Simulating the cellular context in synthetic datasets for cryo-electron tomography

Antonio Martinez-Sanchez, Lorenz Lamm, Marion Jasnin and Harold Phelippeau

Abstract— **Cryo-electron tomography (cryo-ET) allows to visualize the cellular context at macromolecular level. To date, the impossibility of obtaining a reliable ground truth is limiting the application of deep learning-based image processing algorithms in this field. As a consequence, there is a growing demand of realistic synthetic datasets for training deep learning algorithms. In addition, besides assisting the acquisition and interpretation of experimental data, synthetic tomograms are used as reference models for cellular organization analysis from cellular tomograms. Current simulators in cryo-ET focus on reproducing distortions from image acquisition and tomogram reconstruction, however, they can not generate many of the low order features present in cellular tomograms.**

Here we propose several geometric and organization models to simulate low order cellular structures imaged by cryo-ET. Specifically, clusters of any known cytosolic or membrane bound macromolecules, membranes with different geometries as well as different filamentous structures such as microtubules or actin-like networks. Moreover, we use parametrizable stochastic models to generate a high diversity of geometries and organizations to simulate representative and generalized datasets, including very crowded environments like those observed in native cells.

These models have been implemented in a multiplatform open-source Python package, including scripts to generate cryo-tomograms with adjustable sizes and resolutions. In addition, these scripts provide also distortion-free density maps besides the ground truth in different file formats for efficient access and advanced visualization. We show that such a realistic synthetic dataset can be readily used to train generalizable deep learning algorithms.

Index Terms— **Cryo-Electron Tomography, Deep Learning, Image processing, Scientific Computation, Synthetic data generation**

I. Introduction

Cryo-electron tomography (cryo-ET) has emerged as a unique technique to generate tridimensional (3D) representations of the cellular context at sub-nanometric resolution [1]. However, volumetric images obtained by cryo-ET, called tomograms, suffer from high levels of noise and anisotropic distortions, mainly due to the missing wedge [2], which greatly complicate their analysis.

Due to its success in image processing and computer vision in recent years, deep learning (DL) algorithms, and especially convolutional neuronal networks (CNNs) [3], are being adapted to process cryo-ET tomograms at different steps of the workflow: reconstruction [4], denoising [5], [6], segmentation [7], [8] or macromolecular localization [9]. CNNs enable efficient analysis of local features in images because they take advantage of the inherent regular grid structure of digital images. When the training set is representative and large enough, CNN-based methods outperform classical approaches like template matching (TM) [10]. Nevertheless, the major bottleneck to apply DL in cryo-ET is the lack of ground truth. In most cases it is not possible to generate a reliable labeled dataset for a proper supervised training.

Regardless of the method used for segmenting cellular compartments and localizing macromolecules in cryo-tomograms, the next step is usually a quantitative analysis. In cryo-ET, statistical approaches are used to carry out macromolecular censuses [11], [12], analyze the distance between two proteins [13], study protein-membrane distance distribution [14] or prove the co-localization between two macromolecules [15]. The second-order statistics are more sophisticated [16], as they allow to analyze macromolecular organization on different scales, as well as merge and compare information coming from different tomograms. Second-order statistics have recently been used in cryo-ET to evaluate the clustering organization of some protein complexes under different functional states [17], show nanodomains defined by the co-localization of proteins bounded to two aligned membranes [18], describe the liquid-like organization of densely packed proteins by adjusting a statistical mechanics model [19]. In addition, rotations of macromolecules are typically analyzed to determine the rigidity of a bounding, e.g. between two macromolecules [20] or proteins and membranes [21]. Additionally, ad-hoc analysis can be defined to study filament networks [22] and membranes organization [23], [24]. In general, these statistical analyses require a synthetic null-model working as reference [16], which also allows to provide statistical confidence.

Cryo-ET data acquisition is costly as it requires high-end equipment and a lot of human expertise. Unfortunately, determining the applicability of a dataset for studying a specific

Submitted in XX/XX/XXXX. "This work was supported by Ramon y Cajal program, RYC2021-032626-I, of the Spanish State Research Agency (AEI) and Attract-RYC 2023 program of the University of Murcia. L.L. was funded from the Munich School for Data Science (MUDS) and a fellowship from the Boehringer Ingelheim Fonds."

A. M-S. Author is with the Department of Information and Communications Engineering, Faculty of Computers Sciences, University of Murcia, Campus de Espinardo 30100 Murcia, Spain (e-mail:anmartinezs@um.es). Corresponding author.

M. J. Author with the Helmholtz Pioneer Campus, Helmholtz Munich, 85764 Neuherberg, Germany; Department of Chemistry, Technical University of Munich, 85748, Garching, Germany.

H. P. Author is with Thermo Fisher Scientific, Bordeaux, France.

L. L. Author is with Helmholtz AI, Helmholtz Munich, 85764 Neuherberg, Germany.

structure is only possible after tomogram reconstruction, the last step in the workflow. Synthetic data can be used to streamline imaging parameters as well as to ensure sample appropriateness for solving a specific biological question. Simulations have already been used for particle picking and subtomogram averaging assessment [25], [26]. Recently, a tool specifically designed for membrane-embedded proteins has been proposed [27], however, it is limited to a single spherical membrane with uniformly distributed isolated proteins. Simulated annealing and molecular dynamics are used afterward to improve the packing and avoid volume overlapping by approximating molecules to spheres [26].

As shown above, simulating synthetic data containing lower-order structures present in the cellular environment will contribute greatly to the successful application of DL, to analyze quantitatively the organization of structures in cryo-ET, and assist the acquisition and interpretation of experimental data. There are already computational solutions for generating artificial data in transmission electron microscopy (TEM), but they focus on modelling the physics of image formation from atomic representations of macromolecules [28]–[30]. Methods for simulating cellular environments represent either isolated structures or a mixture of purely random distribution of macromolecules with some spherical vesicles [26], [31], [32]. In [33] some membrane-bound proteins are modelled but just by simple blobs in 2D. A concurrent pre-print [34] generates a more diverse scenario than previous approaches, which includes vesicles and membrane-proteins but not curved filaments. More importantly, this pre-print shows that coarse synthetic data already allows to train usable CNN models.

Here, we focus on producing richer contextual information, and directly generate synthetic 3D density maps, also known as phantoms, containing membranes with different combinations of local curvatures, clusters of macromolecules (either cytosolic or membrane-bound), as well as polymers such as microtubules or actin-like networks. Together with these clean 3D images, motif lists in different formats are provided as ground truth, allowing the generation of micrographs using any external TEM simulator to simulate the distortions introduced by the microscope and 3D reconstruction methods. Moreover, the introduction of generative neural networks holds the promise to produce datasets with more realistic noise and distortion [35], they can reproduce noise and distortions directly from the synthetic micrographs [36], but this is outside the scope of this study. In addition, techniques such as data augmentation [37], knowledge transfer [38] or contrastive learning [32] may allow to train DL-based algorithms without having a perfect representation of the experimental data. For statistical organization analysis, null-models do not require simulating data noise and distortions.

The simulation of high-order cellular structures such as organelles (mitochondrion, centriole, ...) or membranes associations (synapses, Golgi, ...) is beyond the scope of this paper. Here we focus on low-order structural features present in almost all cellular contexts. Low-order features are independent of each other and can thus be generated from parametric mathematical models to control their geometry and organization. In addition, we use parametrizable stochastic models to generate a high diversity of different combinations of geometries and organization to simulate representative and generalized datasets.

The rest of this paper is organized as follows: First, the developed methods are described, we detail how to model the geometry of each structure and their organization, and we explain how these models were implemented. Then, we show examples of synthetic data generated for replicating the cellular context in cryo-ET and use these data to train a generalized deep learning algorithm for semantic segmentation. We also evaluate certain properties of the proposed algorithms. Finally, the proposed methods and the obtained results are discussed.

II. METHODS

In this section, we first explain the mathematical models used to parameterize the cellular structures. Second, we describe the algorithms used to generate the synthetic data.

A. Structural Models

The cellular context is represented by a density map $D(\mathrm{x}) : \mathcal{V} \subset \mathbb{R}^3 \mapsto \mathbb{R}$ being $\mathrm{x} \in \mathbb{R}^3$ a point and $\mathcal{V}$ the volume of interest (VOI), typically a cuboid whose dimensions are defined by the output tomogram, but in practice it can have any arbitrary shape.

1) Membranes: Membranes are modelled as parametric surfaces with double Gaussian profile along their local normal. The purpose of this profile is to simulate the lipid bilayer, which, in the TEM, looks like two parallel and very close electron-dense sheets. Although all biological membranes share the same structure, small visual differences may appear due to their distinct molecular composition or dissimilar imaging conditions. To account for these variations, the membrane profile is controlled by two parameters: membrane thickness or distance between both layers, t, and the layer variance, σ_l, which controls the thickness of the layers. Below, we describe the parametric surfaces used to model the membrane geometries and how the double-layered profile is incorporated into them.

Sphere: the outer, L_o, and inner, L_i, layers are defined by

$$L^S = \left\{ \mathrm{x} \in \left[B^S(\mathrm{c}) \geq (l^S - 1) \wedge B^S(\mathrm{c}) \leq (l^S + 1) \right] \right\}. \quad (1)$$

Where $B^S(\mathrm{c}) = \|\mathrm{x} - \mathrm{c}\|$, $\mathrm{c} = (c_x, c_y, c_z)$ is the sphere center, $l_o^S = r + t$ and $l_i^S = r - t$ are the values of l for each layer, and r as the sphere radius.

Ellipsoid: As above, the ellipsoidal layers are defined as

$$L^E = \left\{ \mathrm{x} \in \left[B^E(\mathrm{a}) \geq 1 \wedge B^E(\mathrm{a}) \leq 1 \right] \right\}. \quad (2)$$

Where $B^E(\mathrm{a}) = \frac{x - c_x}{a_x} + \frac{y - c_y}{a_y} + \frac{z - c_z}{a_z}$, with $\mathrm{a} = (a_x, a_y, a_z)$ the ellipsoid semi-axis lengths, for each layer the actual values for these lengths are $\mathrm{a}_o = \mathrm{a} + t$ and $\mathrm{a}_i = \mathrm{a} - t$.

Torus: Given a point $\mathrm{x} = (x, y, z) \in \mathbb{R}^3$ the two toroid layers are defined by

$$L^T = \left\{ \mathrm{x} \in \left[B^T(d, e) \geq 1 \wedge B^T(d, e) \leq 1 \right] \right\}. \quad (3)$$

Where $B^T(d, e) = \left(d - \sqrt{(x - \mathrm{c}_x)^2 + (y - \mathrm{c}_y)^2} \right)^2 + (z - c_z)^2 - e$, with d the major torus radius and e the minor torus

radius, which have different values for each layer: $e_o = e + t$ and $e_i = e - t$.

The layers L described above have a non-smooth profile (hard edge). A 3D Gaussian filter $\mathcal{G}$ is thus applied to obtain a smoothed version $L^{\mathcal{G}} \in C^2$ using the layer thickness σ_l defined above:

$$L^{\mathcal{G}} = \mathcal{G}(L, \sigma_l) \qquad (4)$$

In addition to representing membranes as scalar fields, which is useful to generate the output density map, it is convenient to also represent them using 2-manifolds (surfaces) as

$$\mathcal{L} = \left\{ \mathrm{x} \in \mathbb{R}^3 \mid \mathrm{x} \text{ is a critical point and } \lambda(\mathrm{x}) > 0 \right\}. \qquad (5)$$

Being λ the largest eigenvalue of the Hessian Matrix of the scalar field $L_i^{\mathcal{G}} + L_o^{\mathcal{G}}$ at point x. Since $\mathcal{L}$ is the surface in the middle of both the inner and outer layers of the membrane, it can be used as the reference to localize membrane-bound macromolecules and to measure certain geometric properties of the membrane, such as areas or local curvatures.

2) Filaments: Filaments are modelled as flexible helical polymers. We consider a polymer as a chain of structural units. A structural unit is defined as a density function in 3D $M(\mathrm{x}) : \mathbb{R}^3 \mapsto \mathbb{R}$.

The normalized tangent vector, $\vec{\tau}_i$, at position x_i in the filament is used to determine the position of the next structural unit in the chain as

$$\mathrm{x}_{i+1} = d_M \cdot \vec{\tau}_i + \mathrm{x}_i. \qquad (6)$$

With d_M the distance between two consecutive structural units. Filament centerlines are modelled as helical curves aligned along the Z-axis so the tangent at a point is computed as

$$\vec{\tau}(s) = q \left[g, -f \sin(sq), f \cos(sq) \right]. \qquad (7)$$

Where s is the length of the curve from its origin to the position of the point and $q = (f^2 + g^2)^{-1}$. The parameters of the helix curve are the radius f and the slope f/g. These parameters can be assigned from a given persistence length p, the physical parameter used to characterize polymer rigidity [39], using the following expressions

$$\kappa = \frac{\arccos e^{-d_M/p}}{d_M}, \quad f = \frac{1}{\kappa}, \quad g = h_z \cdot f. \qquad (8)$$

Where κ is the curve curvature in the XY-plane and $h_z \in [0, 1]$ determines the curve elevation along the Z-axis as a fraction of the curvature, thus giving the curve also a torsion.

After being placed on the filament curve, each structural unit can be additionally rotated around the tangent of the curve by an angle of

$$\alpha = \alpha_0 + i \frac{2\pi h_l^I}{d_M}. \qquad (9)$$

With α_0 an arbitrary starting angle to allow randomization, i the position of the structural unit on the filament, and h_l^I the length on the curve required to complete a turn. Consequently, we have two helices, the curve, or centerline, where the structural units are placed to give curvature and torsion, and the helix formed by the angle between two consecutive structural units, which model a possible helical inner structure of the filament as for actin.

After defining the curve used to model the filament centerline from a given persistence length, the different structural units are inserted on the centerline. Below we present the models for the structural units for the two types of filaments used so far:

- **Microtubules (MT):** the structural unit is modelled as a set of spheres with a radius of 40Å. Structural units are composed of 13 smooth-edged 3D solid spheres placed in a circle of radius 125Å on the XY-plane every $2\pi/13$ radians, thus forming a complete circle as described in [40]. The distance between two consecutive structural units is set to 60Å, which is 1.2 times the unit length in the Z-axis (sphere diameter). MTs do not have a helical inner structure so the turn angle is $\alpha = 0$.
- **Actin-like:** the structural unit is modelled as a pair of spheres with a radius of 25Å aligned with the X-axis and separated by 50Å. The distance between two consecutive units is set to $d_M = 60$Å, which is 1.2 times the sphere diameter. The turn angle is $\alpha = 2\pi/12.5$ because $h_l^I = 750$Å. This model approximates and averaged actin filament as described in [41].

3) Macromolecules: Cellular macromolecules can float in the cytoplasmic matrix or be bound to membranes. They can also be associated with filaments, but this case is not yet considered here. In addition, macromolecules can be randomly distributed or aggregated, creating clusters. Here we use the self-avoiding worm-like chain (SAWLC) model for polymers [42] to create these clusters. The SAWLC model also allows for random distribution of macromolecules if the polymer size is equal to that of a single macromolecule. Unlike the filamentous polymers described above, we assume here that the polymers have no rigidity and thus the clusters have no specific shape.

Cytosolic: in the SAWLC model each structural unit (or macromolecule), M_i, is placed at a random point of the next set

$$\mathcal{C}_i = \mathcal{S}(\mathrm{x}_{i-1}, d_i) \cap \mathcal{V}. \qquad (10)$$

Where d_i is a variable distance from the previous macromolecule M_{i-i} in the sequence, and $\mathcal{S}$ a sphere centered at position, x_{i-1}, and radius d_i.

Membrane-bound: in the SAWLC model each macromolecule, M_i, is placed at a random point of the next in the next set

$$\mathcal{M}_i = \mathcal{S}(\mathrm{x}_{i-1}, d_i) \cap \mathcal{V} \cap \mathcal{L}. \qquad (11)$$

In comparison with cytosolic macromolecules, membrane-bound proteins must also be embedded in a membrane set (see Section II-A.1 Equation 5) $\mathcal{L}$.

Heterogeneous clusters: a SAWLC cluster, either cytosolic or membrane-bound, can contain J different types of macromolecules, M^j, with $j \in [1, J] \subset \mathbb{N}$.

B. Stochastic Transformation Models

Each new inserted structural unit has a specific rotation, R, with respect to the reference structural model and a specific translation, T, to its final position in D. Depending on the structure type, the rigid transformations R and T are generated

according to different stochastic models (transformation generators G^T), thus generating a representative variety of structures with different positions and rotations. Similarly, the geometric parameters for membranes and filaments are also generated using stochastic models (structural generators G^S).

1) Membranes: Translations and rotations are generated in the same way for all membrane varieties. Translations are expressed with a coordinates vector generated according to the complete spatial randomness (CSR) model [43], [44]. Rotations are expressed as unit quaternions and are generated using the algorithm described in [45] to approximate a uniform random distribution.

The membrane thickness t and layer thickness σ_l parameters are obtained from a random uniform distribution in a specified range. To ensure that all generated thicknesses are always within a valid physical range [46], [47], the distribution limits must lie within $[2.5, 4.5]$ nm.

Sphere: the radius r is generated from a random uniform distribution within the following range: the minimum corresponds approximately to the smallest vesicles in cells (typically 75Å) and the maximum with the diameter of $\mathcal{V}$.

Ellipsoid: the lengths of the semi-axis a $= (a_x, a_y, a_z)$ are taken from random uniform distributions within the same range as the spheres. In addition, the eccentricities are defined as

$$\epsilon_{xz} = \sqrt{1 - \left(\frac{a_z}{a_x}\right)^2}, \quad \epsilon_{yz} = \sqrt{1 - \left(\frac{a_z}{a_y}\right)^2}, \quad (12)$$

and must be smaller than a given value.

Toroid: the two radii are obtained using a uniform random distributions similarly to sphere radius, then the greater value is assigned to the major radius d and the lower to the minor radius e.

2) Filaments: Unlike membranes, which are independent structures, filaments consist of a sequence of several structural units. Consequently, the rigid transformations for the structural unit, M_i, depend on those applied to the previous unit, M_{i-1}. Therefore a random generator for rigid transformations is only required for the first structural unit, M_0, using the same distributions used for membranes.

For the filament geometry expressed as a helical curve, the only free parameters are the persistence length p, which controls the circular radius (and curvature) in the XY-plane, and the elevation in the Z-axis as a fraction of the circular radius h_z, which controls the torsion. The values of p are derived from a random exponential unit distribution added to a minimum value of persistence length, this value sets a maximum to the filaments' flexibility. The values of h_z are obtained from a random uniform distribution between a given value (≤ 1) and 1.

In the case of actin-like filaments, it is possible to create a branched network, instead of generating isolated filaments. In a branched network, each new filament may start from any location in the VOI, like regular filaments thus starting a new network, but with a certain probability, $P_b \in [0, 1]$, it starts from a random point of an already inserted filament, creating a branch within an already existing network. These parameters thus control the average number of networks to be generated for a given maximum occupancy and the branching density (number of branches per network length).

C. Algorithms

This section contains the description of the different algorithms used to insert the structures into the VOI, $\mathcal{V}$. First, the order chosen to process each structural type seen above is listed. Then, the algorithms for inserting each specific structure are described.

1) Structure type selection: Algorithm 1 details the procedure for processing the structural models. Only after all the structures of one type have been inserted, then we move to the next type. In each type, the order of each entity, e.g. membrane geometry (spherical, ellipsoidal or toroidal), is an input that can be controlled by the user. The criterion used to establish the order specified in Algorithm 1 is to place potentially larger and more rigid structures first. We assume that a cluster of small macromolecules following the flexible SAWLC model will be inserted more easily than a membrane or a filament in a highly fragmented volume.

Algorithm 1 Order of insertion for the different types of structures.

Require: VOI; lists of structural models (one per type)
Ensure: Density map; list of inserted structures; output VOI
    (empty space after inserting of all structures)
    Insert membranes using Algorithm 2
    Insert filaments using Algorithm 3
    Insert cytosolic macromolecules using Algorithm 3
    Insert membrane-bound macromolecules using Algorithm 3

---

2) Structure insertion: The procedure for inserting membranes is described in Algorithm 2, for filaments and macromolecules this procedure is identical and explained in Algorithm 3. Both algorithms are based on the trial-and-error principle: one tries to insert a new instance of a structure according to some geometrical constraints defined by the structural generator G^S into a space defined by the transformations generator G^T, if this is not possible due to overlap with already inserted structures or because a fraction of the structure is outside of the VOI, then the algorithms try again by generating new geometric and transformation parameters. An additional stopping criterion is added to ensure the end of the process after a certain number of failed tries. Information about each structure is entered in the form of an input list Q including: the structural unit M, the target occupancy O and the statistical models for generating the structural parameters G^S and the rigid transformations G^T. Occupancy is the percentage of a volume occupied by one or more structures.

In Algorithm 2, for each membrane model i in the list Q, an instance M with parameters generated by G^S is inserted into the updated VOI, $\mathcal{V}_o$, according to the rigid body transformations R and T generated by the transformation generator G_T. The insertion loop for a membrane model i ends when the target occupancy O_i is reached, and then the algorithm proceeds to the next membrane model in the list $i+1$ until all models have been processed. The input $\mathcal{V}$ is updated

for each inserted membrane to exclude the already occupied volume. Therefore, $\mathcal{V}_o$ will define the final empty space and corresponds to the VOI input for the next algorithm, which will process the next structure type.

Algorithm 2 Insertion of membranes in the VOI.

Require: VOI, $\mathcal{V}$; list of membrane models, Q^M, each entry i is a tuple (O_i, G_i^S, G_i^T); N_t, maximum number of trials for the insertion of a membrane instance.
Ensure: D, density map; Q_o, list of inserted structures; $\mathcal{V}_o$, output VOI.
  $\mathcal{V}_o \leftarrow \mathcal{V}$
  for $i \leftarrow 1$ to length(Q^M) **do** ▷ Loop for membrane models
    $n_f \leftarrow 0$
    $(O_i, G_i^S, G_i^T) \leftarrow Q_i^M$
    $O = 0$
    while $(O < O_i) \wedge (n_f < N_t)$ **do**   ▷ Add membranes with model G_i^S until O_i is reached
      Generate membrane M from G_i^S
      Get T and R from G_i^T
      $M' \leftarrow M(R(\mathrm{x}) + T)$ ▷ Rigid transformation
      if $\mathcal{V}_o \cap \check{M}' = \check{M}'$ **then**   ▷ $\check{M}$ is the volume occupied by M. This ensures volume exclusion
        $\mathcal{V}_o \leftarrow \mathcal{V}_o \cup \check{M}'$
        $O \leftarrow O + 100 \cdot (\text{volume}(M)/\text{volume}(\mathcal{V}_o))$
        Add M' to D
        Add M' to Q_o
        $n_f \leftarrow 0$
      else
        $n_f \leftarrow n_f + 1$ ▷ Update the number of consecutive fails
      end if
    end while
  end for

Algorithm 3 Insertion of filaments or macromolecules in the VOI.

Require: VOI, $\mathcal{V}$; list of models, Q^S, each entry i is a tuple (O_i, G_i^S, G_i^T); N_t, maximum number of trials.
Ensure: D, density map; Q_o, list of inserted structures; $\mathcal{V}_o$, output VOI.
  $\mathcal{V}_o \leftarrow \mathcal{V}$
  for $i \leftarrow 1$ to length(Q^M) **do**   ▷ Loop for models
    $n_f \leftarrow 0$
    $(O_i, G_i^S, G_i^T) \leftarrow Q_i^S$
    $O = 0$
    while $(O < O_i) \wedge (n_f < N_t)$ **do** ▷ Loop for filaments or clusters
      $S_i \leftarrow \emptyset$ ▷ Empty filament or cluster
      Generate seed structure unit S_i^0 from $G_i^S(\emptyset)$
      j = 0
      Get T and R from G_i^T
      $(S_i^j)' \leftarrow S_i^0(R(\mathrm{x}) + T)$
      while $\mathcal{V}_o \cap (\check{S}_i^j)' = (\check{S}_i^j)'$ **do** ▷ Loop for structural units, condition ensures volume exclusion
        $\mathcal{V}_o \leftarrow \mathcal{V}_o \cup (\check{S}_i^j)'$
        $O \leftarrow O + 100 \cdot (\text{volume}(S^j)/\text{volume}(\mathcal{V}_o))$
        Add $(S_i^j)'$ to S_i
        Add $(S_i^j)'$ to D
        $j \leftarrow j + 1$ ▷ move to next structural unit
        Generate next structure unit S_i^j from $G_i^S(S_i^{j-1})$
        $\hat{S}_i^j \leftarrow S_i^j(R(\mathrm{x}) + T)$
      end while
      if $S_i = \emptyset$ **then**
        $n_f \leftarrow n_f + 1$
      else
        Add S_i to Q_o
        $n_f \leftarrow 0$
      end if
    end while
  end for

Regardless of overlapping, membrane instances are independent of each other. However, for filaments and macromolecules the insertion of the next structural unit depends on the previous one, except for the first one. Algorithm 3 is an extension of Algorithm 2 to handle this type of situation. It thus includes three nested loops. The first one iterates over all the introduced structural models. The second one attempts to reach the requested occupancy by inserting structural instances of a filament or cluster. The third one grows the filament or cluster by adding structural units.

Randomization in Algorithms 2 and 3 is introduced through the stochastic structural and transformation models G^S and G^T. The addition of the counter, n_f, in Algorithm 2 and 3, forces a restart of the structure insertion loop after N_t failures. As shown in Section III-C that helps to achieve high concentration of structures and avoids excessively long running times.

D. Implementation details

Although this paper is dedicated to the simulation of cellular structures rather than the reproduction of the 3D image acquisition process in TEM, the implementation presented here is capable of reconstructing 3D tomograms containing some of the basic distortions inherent to cryo-ET such as noise, micrograph misalignment and missing wedge. The micrographs are generated as 2D projections at specified tilt angles using IMOD [48]. The noise is modeled by a Gaussian distribution, which is added to the micrographs, in order to approximate the tomogram contrast for a certain signal-to-noise ratio (SNR). Since the ground truth is known, noise level is adjusted to achieve certain SNR. Target SNRs are randomly generated from a uniform distribution within a specified range. The misalignment of the micrographs is modeled by applying random offsets on the X and Y axes of the micrographs in a sinusoidal model to penalize high tilt angles. Finally, the 3D tomogram is reconstructed from the micrographs using IMOD.

There are more accurate tools than IMOD for reconstructing cryo-tomograms from synthetic data [27], [29]. They have more complex models for noise, which also take into account solvent contribution. Since simulating the TEM process is still an active research field, here we generate enough output infor-

Fig. 1. Simulation of the cellular context. (a) A 2D slice in the XY-plane of the generated density map, the blue box shows an example of how the growth of a microtubule was stopped to avoid overlap with another structure. (b) Polygonal ground truth data, membranes in yellow, actin-like network in blue, microtubules in green, macromolecules (included membrane-bound) in different shades of red.

Fig. 2. Reconstruction of tomograms. (a) A 2D slice in the XY-plane of a reconstructed tomogram corresponding to the density shown in Fig. 1. Scale bar 200 nm. (b) Zoom of the XY-plane showing membrane-bound proteins in the upper left, the label map is in the upper right, the dashed lines show the intersection with the YZ-plane slices shown in the bottom row. In the reconstructed YZ-slice (bottom-left) the membrane disappears at the apex and the macromolecules are elongated in the Z-axis relative to the ground truth (bottom-right) due to the missing wedge distortion introduced during reconstruction. The bilayer profile of the membrane can be appreciated in the magnified images. Scale bars 20 nm.

mation (shape, position and rotation for each structural unit) to feed a TEM simulator. For example, very recently a software based on generative neuronal networks has been published to reconstruct tomograms from the synthetic micrographs [36] having distortions with realistic appearance.

III. RESULTS

Here we present the code implementation and synthetic data generated to show the potential of the proposed approach. All data were generated using a workstation running Ubuntu 22.04 LTS with an Intel® Core™ i7-11700 2.50GHz processor, 128GB RAM and a GeForce RTX™ 3090 Ti GPU.

A. Code

The software implementation of the models and algorithms described here is in a Python 3 package called 'PolNet' open-source available in the public repository

https://github.com/anmartinezs/polnet.git. The scripts and settings for generating the next results are included in this repository.

B. Cellular cryo-ET

The goal of this experiment is to simulate an exemplary tomograms with all currently available types of structures: membranes (spherical, ellipsoidal and toroidal), cytoskeleton (actin-like networks and microtubules), cytosolic and membrane-bound macromolecules. The size of the generated tomograms is 1000x1000x250 voxels, with a voxel size of 10Å, which is the typical configuration used today for tomogram visualization, molecular localization, and segmentation in cryo-ET.

Fig. 1.a shows an example of a generated clean density, free of noise and missing wedge. A ground truth semantic segmentation is also generated, the equivalent information is also stored in VTK format (see Fig. 1.b), which is more suitable for 3D visualization. In addition, the skeletons of the structures

are also stored in VTK format, with filaments represented by their centerlines and clusters of macromolecules represented by a sequence of points connected by lines, with single points consequently corresponding to single macromolecules.

Fig. 2.a shows a 3D reconstructed tomogram from 2D projected micrographs ($\pm 60°$) that were randomly misaligned using a sinusoidal random model to penalize high tilt micrographs, with a mean value of 1 and a maximum misalignment of 1.5 pixels. As a result, the reconstructed tomogram suffers from distortions due to angular sampling and missing wedge. Gaussian noise was added to the micrographs to simulate the low contrast of the cryo-ET data by setting a target SNR randomly in the range $[1, 2]$. Fig. 2.b shows the missing wedge effect by comparing the reconstructed tomogram with the ground truth. The missing wedge effect is particularly visible in the YZ-slice where the membrane disappears and the macromolecules are elongated on the Z-axis. The reconstructed and ground truth tomograms can be used as input information to train/test machine learning algorithms for image segmentation.

This example shows how Algorithm 1 can merge the different types of structures to generate a complex environment, which resembles the low order structures present in the cellular context.

C. Simulation of the crowded environments

The cell cytoplasm and membranes usually have a high concentration of structures. Since structure insertion is based on a trial-and-error algorithm, the execution time increases with the concentration or occupancy. To evaluate structural insertion algorithms, here we run experiments to evaluate the insertion of cytosolic macromolecules, membrane-bound proteins and filaments separately. The membrane insertion algorithm is not evaluated because the total membrane occupancy generally does not exceed 1%. In these experiments, all tomograms had a size of 500x500x200 voxels, with a voxel size 10Å.

Regardless of the type of structure, the input occupancy is a target value. A trial-and-error algorithm cannot guarantee convergence within a certain time, which is especially difficult when the occupancy is very high. For macromolecules, Algorithm 3 terminates after failing N_t consecutive times to initialize a new structure. Increasing N_t therefore allows to obtain tomograms with higher densities at the cost of higher running times due to a greater number of structure insertion failures (see Fig. 3.a-c). An $N_t = 100$ achieves a target occupancy of 25% for cytosolic macromolecules (see Fig. 3.d-f), which already corresponds to a very crowded tomogram. Increasing $N_t = 1000$ allows to reach nearly 30% at the cost of approximately doubling the time.

In contrast to cytosolic macromolecules, for which occupancy is calculated as the volume fraction (in percent) of the occupied VOI, the target occupancy for membrane-bound macromolecules represents the maximum area fraction (in percent) of the membrane surface occupied, as membrane-bound macromolecules can only be inserted on the membrane surface. Therefore, the input occupancy does not match the measured volume occupancy (see Fig. 4.a), although their dependence is approximately linear. Nanodomains, or clusters,

of membrane-bound proteins are evident at low occupancies (Fig. 4.b). It is also possible to generate membranes with a high density of membrane-associated proteins, as in synaptic vesicles, viruses or endoplasmic reticulum (see Fig. 4.c).

As for cytosolic macromolecules, the occupancy is computed from the filament volume. Occupancy values greater than 10% are not necessary because, in general, they do not correspond to realistic conditions in the cell. Fig. 5.b shows an occupancy of 7% which already results in a very dense actin-like network. An important parameter to control the shape of the filament networks is the branching probability for the filament P_b (see Fig. 5.a). The contrast between Fig. 5.b and 5.c shows that a high P_b value generates a densely concentrated filament network in specific regions, while conversely, a low P_b value generates a more evenly distributed network along the VOI.

D. Training a deep learning model

To assess the usability of our synthetic data for training deep learning algorithms, we use a specific implementation of the U-Net architecture known as nnU-Net [49]. This framework is particularly known for its good performance without the need of manual hyper parameter tuning in the field of semantic segmentation for biomedical images. We trained nnU-Net to segment cryo-electron tomograms into the classes of membranes, microtubules, actin-like filaments, ribosomes, membrane-bound proteins and background (i.e., non of the other classes). To this end, we utilized the dataset presented in Subsec. III-B, and used ten reconstructed tomograms (see an example in Fig. 2.a) for training. For the model design, we chose nnU-Net 3D full resolution option, and the option fold all for training. We use the v2 of nnUNet software. As depicted in Fig. 6, the resulting prediction from our network validates our approach. Despite being trained solely on synthetic data, the network was capable of coarsely, but effectively, segmenting the main features of several *in situ* tomograms from different samples.

E. Evaluation of the deep learning model

We have also processed tomograms providing as reference a segmentation to quantitatively validate the model generated in Subsection III-D. Specifically, we have taken publicly available tomogram EMD-11992, which is accompanied with segmentation for microtubules, actin filaments and ribosomes [51], and EMD-13671, with segmentation provided by authors [52] for intermediate filaments, actin filaments and membranes.

The model output is a semantic segmentation of the tomograms, therefore we should use Dice. However, as described in [53], the accurate segmentation evaluation for tubular structures such as filaments requires novel similarity metrics to preserve topological features meanwhile arbitrary features like thickness are neglected. Authors in [53] proposes clDice (centerline Dice) for tube-like structures in 2D images, here we extend this metric to 3D for processing filament like structures. Analogously, we define and compute sDice (surface Dice, described also in [54]) for membranes and pDice (point Dice)

Fig. 3. Simulation of cytosolic macromolecules. (a-c) Macromolecule insertion performance for different target occupancies where two N_t values are evaluated. (a) Occupancy measured as the fraction of voxels belonging to macromolecules relative to the VOI volume. (b) Execution times. (c) Number of unsuccessful attempts to start a new macromolecule cluster. (d) 2D slice of a tomogram with 1% occupancy. (e) 2D slice of a tomogram with a 25% occupancy. (f) 3D view of the tomogram with a 25% occupancy.

Fig. 4. Simulation of membrane-bound proteins. (a) Measured output volume occupancy compared to the input target occupancy on membrane surfaces. (b) 3D visualization of a tomogram with a 7% target occupancy for membrane-bound proteins. (c) Target occupancy of 39%. (b-c) Membrane surfaces in transparent yellow, membrane proteins in different levels of red.

for macromolecules. Consequently, clDice variant is computed for microtubules, actin or any other filamentous structure, pDice for ribosomes because these are macromolecules, and sDice for membranes. These metrics are calculated from the previous computation of the topology precision (TP), the fraction of the predicted segmentation skeleton within the ground truth, and the topology sensitivity (TS), the fraction of the ground truth skeleton within the predicted segmentation. In this paper, all Dice values shown are computed from pDice, clDice or sDice depending on the processed structures as described above.

Table I contains the quantitative evaluation results obtained for the structures with reference segmentation in EMD-13671. For membranes, the segmentation agrees very well the refer-

ence. However, it is important to consider in this case the reference segmentation has been provided by a DL model, MemBrain-v2 [54], specifically designed for segmenting only membranes and trained with a large and diverse experimental dataset. More importantly, the high sensitivity value shows that our model segments almost everything predicted by MemBrain-v2, but as shown in Figure 7 our model seems to partially recover regions missed by missing wedge. Regarding the results for actin filaments, Dice scores are very low, but specificity values are very high. Additionally, Figure 8 shows that besides actin filaments, intermediate filaments are also segmented by our model. To further investigate the capacity of our model to become a general detector for multiple types of filament networks, we processed dataset EMD-13671 with two

Fig. 5. Simulation of filament-like structures. Target occupancy of 10% for microtubules and 90% for the actin-like filament network. (a) Number of actin-like filament branches for different occupancies and P_b. (b) 3D visualization of a tomogram with 7% occupancy and $P_b = 0.1$. (c) 7% occupancy and $P_b = 0.9$. In (b-c) only the centerlines are shown for actin-like filaments to allow a better visualization of the generated crowded network. Actin-like filaments in blue, actin-like filament branches as red spheres and microtubules in green.

Fig. 6. Training a deep learning model. (a) A 2D slice in the XY-plane of an *in situ* tomogram taken from EMD-10439. (b) Segmentation performed by nnU-Net trained solely from the synthetic data, membranes are in yellow, microtubules in green, actin-like filaments in transparent blue and ribosomes in red. (c-d) Similar to (a-b) but the tomogram was provided by [50] authors. It is remarkable the performance of our model to recover membrane areas faded out by missing wedge when compared with segmentation provided in the Fig. 1 of [50]. Scale bars 200 nm.

models: one trained with all features as described in Section III-D (left values in Table I), and a second without the model

for actin-like filament networks but including a high resolution model (PDB-5MVY) of an straight actin filament fragment

(right values in Table I). The second model does not improve significantly the precision but presents a worse sensitivity. Moreover, it mostly fails to recover intermediate filaments (see Figure 8.a-c). The comparison of these models reinforces the necessity of the filament network model presented here to train a general model for processing filaments with different inner structures.

Table II contains the quantitative evaluation results obtained for the structures with reference segmentation in EMD-11992. The ribosome Dice score is around 0.5 with precision above 0.8. This is very interesting because the reference segmentation for ribosomes contains potential false positives depicted in Figure 8.a. Microtubules are almost perfectly segmented. Regarding actin, the situation is similar to EMD-13671, the sensitivity is high but the precision is low because the model segments many other filament-like structures besides actin filaments. Although EMD-11992 does not include a proper reference segmentation for membranes and membrane-bound proteins, Figure 8.d-f shows that our model performs well and is able to recover some top-view membranes and membrane-bound proteins, which are normally elusive because the lipid bilayer vanished due to the missing wedge.

IV. DISCUSSION

Although higher order structures such as organelles (mitochondrion, centriole, ...) or the nucleolus would require specific and very sophisticated models, the simulation of low-order structures proposed here can already be used to study many interesting cellular processes such as membrane morphology [23], [24], filament interaction under pathological conditions [22], macromolecule clusters [17], macromolecules in the liquid-like state [19] and membrane-bound protein nanodomains [18], [55]. In addition, local feature simulation has been shown a path to overcome the current limitations of DL approaches for processing (segmentation, macromolecular localization, ...) cryo-tomograms [32], [34]. However, validating tools analyzing experimental cryo-ET data is particularly challenging. For example, only few public datasets come with a segmentation, typically done by using several computer methods assisted by manual restoration. In practice, subtomogram averaging is used to verify the quality of particle positions: A resulting high-resolution structure indicates that many positions are correct. However, that does not guarantee

a low number in false positives or false negatives. Therefore, quantitative results require a detailed interpretation.

Interestingly, a DL model trained with synthetic data was able to predict some membrane regions faded out by missing wedge. This result opens a new approach to solve the missing wedge problem in cryo-ET. If we are able to reproduce enough contextual information in the synthetic data, then a DL model may recover the missing information. For the case of membranes, we plan to improve membrane models with the intention of generating more complex membrane arrangements. Additionally, our approach also allows to detect top-views for membrane bound proteins, something that is elusive to other specific methods [55].

Another interesting finding is that our model is not specific for a single type of filaments like actin. This model does not replicate specific actin features, but instead, is able to generalize the generic shape of a filament, in addition to its local geometry changes (curvature) and connections with other filaments. As a consequence, we have empirically observed that it can be used to train DL algorithms for parsing generic filament-like structures. In addition to actin, it detects many other filamentous structures such as intermediate filaments, as well as new structures not yet processed. We plan to further investigate how to include higher resolution features in filament structural models to obtain more specific detectors.

Reproducing the high concentration of macromolecules in the cellular cytoplasm or bounded to membranes is a major challenge. So far, only trial-and-error algorithm has been used for that purpose. Considering that we insert macromolecules in clusters (nanodomains), we have extended such approach by adding counters to re-start the insertion and avoid to get stuck in tomogram regions without enough empty space. This solution has allowed to generate synthetic tomograms with a very high concentrations of macromolecules, similarly to the cellular context. Our plan is to investigate the parallelization of the algorithms to speed up simulations.

We demonstrate that our synthetic data approximate well many of the features present in experimental *in situ* tomograms. Specifically, the information contained in the synthetic data is sufficient to train a generalized DL model to segment main features present in real cellular tomograms. Results show that segmenting observable features such as membranes and microtubules, as well as, large macromolecules like ribosomes or some membrane bound complexes does not require sub-nanometric information. Consequently, no external software was required to simulate an accurate model for CTF modulation, although PolNet is prepared to interact with external TEM simulators. Nevertheless, as future work, we plan to include accurate in-house CTF modulation for tomograms to reproduce high resolution (sub-nanometric) features. We foresee such information will become critical to reproduce accurately sub-nanometric features, and these may be necessary to recognize small structures by DL.

ACKNOWLEDGMENTS

We thank Prof Dr Juergen Plitzko at the Max Planck Institute of Biochemistry for sharing computational resources

Fig. 7. Membrane segmentation for EMD-13638. (a) Reference segmentation by MemBrain-v2 [54]. (b) Segmentation of our model, arrows depict top-view membranes partially recovered despite missing wedge.

needed for software development. We thank Tingying Peng from Helmholtz AI for mentoring L.L. We thank Jonathan Schneider for providing and segmenting tomogram in Fig 6.c.

REFERENCES

[1] M. Turk and W. Baumeister, "The promise and the challenges of cryo-electron tomography," *FEBS Letters*, vol. 594, no. 20, pp. 3243–3261, 2020.

[2] V. Lučić, A. Rigort, and W. Baumeister, "Cryo-electron tomography: The challenge of doing structural biology in situ," *The Journal of cell biology*, vol. 202, no. 3, pp. 407–419, 2013.

[3] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.

[4] Y.-T. Liu, H. Zhang, H. Wang, C.-L. Tao, G.-Q. Bi, and Z. H. Zhou, "Isotropic reconstruction for electron tomography with deep learning," *Nature Communications*, vol. 13, no. 1, pp. 6482–, 2022.

[5] T.-O. Buchholz, M. Jordan, G. Pigino, and F. Jug, "Cryo-care: Content-aware image restoration for cryo-transmission electron microscopy data," in *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, 2019, pp. 502–506.

[6] T. Bepler, K. Kelley, A. J. Noble, and B. Berger, "Topaz-denoise: general deep denoising models for cryoem and cryoet," *Nature Communications*, vol. 11, no. 1, pp. 5208–, 2020.

[7] M. Chen *et al.*, "Convolutional neural networks for automated annotation of cellular cryo-electron tomograms," *Nature Methods*, vol. 14, no. 10, pp. 983–985, 2017.

[8] I. de Teresa-Trueba *et al.*, "Convolutional networks for supervised mining of molecular patterns within cellular context," *Nature Methods*, vol. 20, no. 2, pp. 284–294, 2023.

[9] E. Moebel *et al.*, "Deep learning improves macromolecule identification in 3d cellular cryo-electron tomograms," *Nature Methods*, vol. 18, no. 11, pp. 1386–1394, 2021.

[10] T. Hrabe, Y. Chen, S. Pfeffer, L. Kuhn Cuellar, A.-V. Mangold, and F. Forster, "Pytom: A python-based toolbox for localization of macro-molecules in cryo-electron tomograms and subtomogram analysis," *Journal of Structural Biology*, vol. 178, no. 2, pp. 177–188, 2012, special Issue: Electron Tomography.

[11] S. Asano *et al.*, "A molecular census of 26s proteasomes in intact neurons," *Science*, vol. 347, no. 6220, pp. 439–442, 2015.

[12] Q. Guo *et al.*, "In situ structure of neuronal c9orf72 poly-ga aggregates reveals proteasome recruitment," *Cell*, vol. 172, no. 4, pp. 696–705, 2018.

[13] W. Dai *et al.*, "Visualizing individual rubisco and its assembly into carboxysomes in marine cyanobacteria by cryo-electron tomography," *Journal of Molecular Biology*, vol. 430, no. 21, pp. 4156–4167, 2018.

[14] S. Klein *et al.*, "Sars-cov-2 structure and replication characterized by in situ cryo-electron tomography," *Nature Communications*, vol. 11, no. 1, pp. 5885–, 2020.

[15] Y. Fukuda, F. Beck, J. M. Plitzko, and W. Baumeister, "In situ structural studies of tripeptidyl peptidase ii (tppii) reveal spatial association with proteasomes," *Proceedings of the National Academy of Sciences*, vol. 114, no. 17, pp. 4412–4417, 2017.

[16] A. Martinez-Sanchez, W. Baumeister, and V. Lučić, "Statistical spatial analysis for cryo-electron tomography," *Computer Methods and Programs in Biomedicine*, vol. 218, p. 106693, 2022.

[17] M. Jasnin *et al.*, "The architecture of traveling actin waves revealed by cryo-electron tomography," *Structure*, vol. 27, no. 8, pp. 1187–1189, 2019.

[18] A. Martinez-Sanchez *et al.*, "Trans-synaptic assemblies link synaptic vesicles and neuroreceptors," *Science Advances*, vol. 7, no. 10, p. eabe6204, 2021.

[19] E. Rosenzweig *et al.*, "The eukaryotic co 2-concentrating organelle is liquid-like and exhibits dynamic reorganization," *Cell*, vol. 171, no. 1, pp. 148–162, 2017.

[20] F. Brandt, L.-A. Carlson, F. U. Hartl, W. Baumeister, and K. Grünewald, "The three-dimensional organization of polyribosomes in intact human cells," *Molecular Cell*, vol. 39, no. 4, pp. 560–569, 2010.

[21] B. Turoňová *et al.*, "In situ structural analysis of sars-cov-2 spike reveals flexibility mediated by three hinges," *Science*, vol. 370, no. 6513, pp. 203–208, 2020.

[22] V. A. Trinkaus *et al.*, "In situ architecture of neuronal α-synuclein inclusions," *Nature Communications*, vol. 12, no. 1, pp. 2110–20, 2021.

[23] M. Salfer, J. F. Collado, W. Baumeister, R. Fernández-Busnadiego, and A. Martínez-Sánchez, "Reliable estimation of membrane curvature for cryo-electron tomography," *PLOS Computational Biology*, vol. 16, no. 8, pp. 1–29, 08 2020.

[24] B. A. Barad, M. Medina, D. Fuentes, R. L. Wiseman, and D. A. Grotjahn, "A surface morphometrics toolkit to quantify organellar membrane ultrastructure using cryo-electron tomography," *bioRxiv*, 2022.

[25] F. Förster, S. Pruggnaller, A. Seybert, and A. S. Frangakis, "Classification of cryo-electron sub-tomograms using constrained correlation," *Journal of structural biology*, vol. 161, no. 3, pp. 276–286, 2008.

[26] L. Pei, M. Xu, Z. Frazier, and F. Albert, "Simulating cryo electron tomograms of crowded cell cytoplasm for assessment of automated particle picking," *BMC Bioinformatics*, vol. 17, no. 1, p. 405, 2016.

[27] B. Rodríguez de Francisco, A. Bezault, X.-P. Xu, D. Hanein, and N. Volkmann, "Mepsi: A tool for simulating tomograms of membrane-embedded proteins," *Journal of Structural Biology*, vol. 214, no. 4, p. 107921, 2022.

[28] M. Vulović *et al.*, "Image formation modeling in cryo-electron microscopy," *Journal of Structural Biology*, vol. 183, no. 1, pp. 19–32, 2013.

[29] H. Rullgard, L.-G. Ofverstedt, S. Masich, B. Daneholt, and O. Oktem, "Simulation of transmission electron microscope images of biological specimens," *Journal of Microscopy*, vol. 243, no. 3, pp. 234–256, 2011.

[30] B. Himes and N. Grigorieff, "Cryo-TEM simulations of amorphous radiation-sensitive samples using multislice wave propagation," *IUCrJ*, vol. 8, no. 6, pp. 943–953, Nov 2021.

[31] I. Gubins *et al.*, "SHREC 2021: Classification in Cryo-electron Tomograms," in *Eurographics Workshop on 3D Object Retrieval*, S. Biasotti,

Fig. 8. Detailed segmentation evaluation. (a-b) Filamentous networks in a 2D slice of EMD-13671. (a) Manually supervised reference segmentation provided by [52], light blue are actin-like filaments and dark blue intermediate filaments. (b) Segmentation by our DL model trained with synthetic data including actin-like networks. (c) The synthetic training set includes instances of a rigid model for actin (PDB-5MVY), but not the actin-like networks presented here. Dark blue arrows point to intermediate filaments missed, especially in areas of high curvature. (d-f) EMD-11992. (d) Blue arrows point to filamentous structures not analyzed in (e), yellow arrows show membrane top-views where the membrane bilayer vanishes (see (f)), red arrows show potential false positives for ribosomes in (e), and orange arrows point to some membrane proteins including top-views (see (f)). (e) Reference segmentation provided in EMD-11992, red ribosomes, light blue actin filaments and microtubules in green. (f) Segmentation obtained with our model, membranes in yellow and membrane bound proteins in orange. Scale bars 200 nm.

R. M. Dyke, Y. Lai, P. L. Rosin, and R. C. Veltkamp, Eds. The Eurographics Association, 2021.

[32] G. Rice, T. Wagner, M. Stabrin, O. Sitsel, D. Prumbaum, and S. Raunser, "Tomotwin: generalized 3d localization of macromolecules in cryo-electron tomograms with structural data mining," *Nature Methods*, 2023.

[33] L. Zhou, C. Yang, W. Gao, T. Perciano, K. M. Davies, and N. K. Sauter, "A machine learning pipeline for membrane segmentation of cryo-electron tomograms," *Journal of Computational Science*, vol. 66, p. 101904, 2023.

[34] C. Purnell *et al.*, "Rapid synthesis of cryo-et data for training deep learning models," *bioRxiv*, 2023.

[35] X. Wu *et al.*, "Cryoetgan: Cryo-electron tomography image synthesis via unpaired image translation," *Frontiers in Physiology*, vol. 13, 2022.

[36] P. Harar, L. Herrmann, P. Grohs, and D. Haselbach, "Faket: Simulating cryo-electron tomograms with neural style transfer," 2023.

[37] L. Perez and J. Wang, "The effectiveness of data augmentation in image classification using deep learning," *CoRR*, vol. abs/1712.04621, 2017.

[38] E. Moebel and C. Kervrann, "Towards unsupervised classification of macromolecular complexes in cryo electron tomography: Challenges and opportunities," *Computer Methods and Programs in Biomedicine*, vol. 225, p. 107017, 2022.

[39] H. Nagashima and S. Asakura, "Dark-field light microscopic study of the flexibility of f-actin complexes," *Journal of Molecular Biology*, vol. 136, no. 2, pp. 169–182, 1980.

[40] H. Jiang, L. Jiang, J. D. Posner, and B. D. Vogt, "Atomistic-based continuum constitutive relation for microtubules: elastic modulus prediction," *Computational Mechanics*, vol. 42, pp. 607–618, 2008.

[41] A. Jegou and G. Romet-Lemonne, "The many implications of actin filament helicity," *Seminars in Cell & Developmental Biology*, vol. 102, pp. 65–72, 2020.

[42] Y. Pollak, S. Goldberg, and R. Amit, "Self-avoiding wormlike chain model for double-stranded-dna loop formation," *Phys. Rev. E*, vol. 90, p. 052602, Nov 2014.

[43] T. Wiegand and K. Moloney, "Rings, circles, and null-models for point pattern analysis in ecology," *Oikos*, vol. 104, no. 2, pp. 209–229, 2004.

[44] P. Andrey *et al.*, "Statistical analysis of 3D images detects regular spatial distributions of centromeres and chromocenters in animal and plant nuclei," *PLOS Computational Biology*, vol. 6, no. 7, p. e1000853, 2010.

[45] J. Kuffner, "Effective sampling and distance metrics for 3d rigid body path planning," in *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA '04. 2004*, vol. 4, 2004, pp. 3993–3998 Vol.4.

[46] K. Mitra, I. Ubarretxena-Belandia, T. Taguchi, G. Warren, and D. M. Engelman, "Modulation of the bilayer thickness of exocytic pathway membranes by membrane proteins rather than cholesterol," *Proceedings of the National Academy of Sciences*, vol. 101, no. 12, pp. 4083–4088, 2004.

[47] F. A. Heberle, D. Welsch, H. L. Scott, and M. N. Waxham, "Optimization of cryo-electron microscopy for quantitative analysis of lipid bilayers," *Biophysical Reports*, vol. 3, no. 1, p. 100090, 2023.

[48] K. J. R., D. N. Mastronarde, and J. R. McIntosh, "Computer visualization of three-dimensional image data using imod," *Journal of structural biology*, vol. 116, no. 1, pp. 71–76, 1996.

[49] F. Isensee, P. F. Jaeger, S. A. Kohl, J. Petersen, and K. H. Maier-Hein, "nnu-net: a self-configuring method for deep learning-based biomedical image segmentation," *Nature methods*, vol. 18, no. 2, pp. 203–211, 2021.

[50] J. Schneider and M. Jasnin, "Capturing actin assemblies in cells using in situ cryo-electron tomography," *European Journal of Cell Biology*, vol. 101, no. 3, p. 151224, 2022.

[51] J. Mahamid *et al.*, "Visualizing the molecular sociology at the hela cell nuclear periphery," *Science*, vol. 351, no. 6276, pp. 969–972, 2016.

[52] M. Jasnin *et al.*, "Elasticity of podosome actin networks produces nanonewton protrusive forces," *Nature communications*, vol. 13, no. 1, p. 3842, 2022.

[53] S. Shit *et al.*, "cldice-a novel topology-preserving loss function for tubular structure segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 16 560–16 569.

[54] L. Lamm *et al.*, "Membrain v2: an end-to-end tool for the analysis of membranes in cryo-electron tomography," *bioRxiv*, 2024.

[55] A. Martinez-Sanchez *et al.*, "Template-free detection and classification of membrane-bound complexes in cryo-electron tomograms," *Nature Methods*, vol. 17, no. 2, pp. 209–216, 2020.