset, PC-VAE reduces the computational complexity as it

scales linearly with dataset size. This enables analysis of large-scale datasets on the order of a several thousands of particles, with practical computational time. After an unsupervised training, the latent space is used to cluster and reconstruct particles, preserving structural heterogeneity and improving the SNR.

We validated PC-VAE on three experimentally acquired datasets: a 2D Nuclear Pore Complex dataset (1,399 particles) showing radius variation from 63.6 nm to 49.9 nm, a 3D DNA-origami tetrahedron dataset (218 particles) revealing height variation from 42.4 nm to 107.3 nm, and a 3D Nuclear Pore Complex dataset (3,810 particles) with radius variations from 51.9 nm to 56.7 nm together with height variations from 43.7 nm to 51.7 nm. The computational times per epoch were 13 seconds (3D tetrahedron), 24 seconds (2D NPC), and 279 seconds (3D NPC), on a Linux operating system with Intel (R) Xeon (R) Silver 4110 CPU @ 2.10 GHz with 32 GB RAM and NVIDIA RTX A4000 GPU demonstrating significant improvements over previous methods.

In summary, PC-VAE provides a scalable and efficient method for detecting continuous heterogeneity in SMLM data, enhancing both sensitivity and computational performance.

### 5.1.3. CONTINUOUS STRUCTURAL HETEROGENEITY DETECTION IN CRYO-EM DATA

In this study, we presented a method for detecting continuous structural heterogeneity in Cryo-EM data using MDS on orientation-classed particle images. Our approach successfully disentangles orientation from conformational heterogeneity, enabling robust detection of structural states. By analyzing synthetic datasets with inherent conformational heterogeneity, we demonstrated that this method captures subtle conformational transitions and generates 3D reconstructions for different states of a structure.

For high-quality synthetic data, our approach revealed continuous heterogeneity and produces detailed 3D reconstructions. The MDS latent space accurately reflected the underlying structural transition, and polarity matching of the heterogeneity axis for neighbouring orientation classes ensured a consistent representation across orientation classes. This allowed us to visualize and interpret the conformational changes in a manner that provided clear insight into the continuous structural heterogeneity of the sample.

However, the method faced challenges when applied to more realistic synthetic data, particularly due to the effects of radiation damage. In the presence of radiation damage, high-frequency information in the micrographs was lost, making it more difficult to detect continuous heterogeneity. The latent space derived from the MDS did not show as clear a representation of structural variation in these cases, as the loss of high-frequency data reduced the sensitivity of the method to subtle conformational changes. This limitation is inherent in Cryo-EM data when radiation damage is present, which can impact the ability of the method to fully capture structural heterogeneity. The impact of defocus variations with standard Contrast Transfer Function (CTF) correction turned out to have a much smaller impact on the heterogeneity detection.

In conclusion, while the proposed method is effective for analyzing high-quality Cryo-EM data, it faces limitations when applied to datasets affected by radiation damage. The loss of high-frequency information due to radiation damage significantly hampers the ability to detect continuous heterogeneity, resulting in less clear structural insights.

## 5.2. OUTLOOK

### 5.2.1. DETECTING CONTINUOUS STRUCTURAL HETEROGENEITY IN SINGLE-MOLECULE LOCALIZATION MICROSCOPY DATA

**Generalization.** The method developed for detecting continuous heterogeneity in SMLM datasets may be extended to other imaging modalities, both in microscopy such as Cryo-EM[6, 7] implemented in chapter four, and in medical imaging like MRI, CT, as well as other applications such as lung Extracellular Vesicles imaging and in computer vision, pose, orientation detection.

**Resolution estimation.** Unlike cryo-EM SPA, where additional particles improve resolution following a predictable scaling, our method lacks a clear theoretical foundation for resolution scaling. Investigating a formal framework to relate sample size, noise levels, and structural detectability would enhance the robustness of our methodology.

**Visualization bin size and non-geometric variations.** A critical challenge in our approach remains the determination of the necessary number of samples per bin to reliably detect structural heterogeneity or in other words binning of the latent space. Future work could explore how this number is influenced by localization precision, labeling density, and dataset variability. Furthermore, our method struggles to detect variations that lack a clear geometric interpretation, such as differences in the degree of labeling (DOL). Research into integrating additional metrics that account for such variations beyond purely spatial representations would be beneficial.

**Alternative metrics and disentanglement of variation modes** The current method relies on a single dissimilarity metric, which may have varying sensitivity to different structural variations. Investigating alternative or hybrid dissimilarity measures such as Chamfer distance [8] could improve the detection of multiple variation modes and enhance the robustness of our analysis. Additionally, a fundamental challenge is the disentanglement of different modes of variation within the dataset.

### 5.2.2. DETECTING CONTINUOUS STRUCTURAL HETEROGENEITY IN SINGLE MOLECULE LOCALIZATION MICROSCOPY DATA WITH A POINT CLOUD VARIATIONAL AUTO-ENCODER

**Enhanced model interpretability** Future work could focus on improving the interpretability of the latent space by incorporating structured priors or regularization techniques. This would help ensure that each latent dimension corresponds to a distinct mode of variation, making the learned representations more meaningful and useful for downstream analyses. One possible approach is to establish a "feature library" that links variations in the latent space to specific structural or image-based characteristics, such as size, shape, or density. This could be achieved by correlating latent dimensions with manually extracted features or using supervised techniques to guide the model toward learning more interpretable representations.

**Improved disentanglement of latent representations** To achieve a more independent separation of different modes of variation, methods such as the total correlation variational autoencoder (TC-VAE)[9] could be explored. This would allow for a more clear identification of biologically relevant structural changes in complex SMLM datasets.

**Incorporation of localization uncertainty** By modifying the Chamfer cost function to

include localization uncertainties, PC-VAE could provide a more robust representation of SMLM data. This adaptation would need to balance accuracy and stability, ensuring that the model remains computationally efficient while integrating these uncertainties effectively. The modified Chamfer loss is defined as:

$$\mathcal{L}_{\text{recon}} = \frac{1}{N} \left( \sum_{\mathbf{x} \in \hat{\mathbf{X}}} \min_{\hat{\mathbf{x}} \in \hat{\mathbf{X}}} w_{\mathbf{x}} |\mathbf{x} - \hat{\mathbf{x}}|_2^2 + \sum_{\hat{\mathbf{x}} \in \hat{\mathbf{X}}} \min_{\mathbf{x} \in \mathbf{X}} w_{\hat{\mathbf{x}}} |\mathbf{x} - \hat{\mathbf{x}}|_2^2 \right), \quad (5.1)$$

where the weights  $w_{\mathbf{x}}$  and  $w_{\hat{\mathbf{x}}}$  are defined based on localization uncertainties as:

$$w_{\mathbf{x}} = \frac{1}{\sigma_{\mathbf{x}}^2 + \epsilon}, \quad w_{\hat{\mathbf{x}}} = \frac{1}{\sigma_{\hat{\mathbf{x}}}^2 + \epsilon}, \quad (5.2)$$

where  $\sigma_{\mathbf{x}}^2$  and  $\sigma_{\hat{\mathbf{x}}}^2$  represent the localization variances, and  $\epsilon$  is a small constant to prevent numerical instability. This formulation ensures that points with higher uncertainties contribute less to the reconstruction loss, making PC-VAE more robust to localization noise in SMLM data. This implementation requires input layers for uncertainties similar to point clouds to be added to the network architecture. Furthermore, these changes in the cost function might affect the stability and convergence of the network which should be further studied.

**Application to image registration and particle tracking** Beyond structural analysis, PC-VAE could be leveraged for tasks such as image registration and particle tracking. Particle tracking could be performed by detection of the movements of particles within frames as a continuous heterogeneity mode. Its ability to extract meaningful low-dimensional representations makes it a suitable candidate for improving alignment and tracking accuracy in single particle tracking (SPT) experiments.

**Expansion to temporal and dynamic studies** PC-VAE could be extended to analyze temporal variations in SMLM datasets, enabling the study of dynamic structural changes over time. By learning time-dependent representations, this approach could provide new insights into the evolving nature of nanoscale biological structures. Current implementation works on particles but this can be altered with localizations per frames. This manipulation can result in capturing the time dependent modes of variations within frames.

### 5.2.3. CONTINUOUS STRUCTURAL HETEROGENEITY DETECTION IN CRYO-EM DATA

**Overcoming the roadblocks** Radiation damage is a key challenge in detecting continuous structural heterogeneity in cryo-EM using the proposed algorithm. The issue arises because the similarity metric we use (MSE) is highly sensitive to pixel-level variations, and any disruption in high-frequency information leads to poor pairwise comparisons, resulting in less accurate MDS spaces.

To address this problem, we propose investigation of cost functions that are more sensitive to the structural features of the imaged protein or molecule rather than pixel-level differences. While employing a more complex similarity measure could improve the method's sensitivity, it may also reduce computational efficiency. To mitigate this, parallelization can be implemented across pairs and orientation classes using multi threading to handle the increased computational complexity.

**Efficiency and scalability improvements** While our current approach is effective for moderate dataset sizes, scaling it to larger datasets or real-time applications presents computational challenges. Future work should explore more efficient implementations, including GPU acceleration, optimized data structures, and parallelized processing. These improvements would enable the method to handle larger-scale experiments while maintaining accuracy and speed.

**Disentanglement of complex motions and alternative cost functions** A key challenge in our approach is disentangling complex motions within structural variations to better isolate independent factors of change. Future developments could investigate methods for controlling noise levels and leveraging alternative cost functions, such as the Earth Mover's Distance[10], to improve structural classification. These refinements could enhance the accuracy and robustness of the analysis in diverse datasets.

**Improvements in orientation classification and pre-processing pipeline** Although our method compensates for minor particle misalignment through shifting and rotation, further improvements in particle classification and 2D registration could enhance the algorithm's robustness. This can be performed as orientation class refinement leading to better separation of orientations. Additionally, since the proposed approach relies on pixel-level comparisons between images, advancements in CTF estimation and particle picking would contribute to more accurate detection of continuous structural heterogeneity.

**Semi-automatic latent dimension selection** We observed that even with low-quality datasets, heterogeneity information remains preserved within the MDS space, typically above the first three dimensions. However, this distribution makes an automatic 1D representation challenging, as the dominant mode of variation no longer corresponds to the sought-after continuous structural heterogeneity. We tried to find the 1D latent dimension that carries the heterogeneity information automatically when performing polarity matching by comparing neighbor particle stacks based on the confidence terms defined earlier. However, this did not capture the 1D latent dimension properly. In a future ordering and visualization of 2D particles based on different dimensions (first 4-8 dimensions should be sufficient) manual selection of the dimension that carries the heterogeneity information might prove a necessary step.

## REFERENCES

- [1] T. Huijben, H. Heydarian, A. Auer, F. Schueder, R. Jungmann, S. Stallinga, and B. Rieger, *Detecting structural heterogeneity in single-molecule localization microscopy data*, Nature Communications **12** (2021), p. 3791.
- [2] R. Henderson, *Avoiding the pitfalls of single particle Cryo-electron microscopy: Einstein from noise*, Proceedings of the National Academy of Sciences **110** (2013), pp. 18037–18041.
- [3] Y. Wu, P. Hoess, A. Tschanz, U. Matti, M. Mund, and J. Ries, *Maximum-likelihood model fitting for quantitative analysis of SMLM data*, Nature Methods **20** (2023), pp. 139–148.

- [4] A. Mead, *Review of the development of multidimensional scaling methods*, Journal of the Royal Statistical Society: Series D **41** (1992), pp. 27–39.
- [5] J. Broeken, H. Johnson, D. S. Lidke, S. Liu, R. P. Nieuwenhuizen, S. Stallinga, K. A. Lidke, and B. Rieger, *Resolution improvement by 3D particle averaging in localization microscopy*, Meth. Appl. Fluor. **3** (2015), p. 014003.
- [6] M. Kudryashev, D. Castaño-DÁrjez, and H. Stahlberg, *Limiting factors in single particle Cryo electron tomography*. Computational and structural biotechnology journal **1** (2012).
- [7] G. Tang, L. Peng, P. R. Baldwin, D. S. Mann, W. Jiang, I. Rees, and S. J. Ludtke, *EMAN2: An extensible image processing suite for electron microscopy*, Journal of Structural Biology **157** (2007), pp. 38–46.
- [8] M. A. Butt and P. Maragos, *Optimum design of Chamfer distance transforms*, IEEE Transactions on Image Processing **7** (1998), pp. 1477–1484.
- [9] R. Chen, X. Li, R. Grosse, and D. Duvenaud, *Isolating sources of disentanglement in Variational Autoencoders*, Advances in Neural Information Processing Systems **31** (2018).
- [10] Y. Rubner, C. Tomasi, and L. J. Guibas, *The earth mover's distance as a metric for image retrieval*, International Journal of Computer Vision **40** (2000), pp. 99–121.



# ACKNOWLEDGEMENTS

Once there was nothing not even void or abyss. Not even blackness or emptiness one could imagine, Not even time we can comprehend. Pure absence. Story finished. But no, no, no. Let the story begin.

If we trace back all the actions we perform, like reading this very acknowledgment in the "now", it undoubtedly came with chains of causality in time picking up the book, looking at it, and our brain's precise neurochemical activities resulting in comprehension. But what if there was pure nothing? No story began. No love emerged. No leaf fell from a tree. Not even a droplet from a waterfall would drop and all these numbers, 30, 3, 39, 33, 16, 18, 8, 58, 1, 51, 11, 20, 18, could not be even comprehended.

With this introduction, I want to thank the Creator, the single Creator who not only creates but also directs and guides us to exist and to "be". I understand that many philosophical questions might arise in the reader's mind at this point, and I encourage exploring those questions from various perspectives. Perhaps this research topic itself is more beneficial for a mortal human than any other. Yet, as the old saying goes, "Nothing comes from nothing".

The universe began somewhere in time, life began somewhere else in time, and I was born in 1993. I grew up in a nice and warm family and I want to thank my mom for sharing kindness, love, confidence, and teaching me to be a helping hand than someone who gathers everything for himself. I want to thank my father for his support, hard work, and for covering my inner core with the steel of persistence and a "never give up" mindset.

Taking a step back in time, I want to thank my heroes my grandfather, who taught me to pursue light and goodness at all costs, and my grandmother "Pary" whose patience and resilience, despite physical disabilities, brought immense strength to our family. Her absence leaves a part of our soul missing, never to be repaired. I also thank my other late grandfather, who taught me how to think critically and gave me the ambition to achieve my goals, and my grandmother, who gave nothing but an always smiling face, pure love and goodwill.

Moving forward in time, I started my first year of school on a rainy day. I vividly remember my parents bringing me to a school far from our home, and the jasmine bush wet from the rain, located between the yard and the classroom door. I remember shaking the raindrops off that jasmine bush, sending water flying onto other children :). Later on, our teacher, Mrs. Farrokhi, gave us each an almond painted in gold as a to teach us how to write persian alphabets. I still keep it under my desk.

Soon after, I met my best childhood and lifelong friend my brother Sajjad. I want to thank him for all his support, kindness whether directly or in the form of jokes. Worth mentioning one of his favourite jokes was that he said all my educational achievements were only because I did well on the Arabic language test in the university entrance exam!

I have special thanks to my beloved wife, without whose support and strong spirit I could not have completed this thesis. With you, I live two lives instead of one, have two

brains to think, and two hearts to love. Shortly after migrating, alongside our joyful moments, we faced many challenges, from loss of loved ones to doubts about continuing my PhD or the challenges of migration which we overcame together. I also want to thank my mother, father, and brother in law for their kindness and support.

I have always dreamed of traveling the world, and I felt confined by the borders. Yet at the same time, I want to thank my country, Iran, for everything, from the deep culture passed down through millennia of human civilization to the beautiful nature that embraced us like a protective hand. Despite the current difficulties in Iran, there is a Persian saying that translates to: "No place is like my own home."

I would like to thank my supervisors, Prof. Dr. Sjoerd Stallinga and Prof. Bernd Rieger, for their dedication, support, and commitment to science. Like carving a statue, becoming a better researcher requires hard work and endurance. I have learned much from critical thinking to knowing where to dig deeper and where to cut losses which I believe are the hallmarks of a good scientist.

I also thank all the PhD students and postdocs in our group: Peter, Hamidreza, Awoke, Hamed, Wenxiu, Qingru, Vahid, Enya, Isabel, Valleri, and others; and the staff Annelies, Angela, Ronald, Henry and other colleagues as Fatma and others who helped me in this journey.

Throughout my educational journey and life, I have made many friends some of whom remain close despite the distance. I would like to thank my BSc friends Reza and Saeed, as well as my friends in neighborhood: Mohammad Hosein D., Pooyan, Hosein, and Mohammad Hosein M. We not only shared good times but also offered each other constant support. I'm especially grateful to Mohammadreza, whose pure and mischievous spirit brought so much laughter and fun to my life, spending time with him is still one of my favorite things to do.

During my MSc, I was fortunate to have friends like Mahan, Reza, and Saeed. I also fondly remember my friends from the 17th class in high school, and Amirmohsen for his advises and our mid-night discussions.

I would like to express my deep appreciation to Dr. Stefano Rini not only as my internship supervisor but also as a friend who taught me a great deal in a short time.

Since migrating, I have made new friends; Ali A., Mohammad S., Farzad N., Javad D., Reza, and many others that could be in the list and I want to thank them all for their supports. I would also like to thank my friends who helped and supported me directly in writing this thesis: Vicente for his helps and support in understanding a challenging mathematical paper in beginning stage of my PhD. Maarten and Yi who helped alot to know more about Cryo-EM and Machine learning design. I also want to thank Dariana and Gregory for sharing their 2D NPC dataset with us and Ali saleh Bigdeli for his programming tips and Dr. A. Jakobi, Q. Tao for their helps.

I also want to thank all my teachers and school staff throughout the years for their guidance and support.

I would also like to thank my relatives for their supports: Mohaddese, Massomeh, Mahboobe, Nader, Mansooreh, Jamileh, Najmeh, Dae Jaber, Zandae, Masoomah, Asghar Agha, Ghasem Agha, Agha Asheghan, Agha Mehdi, Agha Koroush, Zandae, Mohammad, Mojtaba, Samineh, Mahsa, Mahshid, Mohadeseh, Mohammad Mehdi, Fatemeh, Maedeh, Hosein, Reyhaneh, Yas, Narges, Yeganeh, Abolfazl, Arshian, Amirali, Mehrad, Ho-

sein, Maryam, Negar, Mohsen, Mehdi, and Sana.

In closing, I remind myself and the reader that life is short. Of the approximately 110 billion people who have ever lived on this planet, only about 5% are still alive and not for a long time. I write this thesis in an era of neglect, where daily we see starving, wounded, or perished innocent children, while others remain happy, supportive, or silent. Or we see corrupted politicians controlling us. I believe our lives are too short to be weak or aimless. We will surely leave everything behind, but what remains forever are the kindness and love we share, the good paths we leave, and the helping hands we extend. I would like to finish by Saadi's poem translated to English regarding this:

"Human beings are members of a whole,  
In creation of one essence and soul.  
If one member is afflicted with pain,  
Other members uneasy will remain."



# CURRICULUM VITÆ

**Sobhan HAGHPARAST**

## EDUCATION

- **PhD in Applied Sciences**, TU Delft, Netherlands  
Continuous structural heterogeneity detection in high-resolution microscopy.
- **MSc in Biomedical Engineering**, Sharif university of technology, Tehran, Iran  
A novel frame rate up-conversion method to decrease blocking and blurring in echocardiographical videos.
- **BSc in Electrical Engineering**, TU Noshirvani, IR Analogue to digital converters using FPGA
- **High school**, National Organization for Development of Exceptional Talents (NODET), Sari, IR