

Deep Learning Improves Macromolecule Identification in 3D Cellular Cryo-Electron Tomograms

E. Moebel¹, A. Martinez-Sanchez^{2,3}, L. Lamm^{4,5}, R.D. Righetto⁴, W. Wietrzynski^{6,4}, S. Albert⁶,
D. Larivière⁷, E. Fourmentin⁷, S. Pfeffer⁸, J. Ortiz⁹, W. Baumeister⁶, T. Peng⁵,
B.D. Engel^{4,10}, C. Kervrann^{1,*}

¹ Serpico Project-Team, Centre Inria Rennes-Bretagne Atlantique and CNRS-UMR 144,
Inria, CNRS, Institut Curie, PSL Research University,
Campus Universitaire de Beaulieu, 35 042 Rennes Cedex, France

² Department of Computer Science, Faculty of Sciences, University of Oviedo, Calvo Sotelo s/n, 33071 Oviedo, Spain

³ Health Research Institute of Asturias (ISPA), Avenida Hospital Universitario s/n, 33011 Oviedo, Spain

⁴ Helmholtz Pioneer Campus, Helmholtz Zentrum München, Ingolstädter Landstraße 1, 85764 Neuherberg, Germany

⁵ Helmholtz AI, Helmholtz Zentrum München, Ingolstädter Landstraße 1, 85764 Neuherberg, Germany

⁶ Max Planck Institute of Biochemistry, Am Klopferspitz 18, 82152 Martinsried, Germany

⁷ Fourmentin-Guilbert Scientific Foundation, 2 Avenue du Pavé Neuf, 93 160 Noisy-le-Grand, France

⁸ Zentrum für Molekulare Biologie der Universität Heidelberg, Im Neuenheimer Feld 282, 69120 Heidelberg, Germany

⁹ Ernst Ruska-Centre, Wilhelm-Johnen-Straße, 52425 Jülich, Germany

¹⁰ Department of Chemistry, Technical University of Munich, Lichtenbergstraße 4, 85748 Garching, Germany

Abstract

Cryo-electron tomography (cryo-ET) visualizes the 3D spatial distribution of macromolecules at nanometer resolution inside native cells. While this label-free cryogenic imaging technology produces data containing rich structural information, automated identification of macromolecules inside cellular tomograms is challenged by noise and reconstruction artifacts, as well as the presence of many molecular species in the crowded volumes. Here, we present a computational procedure that uses artificial neural networks to simultaneously localize **multiple classes of macromolecules** in cellular cryo-electron tomograms. **Once trained, the inference stage of DeepFinder** is significantly faster than template matching, and performs better than other competitive deep learning methods at identifying macromolecules of various sizes in both synthetic and experimental datasets. On cellular cryo-ET data, DeepFinder localized membrane-bound and cytosolic ribosomes (~3.2 MDa), Rubisco (~540 kDa soluble complex), and photosystem II (~550 kDa membrane complex) with **an accuracy comparable** to expert-supervised ground truth annotations. Furthermore, we show that DeepFinder is flexible and can be combined with template matching to localize the missing macromolecules not found by one or the other method. The DeepFinder algorithm is therefore very promising for the **semi-automated** analysis of a wide range of molecular targets in cellular tomograms, including macromolecules with weights of 500-600 kDa and membrane proteins.

1

¹ * Corresponding author: charles.kervrann@inria.fr

Deep learning improves macromolecule identification in 3D cellular cryo-electron tomograms

E. Moebel, A. Martinez-Sanchez, L. Lamm, R.D. Righetto, W. Wietrzynski, S. Albert, D. Larivière, E. Fourmentin, S. Pfeffer, J. Ortiz, W. Baumeister, T. Peng, B.D. Engel, C. Kervrann

Cryo-electron tomography (cryo-ET) visualizes the 3D spatial distribution of macromolecules at nanometer resolution inside native cells. While this label-free cryogenic imaging technology produces data containing rich structural information, automated identification of macromolecules inside cellular tomograms is challenged by noise and reconstruction artifacts, as well as the presence of many molecular species in the crowded volumes. Here, we present a computational procedure that uses artificial neural networks to simultaneously localize **multiple classes of macromolecules** in cellular cryo-electron tomograms. **Once trained, the inference stage of DeepFinder is significantly faster than template matching, and performs better than other competitive deep learning methods at identifying macromolecules of various sizes in both synthetic and experimental datasets.** On cellular cryo-ET data, DeepFinder localized membrane-bound and cytosolic ribosomes (~ 3.2 MDa), Rubisco (~ 540 kDa soluble complex), and photosystem II (~ 550 kDa membrane complex) **with an accuracy comparable to expert-supervised ground truth annotations.** Furthermore, we show that DeepFinder is flexible and can be combined with template matching to localize the missing macromolecules not found by one or the other method. The DeepFinder algorithm is therefore very promising for the **semi-automated** analysis of a wide range of molecular targets in cellular tomograms, including macromolecules with weights of 500-600 kDa and membrane proteins.

Decades of research in cell biology have revealed that cellular processes are performed by groups of interacting macromolecules in a crowded environment. Deciphering the underlying mechanisms of these interactions is thus of paramount importance to gain a deeper understanding of the cell. Cryo-electron tomography (cryo-ET) can provide new insights into molecular interactions by producing 3D views of the native cellular interior at sufficient resolution to identify macromolecules. Unlike light microscopy, cryo-ET lacks specific markers, so the structures of the macromolecules themselves must be used for identification. However, cryo-ET has an advantage over light microscopy in that it visualizes everything in the cell, not just the tagged molecules. With enough resolution and powerful computational tools for structural identification, cryo-ET has the potential to build a complete molecular atlas of the cell.

Cryo-ET data is generated by the following steps. First, the samples are vitrified in order to preserve both the native structures and spatial distribution of macromolecules in the cell. For most cells, a thinning step is required, which can be accomplished with focused ion beam milling². Subsequently, the thin biological material is loaded into the transmission electron microscope for acquisition of a tilt-series, which serves as input to 3D reconstruction algorithms^{3;4;5}. During tilt-series acquisition, the specimen is rotated around an axis perpendicular to the electron beam and imaged from multiple perspectives. The tilt-series must be acquired with low electron dose because frozen biological material is easily damaged by the electron beam. Unlike the heavy metal contrasting agents used in conventional electron microscopy, the organic molecules found in frozen cells have low contrast against the water background. Combined with the limited electron dose used for imaging, cryo-ET data has a very low signal-to-noise ratio. Furthermore, due to the geometry of the EM grid and increasing thickness of the sample at high tilts, tilt-series are typically restricted to roughly ± 60 degrees. As a result, the reconstructed tomograms suffer from a **wedge of missing information** in Fourier space. This missing wedge artifact causes anisotropic resolution in the 3D volumes, with delocalization of densities along the Z-direction, as well as loss of information in the X-Y plane along the direction perpendicular to the tilt axis^{6;7}. The missing wedge and low signal-to-noise ratio, combined with the highly crowded environment of the cell, pose significant challenges to the identification of macromolecules in cellular tomograms.

One well-established method for localizing macromolecules in cryo-ET data is template matching⁸, where a low-resolution template depicting the macromolecule of interest is comprehensively scanned through the tomogram. High cross-correlation scores indicate potential particle positions, from which subvolumes are extracted for downstream averaging procedures. While template matching is relatively efficient for localizing large complexes such as ribosomes, it is necessary to apply a series of iterative searching, alignment, and classification steps to identify smaller molecules⁹. Additional difficulties arise when template matching is used to localize several macromolecular species that are structurally similar, or to differentiate specific states of the same macromolecular species (e.g., membrane-bound vs. cytosolic ribosomes). Template matching is applied to separately localize all macromolecules of a single species (mono-class procedure). Nevertheless, dedicated classification steps^{10;11} are required to differentiate true particles from false positives, as well as to subdivide these particles into structurally distinct sub-classes. Classification remains a challenging problem in cryo-ET if the number of considered classes or subclasses is high. Currently, complex and time-consuming processing chains are routinely used to localize macromolecules in cellular volumes, with a single template matching or subtomogram classification round typically taking 10-30 hours of computation on specialized CPU clusters. As a result, this whole procedure is most often used to analyze only a few classes or subclasses of particles in the same volume. However, cryo-electron tomograms contain many more macromolecular species embedded within the crowded cellular environment, hidden by noise and reconstruction artifacts. To address this challenging issue, powerful new pattern recognition techniques are required.

In this paper, we describe a deep learning-based framework to quickly identify macromolecules in cryo-electron tomograms. Deep learning¹² is revolutionizing various fields of data processing, including computer vision¹³, language processing¹⁴, image classification¹⁵ and segmentation¹⁶, and object recognition¹⁷. In bioimage analysis, convolutional neural networks (CNN) have produced spectacular results^{18;19;20;21;22}, including in super-resolution microscopy²³ and in light microscopy image denoising²⁴. Briefly, a CNN is defined as an architecture composed

of successive connected neuron layers. The neurons are small processing units capable of performing linear and non-linear operations. Each neuron is controlled by parameters that are optimized during the learning process. In the case of CNNs, the neurons are applied in a convolutive manner. Thus, a neuron can be thought of as a filter, and a neuron layer as a filter bank. Each layer automatically extracts various and specific features from the data. Applying the layers sequentially enables the network to progressively compute high-level features, which results in a hierarchical or multiscale representation of the data. In human face recognition, the first layers typically encode basic features such as image contours/edges and textures, which allows the next layers to gradually capture more complex shapes (e.g., eyes, ears), and object ensembles (e.g., faces). These multiscale representations are learned from the data, and can be generated faster than conventional handcrafted representations.

CNNs have recently been investigated for learning high-level generic features in electron microscopy. Several algorithms based on deep learning techniques have been developed for 2D particle picking in single-particle cryo-EM²⁵, including DeepPicker²⁶, AutoCryoPicker²⁷, crYOLO²⁸, Topaz²⁹, and Warp³⁰. In cellular cryo-ET, the algorithm proposed by Chen *et al.*, was implemented for supervised segmentation of 2D slices from 3D volumes³¹, but it was not designed to handle complex environments (e.g., crowded cell) and requires an additional classification step to achieve satisfactory results. A promising 3D processing approach for the supervised classification of subtomograms has been proposed in Che *et al.*³². However, this algorithm, which we have evaluated on our dataset, requires that the molecular complexes of interest have already been found by another method. To overcome the limitations of the aforementioned approaches^{26;31;32}, we propose a 3D deep learning method to identify macromolecules within native crowded cellular environments. Our DeepFinder algorithm can handle multiple macromolecular species in one pass, which, as we demonstrate, improves the performance of CNNs in 3D cryo-ET, identifying small particles that template matching struggles to detect. Moreover, complex and time-consuming post-classification steps are no longer required to produce reliable results. We also show that DeepFinder is very flexible, and can be efficiently combined with template matching to improve localization sensitivity on crowded cellular cryo-tomograms.

Results

Overview of our 3D deep learning-based approach. We present DeepFinder, a flexible algorithm based on 3D convolutional neural networks that in one pass can [robustly](#) localize macromolecules of several different species, with various sizes and shapes, within cryo-tomograms. DeepFinder is built upon a multi-class network (see Methods and Supplementary Fig. 7). The algorithm is based on a U-net architecture^{16;20}, as illustrated in Supplementary Fig. 8, and consists of a training stage (Stage I) and an analysis (or inference) stage (Stage II) (see details in Methods and Fig. 1A). The analysis stage is a two-step procedure. In the first step, a supervised 3D CNN-based method is used to classify the 3D tomogram voxels into different categories or sub-categories of macromolecules. In the second step, a clustering algorithm is applied to aggregate voxels into groups and determine the location of particles (gravity center of clusters) in the volume (Fig. 1B). The particles are further used for structure determination through subtomogram averaging. [DeepFinder's Stage I is a supervised approach requiring expert-user inputs, whereas Stage II is nearly unsupervised, as opposed to template matching, which requires more manual input.](#) DeepFinder is a free, open-source program implemented in Python with an accessible graphical user interface (see Supplementary Fig. 6 for illustration).

In the following sections, we demonstrate that DeepFinder successfully identifies macromolecules of varying shape and molecular weight in artificial and cellular cryo-ET datasets. We show that handling several molecular species simultaneously (i.e., multi-class strategy) is especially beneficial for identifying small macromolecules.

DeepFinder converts the 3D positions of annotated macromolecules into segmentation targets for training. In the training stage (Stage I), DeepFinder converts the input 3D coordinates of macromolecules supplied by the experts into voxel-wise annotations (see Methods), avoiding the cumbersome and time-consuming manual annotation of voxels. As converting the positions of macromolecules into voxel-wise annotations is not a trivial task, we propose two approaches with different levels of computing complexity. The first approach exploits the segmentation maps built from well-delineated shapes of macromolecules estimated by a subtomogram averaging procedure⁴¹. In the second approach, the shapes of macromolecules are approximated by 3D spheres (see Fig. 2). The sphere-based representation is appealing since neither information about the shape nor the orientation of the macromolecules are needed, and time-consuming subtomogram averaging steps are avoided. However, using 3D spheres instead of 3D shapes induces more "label noise", which hinders optimization of the network.

In our experiments on the SHREC'19 dataset, the sphere-based representation was able to provide good scores for large and medium sized, as shown in Fig. 3C. For small and tiny macromolecules, it even provided partly better results than the shape-based representation (except for the smallest, 1s3x). [The explanation for this good performance is two-fold: on the one hand, small macromolecules tend to resemble sphere-like structures. On the other hand, the "label noise" induced by the sphere-based representation is compensated by the high quantity of labeled voxels in the segmentation map \(see Rolnick *et al.*⁴² for discussion\).](#) In summary, the computationally cheap sphere-based representation is recommended, provided that the target macromolecule shape is not too topologically different from a sphere (e.g. tube-like structure).

DeepFinder identifies multiple macromolecular species at once. We found that a multi-class strategy is more beneficial than identifying macromolecule classes separately. To demonstrate this point, we evaluated the performance of our multi-class network versus 12 mono-class versions of our network on the SHREC'19 dataset (Dataset #1). The SHREC'19 dataset is composed of 10 synthetic tomograms, each containing 12 different classes of macromolecules, with different sizes and structures³³. These macromolecules have been categorized by size into 4 groups (large, medium, small, tiny) by the organizers. The number of macromolecules per class and per tomogram varies slightly (~ 200 particles per class). As shown in Fig. 3B, the binary networks have a similar performance when compared to multi-class networks for larger macromolecules (e.g., the following Protein Data Bank (PDB) entry identifiers: 4b4t, 4d8q), but perform worse (or fail) for the smaller ones (e.g., 1s3x, 3gl1, 3h84 and 3qm1). Overall, the multi-class training improves classification accuracy over binary classification. These results also suggest that smaller macromolecules and under-represented classes can benefit from features learned from classes of larger macromolecules.

In the SHREC'19 challenge, seven different methods (including DeepFinder) were compared to evaluate the performance of deep learning techniques on challenging synthetic 3D cryo-electron tomograms. Participants had access to 9 out of 10 ground truth tomograms, whereafter





Figure 1: Overview of DeepFinder. (A) The DeepFinder workflow consists of training stage (Stage I) and an analysis (or inference) stage (Stage II). These two stages correspond to five steps (represented by blue boxes) to locate macromolecular complexes within crowded cells. The DeepFinder package provides a graphical interface for each step to increase the accessibility of the algorithm. Alternatively, each step can be executed with scripts using the API. (B) Ribosome localization with DeepFinder in a cryo-electron tomogram of a *Chlamydomonas reinhardtii* cell. Tomographic slice with superimposed segmented cell membrane (gray) and ribosomes classified with respect to their binding states: membrane-bound (blue) and cytosolic (yellow). (C) Tomographic slices showing coordinates of detected ribosomes (colors correspond to (B)). The positions and classes were determined by analyzing the segmentation map shown in (B).

the algorithms were tested on the final unseen tomogram. DeepFinder achieved the best global F_1 -score (see evaluation in Methods) among all competing methods³³. The closest competitor to DeepFinder was the DoG-CB3D method³², which separates the localization and classification tasks. In Fig. 3A, we report the results obtained on the SHREC’19 dataset with DeepFinder, template matching, and DoG-CB3D. **Unlike DoG-CB3D, DeepFinder outperformed template matching for all macromolecular species in the benchmark dataset. The template matching results were obtained with the PyTom software³⁴ described in Appendix A.** Performance is correlated with macromolecular size (Fig. 3A). This is especially true for template matching, where the smallest complexes had F_1 -scores close to 0. Even so, template matching is still a competitive method for detecting the largest macromolecules. Meanwhile, the organizers provided the results obtained with template matching, as described in Appendix A.

On the SHREC’20 dataset³⁵, DeepFinder was ranked second, with essentially the same algorithm from the previous year. It is worth noting that the top algorithm (U-net Multi-task Cascade (UMC)) requires noise-free images for training³⁵. This approach is therefore not applicable on real cellular datasets, although there has been significant progress on using deep learning to denoise tomograms^{29;30;36;37}.

More data annotation is required to localize small macromolecules. We analyzed the influence of the training set size on the performance of our method (see Supplementary Figs. 10-11). It is desirable to achieve good performance even if the training set size is small, as manual annotation of data is time consuming and requires considerable effort in 3D imaging. As shown in Supplementary Fig. 11 (synthetic SHREC’19 dataset), the performance of DeepFinder correlates with macromolecule size. For large and medium macromolecules, the performance is remarkably stable, even when using only one training tomogram (i.e., 206 macromolecules per class). For small and tiny macromolecules, the drop in performance is more significant, and larger training sets are required to get higher score values. This is not surprising since the number

Strategy 1: shape-based targets



Strategy 2: sphere-based targets




Figure 2: Target generation strategies for training. DeepFinder includes a graphical user interface (see Supplementary Fig. 6) to guide users throughout the learning procedure, from selecting the 3D coordinates of macromolecules to voxel-wise annotation. **Strategy 1 (left):** voxel-wise annotations are obtained from position-wise annotations using a subtomogram averaging procedure. Subtomogram averaging is a registration algorithm that produces higher resolution structures by averaging thousands of aligned subvolumes containing the molecular species. Subtomogram volumes are extracted at the annotated positions, aligned and finally averaged. The alignment procedure outputs the object orientations, while the averaging process provides a density map of the macromolecule with greatly reduced noise and missing wedge artifacts. From this density map, it is possible to create a binary mask of the macromolecule by thresholding the averaged subtomogram. The resulting mask is pasted into an empty volume at each annotated position with the estimated 3D orientation, and the resulting volume with well-delineated macromolecules is then used as a target to train the 3D CNN parameters. **Strategy 2 (right):** the macromolecule masks are approximated and replaced by sphere masks. Hence, the subtomogram averaging procedure is bypassed and the target generation process is much faster. However, the training targets contain more "label noise".

of labeled voxels associated with large macromolecules is high in the segmentation maps. Accordingly, less annotated macromolecules are needed.

To carefully examine the performance of DeepFinder on cellular cryo-tomograms, we focused a range of protein complexes: ribosomes, rubisco and PSII. These complexes correspond to different particle sizes and quantities of annotation. Ribosomes (large) and rubisco complexes (small) are relatively well localized with template matching, yielding abundant annotations for training and allowing a direct comparison of DeepFinder with the template matching results. Detecting PSII is a more challenging task, given that in addition to being small complexes, the quantity of annotations is limited. The calculation of F_1 -scores usually relies on the availability of ground truth positions, as is the case on synthetic datasets (SHREC'19 and SHREC'20), which are generally not available with experimental data. Instead of ground-truth, the F_1 scores reported below are calculated with respect to the annotations provided by the experts which use a combination of template matching, visual curation, subtomogram averaging, and classification. An F_1 -score of 1 indicates that the performance of DeepFinder is the same as the approach supervised by the experts.

DeepFinder imitates expert annotations on cellular tomograms depicting large macromolecules. First, we explored the ability of our method to localize two ribosomal states ('membrane-bound ribosomes' and 'cytosolic ribosomes', denoted *mb-ribo* and *ct-ribo*, respectively), and we compared the results to those obtained with template matching (see Appendix A). In this study, we used a cryo-ET dataset (Dataset #2, see Methods) composed of 57 tomograms of *Chlamydomonas reinhardtii* cells and annotations of membrane-bound 80S ribosomes (~3.2 MDa) performed by experts³⁹. In our experiments, we considered four classes to better detect the target macromolecules: membrane-bound ribosome (*mb-ribo*), cytosolic ribosome (*ct-ribo*), membrane, and background. The expert annotations contained positions of 8,795 *mb-ribos*. The training datasets corresponding to the *ct-ribo* and membrane classes were obtained by using semi-automated computational tools, without careful expert supervision (see Methods). The classes were highly imbalanced since the background class represents 99.5% of training voxels within the tomogram. As this scenario is common in cryo-ET, the training procedure (especially the generation of batches) needs to be adapted by integrating techniques such as re-sampling (see Methods). Figure 1B-C illustrates an example of our four-class segmentation and provides insight into the spatial distribution of two ribosomal states. DeepFinder detected *mb-ribos* that were not annotated by the experts. Several of these particles were false positives, but after a detailed inspection (see Supplementary Fig. 14), we confirmed that many others are ribosomes that were missed or discarded during the annotation process.

Analysis of score-maps. In Fig. 4, it is clear that the score-map of template matching (Fig. 4D) is much noisier than the score-maps generated by DeepFinder (Fig. 4B-C). Template matching produced high cross-correlation scores at ribosome locations but also at false-positive locations containing other high contrast structures (e.g., cell membrane in Fig. 4D). Consequently, template matching is mainly used to crudely discard




Figure 3: Analysis of algorithm performance on the synthetic dataset (SHREC'19 challenge). (A) Performance (F_1 -score) of DeepFinder, DoG-CB3D and template matching algorithms and ability to discriminate between 12 classes/subclasses of macromolecules. The tiny macromolecules are 1s3x (human Hsp70 ATPase domain), 3qm1 (Lactobacillus Johnsoni Cinnsmoyl esterase LJ0536 S106A mutant in complex Ethylferulate Form II), and 3gl1 (Saccharomyces cerevisiae ATPase domain of Ssb1 chaperone). The highest possible value of an F_1 -score is 1.0 and the lowest possible value is 0. The scores of template matching were provided by the SHREC'19 challenge organizers (Utrecht University, Department of Information and Computing Sciences and Department of Chemistry). (B) Performance of DeepFinder implemented as a multi-class network architecture and as an architecture made of 12 binary networks. These two architectures differ only by the number of output neurons. (C) Influence of the training target generation method ("shapes" versus "spheres"). In the case of "shapes", the exact shapes of the macromolecules have been used to annotate the tomograms. In the case of spheres", the shape and the orientation of macromolecules are not needed to generate the training targets.




Figure 4: Comparison of score-maps obtained with template matching and DeepFinder and analysis of structural resolution through subtomogram averaging (Dataset #2). (A) Experimental cryo-electron tomogram depicting a *Chlamydomonas reinhardtii* cell (Dataset #2, see Methods). (B,C) Localizations of *mb-ribos* and *ct-ribo* particles with DeepFinder, and histograms of score-map local maxima. (D) Localizations of ribosomes with template matching (TM), and histogram of score-map local maxima. (E) Gold-standard FSC curves in each case with estimated resolutions. The resolution (24Å) corresponding to expert annotations (dedicated to *mb-ribo* particles) is very comparable to the resolution (24.7 Å) obtained with *mb-ribo* particles localized with DeepFinder. (F) Subtomograms obtained from expert annotations (left) and from particles localized with DeepFinder (*mb-ribos* (middle) and *ct-ribos* (right)). The average densities have been obtained with the same alignment procedure and parameter settings. For comparative visualization, the averages have all been low-pass filtered to 40Å resolution.

the voxels with no significant structural information. In the next step, the experts must apply post-processing techniques to separate the desired macromolecules from the false positives. Unlike template matching, DeepFinder produces less noisy score-maps, with meaningful score values in well-localized blobs.

Furthermore, DeepFinder allows accurate localization of two ribosomal states in cellular cryo-ET data, as illustrated in Supplementary Fig. 6A. The *mb-ribos* (in blue) are primarily located close to cell membranes, whereas the *ct-ribos* (in yellow) are mainly located in the cytosol. This is confirmed by analyzing the Euclidean distances of *mb-ribo* and *ct-ribo* particles to the nearest cell membranes (see Supplementary Fig. 12).

Analysis of consensus responses. For the next step of our evaluation, we compared the template matching and DeepFinder procedures by using the expert annotations as ground truth. In Supplementary Fig. 13, we plotted the *Recall*, *Precision* and *F₁*-scores with respect to DeepFinder and template matching algorithm parameters. We focused on the thresholds used to detect objects with template matching (thresholding of score values) and DeepFinder (thresholding of object sizes). We obtained a maximum *F₁*-score of 0.86 for DeepFinder and an *F₁*-score of 0.50 for template matching (with no post-classification step, see Supplementary Fig. 7A) (Dataset #2). Furthermore, we also examined the complementarity between the two sets of *mb-ribo* macromolecules found by the experts and DeepFinder. In the following analysis, we denote the sets obtained by experts and DeepFinder as S_E and S_{DF} , respectively. While the overlap $S_E \cap S_{DF}$ between both sets is substantial (1,516 particles), there is also a significant number of particles belonging to $S_E \setminus S_{DF}$ (220 particles), i.e., the particles annotated by the expert but not found by DeepFinder, and to $S_{DF} \setminus S_E$ (356 particles), i.e., particles found by DeepFinder but missed by the expert. We can benefit from the two complementary sets of particle positions to improve the overall validation rates. The union $S_E \cup S_{DF}$ of the two sets increases the list of potential *mb-ribo* macromolecules, for which a confidence level can be assigned to each set member depending on whether it belongs to $S_E \cap S_{DF}$, $S_{DF} \setminus S_E$ or $S_E \setminus S_{DF}$. The particles belonging to $S_E \cap S_{DF}$, i.e., found by both methods, are very likely to be true positives. Meanwhile the particles belonging to $S_E \setminus S_{DF}$ and $S_{DF} \setminus S_E$ can be labeled as “suspicious” and require more investigation. These

two non-union sets are relatively small, enabling assignment of the bulk high-confidence particles so the expert can focus on validating the remaining low-confidence particles. In this manner, it is possible to uncover inaccuracies in the expert annotations and refine the true-positive particle class, which can further improve the training performance of DeepFinder.

Analysis of structural resolution. DeepFinder achieves 3D structural resolution comparable to the expert annotations, as determined through subtomogram averaging. We analyzed the particles by computing subtomogram averages for each ribosomal state (*ct-ribos* and *mb-ribos*) (Fig. 4F). The input sub-volumes which served to compute subtomogram averages are extracted from tomograms taken in the testing set. We applied the same registration procedure (rotational matching⁴³) before computing the averages. The resulting density maps of *mb-ribos* are composed of two parts: the ribosome density and the membrane density. As expected, the membrane density computed from the set S_{DF} of *mb-ribo* particles is well defined, while there is no membrane density in the average computed from the set of *ct-ribo* particles. We noticed that the ribosome density is similar in the two ribosomal states, with a difference at the interface with the membrane.

Furthermore, we calculated the Fourier shell correlation (FSC) from the subtomogram averages for each ribosomal state (Fig. 4E). For *mb-ribos*, the resolution was comparable between the expert annotations and DeepFinder (24Å and 25Å, respectively). The resolution was lower for the *ct-ribo* average (33Å) presumably because the *ct-ribo* annotations used for training, were not accurately curated by the experts. It is worth noting here that, although it is 8Å lower in resolution, the DeepFinder *ct-ribo* average contains 23% less particles than the DeepFinder *mb-ribo* average (1,394 vs. 1,805 particles). Additionally, the two particle classes are not precisely of the same structure; it is not well established in the literature how the presence of a membrane density influences the calculated resolution. This being said, the goal of this exercise was not to produce the highest resolution average (which could require >10,000 particles), but rather to test how well DeepFinder can distinguish subpopulations, even when particle number is limited.

Additionally, we questioned the nature of DeepFinder detections that are not part of the annotations (i.e. $S_{DF} \setminus S_E$). We therefore carefully analyzed the difference between the sets S_{DF} and S_E that were used to compute the subtomogram averages. The study amounts to considering the two following hypotheses:

H_1 : The set S_{DF} contains too many false positives.

H_2 : The set contains particles with a low quality (due to noise and blur), suggesting that our method found supplementary true *mb-ribo* particles that were missed or discarded during the supervised annotation.

To test these hypotheses, we aligned and averaged *mb-ribo* particles in the set $S_{DF} \setminus S_E$, i.e., found by DeepFinder but not present in the set of expert annotations. We observed that the corresponding density map depicts a ribosome structure (see Supplementary Fig. 14). Thus, hypothesis H_1 is unlikely. As shown in Supplementary Fig. 15, DeepFinder finds additional membrane-bound ribosomes, missed or discarded by experts (hypothesis H_2). Therefore, it appears that the number of actual *mb-ribos* is higher than expected: in our test set, we detected +20.5% *mb-ribos* when compared to the S_E set. This result confirms the benefit of combining several analysis methods in cryo-EM⁴⁴. The consensus between the expert processing chain and DeepFinder decreases the overall false negative rate and generates a larger set $S_E \cup S_{DF}$ of membrane-bound ribosomes. The extended set of particles is probably less homogeneous, as it contains supplementary particles which are noisier. Adding this set of particles in the average may blur the result and therefore degrade the structural resolution. Nevertheless, these particles not found by template matching should not be discarded *a priori*, as they provide a more complete picture of the cellular environment. In summary, DeepFinder, as well as human experts (who used template matching followed by CPCA classification with visual inspection), do miss some true positives. Our results illustrate the collaborative strength of the two particle picking procedures which are not perfect. Even though we demonstrated on the SHREC 2019 and SHREC 2020 datasets that DeepFinder has better precision, recall and F_1 -score than template matching, the global performance can be boosted with consensus analysis of competing methods⁴⁴.

DeepFinder robustly finds small macromolecules in cellular tomograms. On cellular cryo-electron tomograms, the F_1 -score is not significantly worse if we use a *mb-ribo* training dataset consisting of one-fifth of the complete annotated dataset (8,795 *mb-ribos*) (see Fig. 11). Together, these results suggest that DeepFinder does not require a large quantity of annotations for localizing large macromolecules, but more annotations are necessary to localize small macromolecular species. This is exemplified below on smaller macromolecules, including both soluble and membrane proteins.

To demonstrate this point, we conducted experiments on two challenging datasets depicting pyrenoids (Dataset #3) (Fig. 5A-C) and thylakoid membranes (Dataset #4) (Fig. 5E-H) within *Chlamydomonas reinhardtii* cells^{38,40}. In Dataset #3, DeepFinder was able to identify Rubisco holoenzymes with remarkable performance (Fig. 5A-B). The achieved F_1 -score (0.83) is similar to what was obtained for *mb-ribos* (0.86) (Dataset #2), even though the molecular mass of Rubisco (~540 kDa) is much lower than that of the 80S ribosome (~3.2 MDa). Subtomogram averaging performed with the top 30,000 DeepFinder hits yielded a map with resolution of 15 Å (see Methods). Template matching picks with CPCA classification post-processing produced a nearly-identical map using the same alignment and averaging pipeline, but required more computing time and tedious classification, as described previously¹⁰. DeepFinder inference is faster and yielded a "clean" set of particles without any additional post-processing. The two averages are also practically identical to the previously published map at 16.5 Å (see EMD-3694). This result suggests that DeepFinder is able to identify small macromolecular species with similar performance to that seen for larger complexes, provided that the training set is large enough.

In our experiments on Dataset #4, DeepFinder was able to detect Photosystem II (PSII) dimers embedded within the thylakoid membranes (Fig. 5E-H). We trained DeepFinder on three tomograms using four particle classes that were manually annotated³⁸: PSII dimers (~550 kDa), cytochrome b6f dimers (~220 kDa), unknown densities, and background. In this study, A very important distinction is that, while the Rubisco annotations were generated by supervised template matching and classification, the PSII ground truth was assigned completely manually through "membranogram" visual inspection³⁸ (see Fig. 5G) and served further as a reference to compute the F_1 scores. In Supplementary Table 1, we report the F_1 -scores for template matching and DeepFinder (when using the mono-class and multi-class strategies) to detect PSII in the unseen test tomogram; the multi-class approach produced better results than the mono-class approach and template matching. PSII complexes were detected quite accurately in some membranes (as illustrated in Fig. 5F-G), which was confirmed by the high F_1 -score (0.737 for membrane 4 in Supplementary Table 1), whereas in other membranes they were not. Note that the number of particles per membrane used




Figure 5: DeepFinder localizes small macromolecules in cellular tomograms. (A) Segmentation of a tomogram of the native *Chlamydomonas pyrenoid* (Dataset #3). Rubisco holoenzymes segmented by DeepFinder are displayed in transparent orange. For visualization purposes, hand-segmented pyrenoid tubule membranes (gray and yellow) have been added. (B) Tomogram slice with corresponding DeepFinder segmentation and **template matching** (post-processed with CPCA classification¹⁰). (C) Rubisco averages (top and side views) obtained with the PyTOM software³⁴ from the top 30,000 DeepFinder and template matching particles, with PDB 7JN4 fitted into the volumes. (D) Fourier shell correlation (FSC) curves from the averages obtained with 30,000 particles each. Both methods yield a final resolution of 15 Å (FSC > 0.143) (see Methods). (E) Slice through a tomogram of the chloroplast within an intact *Chlamydomonas* cell (Dataset #4). (F) Zoomed-in view of the tomogram slice, with the associated score-maps produced by DeepFinder and template matching. (G) The membranogram is a visualization approach where densities from the tomogram are projected onto the surface of a segmented thylakoid membrane, resulting in a topological view of the membrane surface³⁸. Below are the associated DeepFinder and template matching score-maps and a visualization comparing the positions of particles found by the expert (blue) versus DeepFinder (yellow) and template matching (pink), respectively. (G) Zoomed-in view of the tomogram slice, with the associated score-maps produced by DeepFinder and template matching. (H) Subtomogram average (white) of photosystem II complexes, calculated from 246 particles identified by DeepFinder on a single test tomogram, fit with a PSII molecular structure (blue).

to compute scores is not the same in Supplementary Table 1. Overall, the mono-class approach with only PSII performed worse, in particular with respect to Recall, than the multi-class setting (0.505 versus 0.638 in Supplementary Table 1). **Template matching yielded lower values (F_1 scores, Recall, Precision, see Table 1) than those obtained with the mono-class and multi-class DeepFinder strategies.** In contrast to the specificity of DeepFinder, template matching produced high cross-correlation scores along most of the thylakoid membranes and struggled to distinguish PSII particles from the surrounding membrane density (Fig. 5F-G). The results obtained with DeepFinder are promising but limited due to the small training dataset size (298 PSII particles from 18 membranes) as well as the variable quality of the experimental data. Membranes where PSII complexes were poorly resolved had the most missed picks and lowest F_1 -scores. Nevertheless, **the results illustrated in Fig. 5E-H are encouraging**, and we are convinced that adding more particles to the training sets (as exemplified in the Dataset #3 experiment) will further improve the capability of DeepFinder to reliably detect PSII and other membrane proteins.

Collectively, these experiments on soluble and membrane complexes of different sizes consistently indicate that the multi-class approach is a flexible and robust method to identify target macromolecules. Moreover, spurious components such as gold beads or ice chunks can be typically included in the negative class (i.e., background) during the training task in order to reduce false positives, as illustrated in Supplementary Fig. 9.

Discussion

Currently, only a handful of different molecular species have been analyzed in cellular cryo-tomograms. These volumes contain rich information on the whole proteome, but many macromolecules are hardly discernable from noise and artifacts. To realize the potential of visual proteomics⁴⁹ and study the interaction between several macromolecular species, more robust pattern recognition techniques are required. We believe that deep learning offers a powerful framework to address this challenge. **Therefore, we have developed DeepFinder** to efficiently identify molecular complexes with variable shapes and molecular weights. We showed quantitatively that we can obtain a structural resolution with this deep

learning-based method that is similar to that obtained by expert curation, for example as demonstrated for ~540 kDa Rubisco particles in Fig. 5. Moreover, once trained, DeepFinder is relatively fast compared to the common template matching and subtomogram classification pipeline, as demonstrated in Supplementary Table 2. Larger datasets can be processed in one day with DeepFinder, while several macromolecular species are simultaneously identified. Furthermore, we provide an open-source implementation of DeepFinder, with a graphical user for both training and inference (see Supplementary Fig. 6), aimed to be accessible for a broad base of users.

On the evaluation side, in the case of experimental data, annotations are a sub-set of the unknown ground truth. This is why no existing algorithm produces an F_1 -score of 1 with respect to the expert annotations. Indeed, most of the expert annotations can be considered as true positives, but there will almost certainly be a percentage of false negatives in the labeled dataset. An F_1 -score of 1 during training would indicate that the network is rather overfitting to the expert's annotations, instead of really learning the structures of particles. Therefore, for a properly working detection algorithm, an F_1 -Score of 1 is not to be expected. In particular, a low precision value does not necessarily mean that the picked particles are false positives.

To cope with the missing wedge effect (i.e., delocalization along Z-direction), the usual template matching procedure uses an isotropic template, along with missing wedge-constrained cross-correlation as a criterion. In contrast, DeepFinder is trained to predict a segmentation mask depicting particle shapes which are not affected by the missing wedge (see Fig. 2). Later on, DeepFinder applies a series of linear (convolution) and non-linear operations (max-pooling, ReLU activation). All these cascaded operations are performed in real space and progressively modify the results in Fourier space, thereby reducing the influence of the missing wedge. A way to interpret the cascade mechanism is to consider the simple case of generating a binary mask from an image or volume. The thresholding applied for binarization is a nonlinear operation. The resulting binary mask will have its corresponding Fourier representation completely filled with information, even if the input data has a missing wedge. Meanwhile, applying the segmentation mask to the input data amounts to multiplying them in real space, and is equivalent to applying a convolution operation in Fourier space, hence partially filling the missing wedge. This principle has been used to fill the missing wedge in cryo-ET and electron crystallography^{46;47;48}. DeepFinder generates an output like this, except it can have multiple discrete objects in the segmentation. Finally, even if some features are missing or distorted due to the missing wedge, particles can still be recognized based on the remaining intact features of the particles.

It is worth noting that annotating the training set requires some time. However, when processing multiple tomograms (as is normally the case in cryo-ET), the benefits from the much faster inference by DeepFinder outweigh the annotation time. Importantly, the goal is to achieve generalizability to all kinds of tomograms, which already works to some extent, as we demonstrate with our network trained on the *Chlamydomonas* (algae) data and applied to previously unseen mouse cell tomograms (see an example in Supplementary Fig. 16). The generalizability will only get better as new training data is added, decreasing the amount of necessary annotation for new analysis and making DeepFinder even more efficient. In contrast, the long computational time of template matching remains fixed. As it is not possible to precisely quantify the time of the annotation task, which depends on human performance and expertise, we focused on algorithm run-time with given annotations, as is common in machine learning (see Table 2).

In conclusion, we believe that DeepFinder is a promising algorithm for cryo-ET image analysis, which will contribute to developing visual proteomics in the forthcoming years. While unsupervised learning holds great promise for the future, the current price to be paid for better results is to resort to supervised machine learning algorithms guided by time-consuming expert annotations; PDB (Protein Data Bank) and EMDB (Electron Microscopy Data Bank) databases are currently analyzed to help expand the range of molecular species that can be identified in cellular tomograms. Future challenges include developing unsupervised or self-supervised (reference) deep learning methods in order to decrease the amount of required training data.

References

- [1] Vendeville, A., Larivière, D. & Fourmentin, E. An inventory of the bacterial macromolecular components and their spatial organization. *FEMS Microbiology Reviews* **35**, 395–414 (2011).
- [2] Schaffer, M. et al. Optimized cryo-focused ion beam sample preparation aimed at in situ structural studies of membrane proteins. *J. Struct. Biol.* **197**, 73–82 (2017).
- [3] Frank, J. Approaches to large-scale structures. *Curr. Opin. Struct. Biol.* **5**, 194–201 (1995).
- [4] McEwen, B., Renken, C., Marko, M. & Mannella, C. Principles and practice in electron tomography. *Methods Cell Biol.* **89**, 129–168 (2008).
- [5] McIntosh, R., Nicastro, D. & Mastrorade, D. ew views of cells in 3d: an introduction to electron tomography. *Trends Cell Biol.* **15**, 43–51 (2005).
- [6] Nicastro, D., Frangakis, A., Typke, D. & Baumeister, W. Cryo-electron tomography of neurospora mitochondria. *J. Struct. Biol.* **129**, 48–56 (2000).
- [7] Guesdon, A., Blestel, S., Kervrann, C. & Chrétien, D. Single versus dual-axis cryo-electron tomography of microtubules assembled in vitro: limits and perspectives. *J. Struct. Biol.* **181**, 169–78 (2013).
- [8] Best, C., Nickell, S. & Baumeister, W. Localization of protein complexes by pattern recognition. *Methods Cell Biol.* **2007**, 615–638 (2007).
- [9] Albert, S. et al. Direct visualization of degradation microcompartments at the er membrane. *Proc. Natl. Acad. Sci.* **117**, 1069–1080 (2020).
- [10] Förster, F., Pruggnaller, S., Seybert, A. & Frangakis, A. S. Classification of cryo-electron sub-tomograms using constrained correlation. *J. Struct. Biol.* **161**, 276–286 (2008).
- [11] Martinez-Sanchez, A. et al. Template-free detection and classification of membrane-bound complexes in cryo-electron tomograms. *Nat. Methods* **17**, 209–216 (2020).
- [12] LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
- [13] LeCun, Y., Kavukcuoglu, K. & Farabet, C. Convolutional networks and applications in vision. In *Proc. IEEE Int. Symp. on Circuits and Systems*, 253–256 (2010).
- [14] Hinton, G. et al. Deep neural networks for acoustic modeling in speech recognition. *IEEE Signal Process. Mag.* **29**, 82–97 (2012).
- [15] Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet classification with deep convolutional neural networks. In *Proc. Neural Inf. Processing Systems (NIPS)*, 1–9 (2012).

- [16] Long, J., Shelhamer, E. & Darrell, T. Fully convolutional networks for semantic segmentation. In *Proc. Conf. Comput. Vis. Pattern Recognition*, 3431–3440 (2014).
- [17] Szegedy, C., Toshev, A. & Erhan, D. Deep neural networks for object detection. In *Proc. Conf. Neural Inf. Processing Systems (NIPS)*, 1–9 (2013).
- [18] Sullivan, D., Winsnes C., L., and Akesson & Lundberg, E. Deep learning is combined with massive-scale citizen science to improve large-scale image classification. *Nat. Biotechnology* **36**, 820–828 (2018).
- [19] Moen, E. et al. Deep learning for cellular image analysis. *Nat. Methods* **16**, 1223–1246 (2019).
- [20] Falk, T. et al. U-net – deep learning for cell counting, detection, and morphometry. *Nat. Methods* **16**, 67–70 (2019).
- [21] Gupta, A. et al. Deep learning in image cytometry: A review. *Cytometry Part A* **95**, 366–380 (2019).
- [22] Belthangady, C. & Royer, L. Applications, promises, and pitfalls of deep learning for fluorescence image reconstruction. *Nat. Methods* **16**, 1215–1225 (2019).
- [23] Ouyang, W., Aristov, A., Lelek, M., Hao, X. & Zimmer, C. Deep learning massively accelerates super-resolution localization microscopy. *Nat. Biotechnology* **36**, 460–468 (2018).
- [24] Weigert, M. et al. Content-aware image restoration: pushing the limits of fluorescence microscopy. *Nat. Methods* **12**, 1090–1097 (2018).
- [25] Wu, X., Zeng, X., Zhu, Z., Gao, X. & Xu, M. Template-based and template-free approaches in cellular cryo-electron tomography structural pattern mining. *Comp. Biology Chapter 11*, 1146–1152 (2019).
- [26] Wang, F. et al. DeepPicker : A deep learning approach for fully automated particle picking in cryo-EM. *J. Struct. Biol.* **195**, 325–336 (2016).
- [27] Al-Azzawi, A., Ouadou, A., Tanner, J. J. & Cheng, J. Autocrypicker: an unsupervised learning approach for fully automated single particle picking in cryo-em images. *BMC Bioinformatics* **20**, 326 (2019).
- [28] Wagner, T. et al. Sphire-cryolo is a fast and accurate fully automated particle picker for cryo-em. *Communications Biology* **2**, 218 (2019).
- [29] Bepler, T. et al. Positive-unlabeled convolutional neural networks for particle picking in cryo-electron micrographs. *Nat. Methods* **16**, 1153–1160 (2019).
- [30] Tegunov, D. & Cramer, P. Real-time cryo-electron microscopy data preprocessing with warp. *Nat. Methods* **16**, 1146–1152 (2019).
- [31] Chen, M. et al. Convolutional neural networks for automated annotation of cellular cryo-electron tomograms. *Nat. Methods* **14**, 983–985 (2017).
- [32] Che, C. et al. Improved deep learning based macromolecules structure classification from electron cryo tomograms. *Mach. Vis. Appl.* **29**, 1227–1236 (2018).
- [33] Gubins, I. et al. SHREC’19 Track : Classification in cryo-electron tomograms. In *Eurographics Workshop on 3D Object Retrieval, SHREC – 3D Shape Retrieval Contest* (2019), 1–6 (2019). URL <https://www2.projects.science.uu.nl/shrec/cryo-et/2019/>.
- [34] Hrabe, T. et al. PyTom: A python-based toolbox for localization of macromolecules in cryo-electron tomograms and subtomogram analysis. *J. Struct. Biol.* **178**, 177–188 (2012).
- [35] Gubins, I. et al. Shrec 2020: Classification in cryo-electron tomograms. *Computers & Graphics* **91**, 279–289 (2020).
- [36] Buchholz, T., Jordan, M., Pigino, G. & Jug, F. Cryo-care: Content-aware image restoration for cryo-transmission electron microscopy data. In *IEEE Int. Symp. Biomedical Imaging (ISBI)*, 502–506 (2019).
- [37] Moebel, E. & Kervrann, C. A monte carlo framework for missing wedge restoration and noise removal in cryo-electron tomography. *J. Struct. Biol.;* **X 4**, 100013 (2020).
- [38] Wietrzynski, W. et al. Charting the native architecture of chlamydomonas thylakoid membranes with single-molecule precision. *eLife* **9**, e53740 (2020).
- [39] Pfeffer, S. et al. Dissecting the molecular organization of the translocon-associated protein complex. *Nat. Communications* **8**, 14516 (2017).
- [40] Freeman-Rosenzweig, E. et al. The eukaryotic *co2*-concentrating organelle is liquid-like and exhibits dynamic reorganization. *Cell* **171**, 148–162 (2017).
- [41] Förster, F. & Hegerl, R. Structure determination *in situ* by averaging of tomograms. *Cell. Electron Microsc.* **79**, 741–767 (2007).
- [42] Rolnick, D., Veit, A., Belongie, S. & Shavit, N. Deep learning is robust to massive label noise. *arXiv Prepr.* (2017). [arXiv:1705.10694v2](https://arxiv.org/abs/1705.10694v2).
- [43] Chen, Y., Pfeffer, S., Hrabe, T., Schuller, J. M. & Förster, F. Fast and accurate reference-free alignment of subtomograms. *J. Struct. Biol.* **182**, 235–245 (2013).
- [44] Sanchez-Garcia, R., Segura, J., Maluenda, D., Carazo, J. & C.O.S., S. Deep consensus, a deep learning-based approach for particle pruning in cryo-electron microscopy. *IUCrJ.* **5**, 854–865 (2018).
- [45] Comaniciu, D., Meer, P. & Member, S. Mean Shift : a robust approach toward feature space analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**, 603–619 (2002).
- [46] Gipson, B. R. et al. Automatic recovery of missing amplitudes and phases in tilt-limited electron crystallography of two-dimensional crystals. *Phys. Rev. E* **84**, 011916 (2011).
- [47] Deng, Y. et al. Icon: 3d reconstruction with “missing-information” restoration in biological electron tomography. *J. Struct. Biol.* **195**, 100–112 (2016).
- [48] Biyani, N. et al. Image processing techniques for high-resolution structure determination from badly ordered 2d crystals. *J. Struct. Biol.* **203**, 120–134 (2018).
- [49] Förster, F., Han, B. G. & Beck, M. Visual proteomics. *Methods Enzymol.* **483**, 215–243 (2010).
- [50] Yosinski, J., Clune, J., Bengio, Y. & Lipson, H. How transferable are features in deep neural networks ? In *Proc. Conf. Neural Inf. Processing Systems (NIPS)*, 3320–3328 (2014).
- [51] Kingma, D. P. & Ba, J. L. Adam: a method for stochastic optimization. *arXiv Prepr.* (2014). [arXiv:1412.6980v9](https://arxiv.org/abs/1412.6980v9).
- [52] Salehi, S. S. M., Erdogmus, D. & Gholipour, A. Tversky loss function for image segmentation using 3D fully convolutional deep networks. *arXiv Prepr.* (2017). [arXiv:1706.05721v1](https://arxiv.org/abs/1706.05721v1).
- [53] Ronneberger, O., Fischer, P. & Brox, T. U-Net: Convolutional networks for biomedical image segmentation. In *Proc. Med. Image Comput. Comput. Assist. Interv. (MICCAI)*, vol. 9351, 234–241 (2015).
- [54] Milletari, F., Navab, N. & Ahmadi, S.-a. V-Net : fully convolutional neural networks for volumetric medical image segmentation. *arXiv Prepr.* 1–11 (2016). [arXiv:1606.04797v1](https://arxiv.org/abs/1606.04797v1).
- [55] Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. In *Proc. Int. Conf. Learn. Representation*, 1–14 (2015).

- [56] Martinez-Sanchez, A., Garcia, I., Asano, S., Lucic, V. & Fernandez, J.-j. Robust membrane detection based on tensor voting for electron tomography. *J. Struct. Biol.* **186**, 49–61 (2014).
- [57] Kremer, J. R., Mastronarde, D. N. & McIntosh, J. R. Computer visualization of three-dimensional image data using IMOD. *J. Struct. Biol.* **116**, 71–76 (1996).
- [58] Zivanov, J. *et al.* New tools for automated high-resolution cryo-em structure determination in relion-3. *Elife* **7**, e42166 (2018).
- [59] Bharat, T. B. & Scheres, S. Resolving macromolecular structures from electron cryo-tomography data using subtomogram averaging in relion. *Nat. Protoc* **11**, 2054–2065 (2016).
- [60] Harauz, G. & van Heel, M. Exact filters for general geometry three dimensional reconstruction. *Optik (Stuttgart)* **78**, 146–156 (1996).
- [61] Rosenthal, P. B. & Henderson, R. Optimal determination of particle orientation, absolute hand, and contrast loss in single-particle electron cryomicroscopy. *J. Mol. Biol.* **333**, 721–745 (2003).
- [62] Chen, S. *et al.* High-resolution noise substitution to measure overfitting and validate resolution in 3d structure determination by single particle electron cryomicroscopy. *Ultramicroscopy* **135**, 24–35 (2013).
- [63] de la Rosa-Trevín, J. *et al.* Scipion: A software framework toward integration, reproducibility and validation in 3d electron microscopy. *Journal of Structural Biology* **195**, 93–99 (2016).
- [64] Pettersen, E. F. *et al.* UCSF Chimera - A visualization system for exploratory research and analysis. *J. Comput. Chem.* **25**, 1605–1612 (2004).
- [65] Goddard, T. *et al.* Ucsf chimeraX: Meeting modern challenges in visualization and analysis. *Protein Sci.* **27**, 14–25 (2018).
- [66] Albert, S. *et al.* Proteasomes tether to two distinct sites at the nuclear pore complex. *Proc. Natl. Acad. Sci.* **114**, 201716305 (2017).
- [67] Henderson, R. Avoiding the pitfalls of single particle cryo-electron microscopy : Einstein from noise. *Proc. Natl. Acad. Sci.* **110**, 18037–18041 (2013).

Methods

Description of DeepFinder algorithm – Stage I (Training). In the training stage (see Fig. 1A), DeepFinder parameters are learned from pairs of tomograms and their corresponding voxel-wise annotations. The underlying 3D CNN requires that each voxel is annotated as a member of a given macromolecular species, or as background. While voxel-wise classification examples are naturally available for synthetic data, this is often not the case for experimental data. In our case, the experts only supplied the 3D coordinates of macromolecules of interest without labeling voxels, in several training tomograms. In practice, voxel-wise annotation cannot be performed manually in cryo-ET for two reasons: *i*) it is time consuming to individually label each voxel belonging to a 3D macromolecule; *ii*) noise and artifacts in the data make it difficult to accurately distinguish macromolecule borders.

To get voxel-wise annotations, we propose two strategies starting from 3D spatial coordinates of the particles:

- *Shape-based annotation:* To reduce the amount of "label noise", we propose a second strategy to convert macromolecule coordinates into voxel-wise annotations. Instead of using spheres, we generate the target by using the structural shape of the macromolecule. These targets are more difficult to obtain, because information on the shape and the orientation of the macromolecules are needed. In order to obtain this information, we exploit the commonly used subtomogram averaging procedure⁴¹ (see Fig. 2). While this strategy provides more accurate targets, it still induces "label noise", albeit less than when using spheres. Indeed, as we project the same average shape (computed with subtomogram averaging) several times at different positions into the tomogram for training, not all structural details of individual macromolecules are not captured.
- *Sphere-based annotation:* 3D spheres are positioned at each annotated location in the tomogram. The sphere radius should be similar to the size of the target macromolecule. As the experts generally know the size of the target macromolecule (e.g., 15 nm radius for ribosomes to 5 nm radius for small particles like Rubisco), this parameter can be readily input into DeepFinder. This strategy is easy to implement and is appealing in terms of computing time. The main limitation is that spheres may introduce significant "label noise". The resulting errors in the annotations hinder the training procedure, and the segmentation network may have sub-optimal performance. Nonetheless, it has been shown that CNNs have a natural robustness to moderate levels of "label noise"⁴², which is also confirmed in our experiments with DeepFinder.

In our experiments, all results have been obtained by using shape-based annotations, except for the comparison between sphere-based and shape-based annotations (Fig. 3C).

Training datasets and batch generation. Due to memory limitations, it may not be feasible to load the whole tomogram set with the corresponding targets during training. Instead, for each training iteration, we sample a group of smaller 3D patches to constitute a "batch". The patch size should be large enough to capture sufficient context information. For a macromolecule with a 10 voxel radius, we choose a patch size of $56 \times 56 \times 56$ voxels. Our implementation is such that only the current batch is loaded into memory. This helps to handle large datasets with limited computational resources, allowing DeepFinder to be applied on lower-end consumer GPUs with less RAM.

Another problem is the high imbalance between class labels. Because of the small size of the macromolecules compared to the tomogram size, more than 99% of voxels are annotated as "background". This causes the trained network to be skewed towards the over-represented class. Therefore, we "guide" the sampling procedure by selecting patches only at annotated positions such that each patch contains at least one macromolecule. An additional benefit of sampling only at annotated positions is that the amount of false negatives is reduced (here the negative class is the background). It is indeed common that annotations are not exhaustive. Therefore false "background" labels remain at missed macromolecule positions, contributing to increased "label noise". The proposed patch sampling procedure does not discard all false negatives (e.g., a false positive could be neighboring a true positive), but the number should be relatively small at the end.

An additional problem is the imbalance between competing macromolecule classes: for instance, 20 macromolecules of class #1 (e.g., *mb-ribo* particles) versus 100 macromolecules of class #2 (e.g., *ct-ribo* particles). To reduce this imbalanced class effect, we apply a bootstrapping algorithm (i.e., resampling), so that the distribution of the positive classes in a batch will be uniform. This stochastic resampling procedure is effective for sampling the under-represented classes more frequently than the over-represented classes.

It is also common in deep learning to use data augmentation to artificially increase the size of training sets. The idea is to apply geometric transforms to training images. In our approach, we implement "data augmentation" by applying a 180° rotation with respect to the microscope tilt-axis, to each training example. Nevertheless, we do not use "mirror" operations or geometric deformations because the macromolecule structure (including its chirality) is the main clue for detection. Also, we do not use random rotations because of the well-determined orientation of missing wedge artifacts, which is preserved when applying 180° rotations with respect to tilt-axis. Finally, we apply random shifts to improve invariance to translations.

Optimization. In our experiments, DeepFinder has been computationally trained with the ADAM algorithm, chosen for its good convergence rate⁵¹, by setting the learning rate to 0.0001, the exponential decay rate to 0.9 for the first moment estimate, and to 0.999 for the second moment estimate. The batch size was set to 25 and the patch size to $56 \times 56 \times 56$ voxels. No regularization (e.g., L_2 regularizer, "drop out") was needed for processing the datasets. We implemented both the categorical cross-entropy (as implemented in Keras) and the Dice⁵² loss functions. On the synthetic SHREC'19 dataset, Dice loss function was able to better localize the smallest macromolecules. However, both loss functions yielded similar results on real cryo-ET data (localization of *mb-ribo* and *ct-ribo* particles).

Description of DeepFinder algorithm – Stage II (Analysis). We describe below the two steps of the analysis stage of DeepFinder algorithm (see Fig. 1A).

Analysis stage (Stage II): Step #1 – Multi-class voxel-wise classification. The objective is to provide a segmentation map for which each voxel is assigned to a class label (representing a macromolecular species). The architecture of DeepFinder is based on U-net^{16,20}, an "encoder-decoder" type of network designed for segmenting images in an end-to-end manner. U-net is an extension of the fully convolutional network, which achieves multi-resolution feature representation and produces high-resolution label maps. The DeepFinder architecture consists of a down-sampling path (i.e., encoder) needed to generate global information and an up-sampling path (i.e., decoder) used to generate high-resolution outputs^{53,54}, i.e., local information (see Supplementary Fig. 8). Down-sampling is performed with max-pooling layers (factor 2) and up-sampling with up-convolutions¹⁶ (sometimes called "backward convolution"), which is essentially a trained and non-linear up-sampling operation. Combining global and local information is performed by concatenating features at different spatial resolutions. The features are then processed with the convolutional layers of the up-sampling path. The fully convolutional nature of DeepFinder allows one to use various input patch shapes, with the constraint that the patch size must be a multiple of four, because of the two down-sampling stages. Large tomograms can be processed by using an overlap-patch strategy. Unlike Milletari *et al.*⁵⁴, DeepFinder is not that "deep", since we found that using more than two down-sampling stages does not increase the segmentation results. Also, we used only $3 \times 3 \times 3$ filter sizes as described by Simonyan *et al.*⁵⁵. The rationale behind this choice is that two consecutive $3 \times 3 \times 3$ filters mimic a larger $5 \times 5 \times 5$ filter but with fewer parameters. Training is then faster and easier and requires less memory. An important concept in neural architectures is the receptive field of deepest neuron layers. This determines the size of the spatial context to be used to make decisions. Considering a large spatial context is essential to handle macromolecule classes involving interactions with the environment, for instance interactions with cellular membranes. It is established that adding convolutional layers after down-sampling operations is appropriate to enlarge the spatial context⁵⁴. Accordingly, we added two supplementary convolutional layers in the lowest stage of our architecture. In the end, the receptive field size of our network is $48 \times 48 \times 48$ voxels. To complete the description, we use rectified linear units (ReLU)¹⁵ as activation functions for every layer except the last one which uses a *soft-max* function. While ReLU is

a popular choice to tackle non-linearities in the network, the *soft-max* function is necessary in order to interpret the network outputs as pseudo-probabilities for each class. In summary, step #1 is capable of robustly classifying the cryo-ET tomogram into N classes with a high accuracy.

Analysis stage (Stage II): Step #2 – Clustering for macromolecule localization. Given the multi-class voxel-wise classification map and classification errors, the objective of step #2 is to estimate the spatial coordinates of each particle of a given macromolecular species. The voxel labels should be spatially clustered into distinct 3D connected components, each cluster corresponding to a unique particle. Because of noise, non-stationarities in the background, and artifacts in the tomogram, step #1 generates isolated labels or very small groups of voxels, as well as groups that contain different labels. Post-processing is then necessary to aggregate neighboring voxels into 3D connected components and to assign a unique label to each component (or cluster). The clusters significantly smaller than the size of target particles are considered as false positives and are discarded. The centroids of meaningful clusters are computed to determine the location of the particles. As the centroid is computed by uniformly averaging the coordinates of cluster voxels, we are able to numerically produce positions with sub-voxel precision. As several voxel labels can be spatially grouped in a given cluster, the most frequent label is assigned to the detected particle. To address this task, we used the popular "mean-shift" clustering algorithm⁴⁵. The main advantage of "mean-shift" is that it is controlled by only one parameter, commonly called the bandwidth, which is directly related to the average size of macromolecules. The K-means algorithm was not considered further since the number of clusters must be provided as input parameter by the user.

Computing times for Stage I and Stage II are reported in Supplementary Table 2, and are relatively small when compared to template matching. For instance, if we consider a set of 10 tomograms, the computing times are 300 hours (10×30 hours) with template matching, while it is 23.30 hours (20 hours (training) + 10×20 min (inference)) with DeepFinder. This estimation does not take into account the time needed to annotate the data (for DeepFinder) nor the time to generate the template and subsequent subtomogram classification (for template matching). Moreover, as the DeepFinder training expands and becomes more generalized, less training is required to analyze new data, further increasing its speed advantage over template matching. Note that the shape-based target generation takes more computing time than the sphere-based strategy, because it relies on a sub-tomogram alignment routine (with no classification procedure) performed on sub-resolved (binned) tomograms (for dataset #2 we used subtomograms of size $47 \times 47 \times 47$ voxels).

Description of synthetic data (Dataset #1). The synthetic dataset was generated by the cryo-Electron Microscopy group of Utrecht University for the challenge "SHREC 2019: Classification in Cryo-Electron Tomograms", organized for the Eurographics Workshop on "3D Object Retrieval"³³. The dataset has been especially designed to objectively evaluate the performance of localization and classification algorithms in cryo-electron tomography. The dataset was created as follows³³: i) tomogram density maps were generated using Protein Data Bank (PDB) structures (identifiers: 1bxn, 1qvr, 1s3x, 1u6g, 2cg9, 3cf3, 3d2f, 3gl1, 3h84, 3qm1, 4b4t and 4d8q); ii) a series of projection images was generated with a $\pm 57^\circ$ tilt-range and a 3° tilt-step; iii) projection images were degraded by adding noise and contrast transfer function such that signal-to-noise ratio was 0.02; iv) the tomogram was reconstructed using weighted back-projection. The resulting tomograms have a size of $512 \times 512 \times 512$ voxels and a voxel size of 1 nm.

In our experiments, we split the dataset #1 into training, validation and test sets. The training set was composed of 8 tomograms annotated with a total of 19,956 macromolecules. The validation set was composed of one tomogram annotated with a total of 2,490 macromolecules. The test set was composed of one tomogram annotated with a total of 2,540 macromolecules. The training was performed until convergence, which took 10,000 iterations and 33 hours of computation using the Dice loss function on a Nvidia Tesla K80 GPU.

Description of experimental cryo-ET data (Dataset #2). The real cryo-ET dataset is composed of 57 tomograms of *Chlamydomonas reinhardtii* cells, and was used for membrane-bound 80S ribosome (*mb-ribo*) annotation. To produce the annotations, the experts used template matching, and then filtered out the false positives by applying the CPCA subtomogram classification algorithm¹⁰, and careful visual inspection (see illustration in Supplementary Fig. 7). To reduce the computational cost, the tomograms were binned to a voxel size of $13.68\text{\AA}$ with tomogram dimensions of $928 \times 928 \times 464$ voxels. Tilt range was $\pm 60^\circ$ with an increment of 2° .

As starting point, Dataset #2 was annotated by localizing 8,795 *mb-ribo* particles under expert supervision. Then, we introduced two additional classes (without expert supervision): the cytosolic ribosomes (*ct-ribos*) class and the *membrane* class. Annotations for these two additional classes were obtained by applying semi-automatized computational tools to tomograms. First, we selected members of the *membrane* class by employing an algorithm (TomoSegMemTV) dedicated to cell membrane segmentation⁵⁶. Meanwhile, we obtained training (*ct-ribo*) particles by applying template matching and selecting the most isolated candidates, located further than $273.6\text{\AA}$ (i.e., the ribosome diameter) from the segmented membranes. The motivation behind adding new classes to the available annotations was twofold: on the one hand, we wanted to demonstrate the ability of DeepFinder to localize and identify multiple molecular species on real data with one pass; on the other hand, we noticed that the multi-class approach tends to improve the discriminating power of CNN, therefore better finding *mb-ribo* particles in experimental tomograms. When trained with only the *mb-ribo* examples, DeepFinder additionally finds unwanted cytosolic ribosomes. By considering the *ct-ribo* class in the training, we encourage the network to better discriminate both ribosome subclasses (corresponding to binding states). Note that these annotations of membrane components and cytosolic ribosomes have been obtained without the supervision of an expert. Therefore more errors are expected for these two classes when compared to the *mb-ribo* examples, reliably annotated by the experts.

Dataset #2 was arbitrarily split into training, validation and test sets. The training set was composed of 48 tomograms annotated with 6,834 *mb-ribos* and 6,687 *ct-ribo* particles. The validation set was one tomogram annotated with 222 *mb-ribos* and 254 *ct-ribo* particles. The test set was composed of eight tomograms annotated with 1,736 *mb-ribos* and 2,594 *ct-ribo* particles. In our experiments, we considered "background" and three classes: *mb-ribo*, *ct-ribo* and "membrane". The "background" voxels are those that do not belong to any of the other three classes. Training was achieved using the categorical cross-entropy as a loss function (similar results were obtained with the Dice loss function). The training was performed for 6,000 iterations and took 20 hours of computation on a Nvidia Tesla K80 GPU (see plots of loss evolution during training in Supplementary Fig. 17).

Description of experimental pyrenoid cryo-ET data (Dataset #3). The dataset is composed of five tomograms of the *Chlamydomonas* pyrenoid, which have been annotated with a total of 176,229 Rubisco holoenzymes. The tomogram size is $928 \times 928 \times 464$ voxels, the voxel size is $13.68\text{\AA}$, and the tilt range is $\pm 60^\circ$ with an increment of 2° . The expert annotations were obtained with template matching (excluding hits outside the pyrenoid matrix using a manually segmented mask), followed by subtomogram classification using CPCA clustering¹⁰. For more details about the dataset and the annotations, see⁴⁰.

Dataset #3 was split into training and test sets. The training set was composed of four tomograms annotated with 129,662 Rubisco complexes. The test set was composed of one tomogram annotated with 46,567 Rubisco complexes. The training was performed for 20,000 iterations and took 35 hours on a Tesla K80 GPU. The loss curve converged after 10,000 iterations, suggesting that only half the time (i.e., 17.5 hours) would have been sufficient to obtain the presented results.

Description of experimental thylakoid cryo-ET data (Dataset #4). The dataset is composed of four tomograms of *Chlamydomonas* thylakoid membranes, annotated with a total of 637 photosystem II (PSII) complexes. The bin4 tomogram size is $928 \times 928 \times 464$ voxels with a voxel size of $13.68\text{\AA}$, and the tilt range is $\pm 60^\circ$ with an increment of 2° . The expert annotations were obtained by manual particle picking using the "membranogram" approach³⁸. For more details about the dataset and the annotations, see³⁸.

Dataset #4 was split into training, validation and test sets. The training set was composed of three tomograms annotated with 298 PSII complexes. The validation set was a region from one of the training tomograms annotated with 143 PSII complexes. The distance between training and validation particles was

large enough to avoid overlap. The test set was composed of one tomogram annotated with 196 PSII complexes. The training was performed for 5,000 iterations and took 3.5 hours of computation on a Nvidia Quadro RTX 5000 GPU.

Ribosome subtomogram averaging. From the test tomograms of dataset #2, we extracted 1,805 membrane-bound ribosomes and 1,394 cytosolic ribosomes. These numbers were obtained by choosing the cluster size that maximizes the F_1 -score w.r.t. the expert annotations (see Supplementary Fig. 13). Binary masks have been used to focus on regions of interest. Volumes of size $184 \times 184 \times 184$ at a voxel size of $3.42 \text{ \AA}$ were extracted and reconstructed from the unbinned tilt series using the IMOD program subtomosetup⁵⁷. The published subtomogram average of the ER membrane-bound ribosome (EMD-4145)³⁹ has been used as a reference for subtomogram alignment and averaging. The membrane density of EMD-4145 has been manually erased, so that the alignment is focused on the ribosome structure. Subtomogram alignment was achieved by carrying out one iteration of fast rotational matching in PyTOM³⁴. The same process has been applied to obtain the expert subtomogram average in Fig. 4F.

PSII subtomogram averaging. *DeepFinder particles:* From our test tomogram of dataset #4, we extracted the 329 picked positions that were close to segmented membranes. For each of these positions, we computed its corresponding normal vector to the membrane. We incorporated this normal vector as a prior on the particles' orientation angles in the first subtomogram averaging step in RELION⁵⁸, which we also used for all subsequent averaging tasks. The protocol is based on previous averaging protocols^{11,59}, which in turn uses RELION. Hereby, we averaged subvolumes of size $90 \times 90 \times 90$ and enforced C10 symmetry in order to only align the membranes in the subvolumes and postpone the refinement of the structures to later steps. Next, having aligned the membranes for all the picked positions, we performed the first refinement step, enforcing C2 symmetry, whereby C2 is PSII's symmetry group. In order to purify the PSII picks, we performed 5-fold subtomogram classification and separated those classes that did clearly look different from PSII. The separated classes were then examined in another 5-fold classification step. This procedure was repeated three times in order to avoid discarding true picks. In the next step, the remaining picks were refined again using C2 symmetry. A final 5-fold classification gave us classes showing different membrane proteins. We merged the classes that were similar to PSII and performed a final refinement step for the remaining 246 positions, giving us the final subtomogram average.

Rubisco subtomogram averaging. *DeepFinder particles:* A list containing the coordinates of the top 30,000 hits obtained by DeepFinder on the previously published tomogram of the *Chlamydomonas reinhardtii* pyrenoid matrix⁴⁰ (EMD-12749) were used to extract boxes for subtomogram averaging. Volumes comprising 64×64 voxels at a pixel size of $3.42 \text{ \AA}$ were extracted and reconstructed from the unbinned tilt series (EMPIAR-10694) using the IMOD program subtomosetup⁵⁷. The tilt series was corrected for the CTF by phase flipping in IMOD prior to subtomogram extraction. The previously published subtomogram average of rubisco (EMD-3694) was low-pass filtered to $40 \text{ \AA}$ and used as initial reference for the iterative subtomogram aligning and averaging in PyTOM³⁴. Using this reference, one iteration of global search was carried out in real space with an angular step of 18 degrees using the PyTOM script align.py. Subsequently, four iterations of local real space alignment were carried out with the PyTOM script GLocalJob.py, starting with an angular step of 9 degrees. Particles were randomly assigned into two half-sets which were independently refined ("gold standard"). Adaptive alignment resolution and angular sampling were determined automatically based on the Fourier shell correlation⁶⁰ (FSC) curve between the two half-maps at each iteration. During refinement C4 symmetry was imposed. Additional D1 symmetry was imposed to the final half-maps, amounting to the known D4 symmetry of Rubisco. Finally, the half-maps were postprocessed in RELION⁵⁸ using a soft mask, yielding a final resolution of $15 \text{ \AA}$ (FSC > 0.143⁶¹) after correcting for artificial correlations induced by the mask⁶².

Template matching particles: We extracted subtomograms from the same list of coordinates from the top 30,000 template matching hits in the "positive class" which had been used for the previously published subtomogram average of Rubisco (EMD-3694)⁴⁰. Subtomogram extraction, averaging and postprocessing were carried out exactly as described above for the DeepFinder particles. This set of particles yielded a reconstruction at $15 \text{ \AA}$ (FSC > 0.143).

Description of template matching. The template matching results on the synthetic SHREC'19 dataset were provided by courtesy of Utrecht University (Department of Information and Computing Sciences, Department of Chemistry) and were obtained as follows: *i*) template matching was applied for each class. As the particles have been simulated, we know the true number of candidates for each tomogram: 2,500 particles/12 classes = 211. Then, the top 211 candidates for each class were sequentially selected and extracted, from the largest macromolecules to the smallest ones, in a way that all these candidates do not overlap with the already extracted candidates.

We used the template matching algorithm implemented in the PyTom toolbox³⁴ to analyze Dataset #2 (*Chlamydomonas reinhardtii* cells). The template has been provided by the author of the ribosome annotations, who used a low-resolution subtomogram average generated from the dataset, by using manually selected ribosomes and FRM alignment.

Evaluation. We used the F_1 -score to assess localization performance. The F_1 -score (also known as Dice coefficient) is the harmonic mean of *Precision* and *Recall* that depend on the number of true positives (TP):

- *Recall* (also known as sensitivity):

$$R = \frac{\text{Number of TP}}{\text{Number of particles in the ground truth}},$$

- *Precision* (also known as positive predictive value):

$$P = \frac{\text{Number of TP}}{\text{Number of detected particles}}.$$

The F_1 -score, defined as follows

$$F_1 = 2 \frac{RP}{R + P},$$

allows one to evaluate performance by considering a single value.

Moreover, for each test tomogram, bounding boxes were placed at each true location. These bounding boxes were of same size as the target macromolecule, and a detected particle was considered to be a TP if the estimated centroid was located within the bounding box.

Code availability and implementation details

The code can be downloaded for free from our GitLab website (<https://gitlab.inria.fr/serpico/deep-finder>) along with accompanying documentation (<https://deepfinder.readthedocs.io/en/latest/>). DeepFinder is embedded into the new release of Scipion⁶³ (<https://github.com/scipion-em/scipion-em-deepfinder>), an open-source image processing framework for cryo-electron microscopy (<http://scipion.i2pc.es/>).

Each step of DeepFinder shown in Fig. 1A can be executed with scripts using the API (examples are provided) or with a graphical user interface. These steps may also be embedded in other workflows, e.g., if the user needs only the segmentation step. To implement DeepFinder, we used Keras [http://keras.io], an open-source toolbox written in Python and using the TensorFlow framework. The code can be downloaded for free from the following GitLab website: [https://gitlab.inria.fr/serpico/deep-finder] and includes online documentation: https://deepfinder.readthedocs.io/en/latest/.

All training procedures were achieved using a Nvidia Tesla K80 GPU, running Cuda 8 and cuDNN 6. Below, we display the memory consumption of DeepFinder for different training parameters.

Patch size (voxels)	Batch size	
	15	25
40 × 40 × 40	2.56 GB	3.76 GB
56 × 56 × 56	6.16 GB	9.86 GB

We used Chimera⁶⁴ and ChimeraX⁶⁵ for 3D visualization purposes.

Data availability

The synthetic dataset is available on the website of the SHREC'19 challenge (http://www2.projects.science.uu.nl/shrec/cryo-et/2019/). A tomogram from the experimental dataset of *Chlamydomonas reinhardtii* cells^{39;66} can be found in the Electron Microscopy Data Bank (EMDB) (https://www.ebi.ac.uk/pdbe/entry/emdb/EMD-3967) under accession number EMD-3967. Additional tomograms depicting Thylakoids can be downloaded from EMDB under accession numbers EMD-10780, EMD-10781, EMD-10782, and EMD-10783.

The subtomogram average of Rubisco obtained from the set of DeepFinder particles was deposited at the Electron Microscopy Data Bank (EMDB) under accession code EMD-YYYY. The average obtained from the set of particles identified by template matching was deposited at the Electron Microscopy Data Bank (EMDB) under accession code EMD-ZZZZ.

Acknowledgements

This work was jointly supported by the Fourmentin-Guilbert Foundation and Région Bretagne (Brittany Council). Calculations were performed on the Inria Rennes computing grid facilities partly funded by France-BioImaging infrastructure (French National Research Agency - ANR-10-INBS-04-07, "Investments for the future") and at the Max Planck Institute for Biochemistry computing cluster, Martinsried, Germany. L. Lamm, R.D. Righetto, W. Wietrzynski, T. Peng, and B. Engel were supported by DFG grant EN 1194/1-1 as part of FOR 2092, The Munich School for Data Science (MUDS), and Helmholtz Association.

We would like to thank F. Förster and M. Killinger for fruitful discussions about cryo-ET data analysis, and deep learning applied to large 3D volumes analysis, respectively.

We thank the organizers of the SHREC'19 challenge for helpful assistance and for providing the template matching results: I. Gubins and R.C. Veltkamp (Utrecht University, Department of Information and Computing Sciences), G. van der Schot and F.G. Forster (Utrecht University, Department of Chemistry).

Finally, we thank S. Prima for careful reading of the paper and valuable suggestions and comments.

Author contributions

E. Moebel designed and implemented the presented DeepFinder method, carried out the biocomputing experiments, and wrote the manuscript. C. Kervrann supervised the project, was in charge of overall direction and planning, and co-wrote the manuscript. D. Larivière, E. Fourmentin, and C. Kervrann devised the project and the main conceptual ideas. A. Martinez-Sanchez also assisted the conceptualization. B. Engel and W. Baumeister facilitated access to datasets. B. Engel, S. Albert, W. Wietrzynski, and S. Pfeffer provided the *Chlamydomonas reinhardtii* datasets and annotations (Datasets #2, #3 and #4). A. Martinez-Sanchez, J. Ortiz and B. Engel conceived experiments on real datasets. L. Lamm, R.D. Righetto, W. Wietrzynski, and T. Peng performed experiments on datasets depicting thylakoid membranes and pyrenoid matrices within vitreously-frozen *Chlamydomonas reinhardtii* cells. All authors provided critical feedback and helped shape the research, analysis and manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at ...

Correspondence and requests for materials should be addressed to C.K.

Note

Through the LifeExplorer project, the Fourmentin-Guilbert Foundation has been a pioneer in approaching the structural modeling and visualization of entire cells. It is expected from the creation of interactive 3D avatars of cellular environments, bridging from the level of atoms to the level of cells, that the rules governing the spatiotemporal organization of the cytoplasm could be revealed. Such an approach requires to make an inventory and a cartography of all the components constituting a single cell.

The technique of choice for such a mapping is cryo-electron microscopy applied on frozen but intact cells. For years, the Fourmentin-Guilbert Foundation has supported the Max Planck Institute of Biochemistry, headed by W. Baumeister whose team has been capable of delivering whole cryo-electron tomograms of *E. coli* cells at an unprecedented resolution.

The next big challenge, still an ongoing challenge, was to recognize macromolecular components within the tomograms. Most of the effort of the scientific community was put on the identification of the ribosomes thanks to a methodology relying on single-particle analysis and template matching and giving impressive results. However, it is likely that such an approach will be limited to large macromolecules like the ribosomes. Facing this challenge, the Fourmentin-Guilbert Foundation has solicited the research group headed by C. Kervrann to develop alternative recognition methods having the potential to help the *in-situ* identification of the thousands of proteins left in the dark. The Serpico Project-Team has then developed and compared with well-established methods new approaches based on deep learning and capable of identifying and counting the "gold standard" ribosomes within a tomogram. These methods, as an alternative to template matching, should also have the potential to apply on particles smaller and rounder than ribosomes.

Supplementary information

Deep learning improves macromolecule identification in 3D cellular cryo-electron tomograms

E. Moebel, A. Martinez-Sanchez, L. Lamm, R.D. Righetto, W. Wietrzynski, S. Albert, D. Larivière,
E. Fourmentin, S. Pfeffer, J. Ortiz, W. Baumeister, T. Peng, B.D. Engel, C. Kervrann

Membranes		M1	M2	M3	M4	M5	M6	M7	M8	M9	Global
Mono-class PSII	F_1 -score	0.557	0.465	0.533	0.625	0.789	0.516	0.571	0.588	0.286	0.566
	Precision	0.710	0.476	0.800	0.625	0.789	0.727	0.556	0.500	0.400	0.643
	Recall	0.458	0.455	0.400	0.625	0.789	0.400	0.588	0.714	0.222	0.505
Multi-class PSII	F_1 -score	0.632	0.582	0.714	0.737	0.696	0.250	0.562	0.182	0.571	0.619
	Precision	0.638	0.485	0.625	0.636	0.593	0.750	0.600	0.250	0.800	0.601
	Recall	0.625	0.727	0.833	0.875	0.842	0.150	0.529	0.143	0.444	0.638
Template Matching PSII	F_1 -score	0.400	0.200	0.355	0.258	0.350	0.279	0.318	0.000	0.286	0.313
	Precision	0.556	0.375	0.344	0.571	0.333	0.261	0.259	0.000	0.400	0.353
	Recall	0.312	0.136	0.367	0.167	0.368	0.300	0.412	0.000	0.471	0.281

Table 1: Comparison of F_1 -scores computed from detected particles in the test tomogram, for template matching (508 particles) and the mono-class (470 particles) and multi-class (508 particles) (DeepFinder) strategies applied to detect PSII membrane complexes embedded within native thylakoid membranes. For comparison reasons, we chose the threshold for template matching hits to match the number of picks for the multi-class approach. The scores are measured after masking the picks to different membranes (M1, M2...) of the test tomogram. These membranes vary in resolution and in number of PSII complexes they host.

		Computing time	Hardware
Stage I	Target generation, sphere-based	3sec / tomogram	2.9GHz Intel Core i7
	Target generation, shape-based	35min (low resolution subtom. averaging) + 3sec / tomogram	32-core CPU cluster
	Training (until convergence)	15 hours	Tesla K80 GPU
Stage II	Analysis		
	Step 1: segmentation	20min / tomogram	Tesla K80 GPU
	Step 2: clustering	5min / tomogram	2.9GHz Intel Core i7

Table 2: Computing times obtained for processing Dataset #2, composed of 57 tomograms of size $928 \times 928 \times 464$ voxels.




Figure 6: DeepFinder graphical user interface. (A) Training interface composed of a first window for parametrizing the procedure and a second window for displaying the training metrics in real-time. (B) Segmentation interface which also opens a data visualization tool. This tool allows one to explore the tomogram with superimposed segmentations. In addition, DeepFinder also incorporates interfaces for tomogram annotation, target generation and clustering (see the documentation at <https://gitlab.inria.fr/serpico/deep-finder> for more information).




Figure 7: Two workflows for macromolecule localization in cryo-ET. (A) Conventional processing pipeline based on template matching. (B) DeepFinder (analysis stage): a multi-class approach able to localize several particles of different macromolecular species in one pass. [The present figure highlights why DeepFinder is more computationally efficient than template matching. Indeed, unlike DeepFinder, the template matching pipeline needs to be repeated for each class. Also, DeepFinder does not need subsequent subtomogram classification to achieve good performance, as demonstrated in the main text \(see Results\).](#)




Figure 8: **U-net architecture used in DeepFinder.** The architecture adopts the encoder-decoder paradigm, which allows to produce an output volume having the same size as the input volume. Each green box represents a convolutional layer. The number of filters n and the filter size s is labeled as $n \times (s \times s \times s)$. All convolutional layers are followed by a ReLU activation function, except the last layer, which uses a soft-max function. The up-sampling is achieved with up-convolutions (also called “backward-convolution”). Combining feature maps from different scales is performed by concatenation along channel dimension. In the end, the total number of architecture parameters is approximately 903k. More precisely, this number depends slightly on N_{cl} , the number of classes: $902928 + N_{cl} \times 33$.




Figure 9: **DeepFinder handles ice contamination on the lamella surface.** (A) Tomogram slice depicting the border of a FIB-milled lamella. The lamella contains a *Chlamydomonas reinhardtii* cell, with a lamella border suffering from ice contamination. (B) Tomogram slice with superimposed DeepFinder segmentation. Most of the ice contamination artifacts have been correctly classified as "background". Nonetheless, some miss-classifications exist, as can be observed in the zoomed-in boxes (in dashed red). In boxes 1 and 2, DeepFinder confuses some artifacts with membranes, and some features are wrongly classified as membrane-bound ribosomes. Such miss-classifications can be filtered out, either by masking out the outside of the lamella, or by ejecting segmented object that are too small (using the "cluster size" attribute given by the clustering step of the DeepFinder analysis stage).




Figure 10: Evolution of F_1 -scores with respect to sizes of the training sets (number of tomograms) on the synthetic *SHREC'19* dataset (12 classes). This figure gives an estimation of the amount of annotated data needed to identify macromolecules. This amount depends on the size of the target macromolecule: the smaller the target, the more annotations are required. Each tomogram contains in average 208 macromolecules per class. The macromolecules have been categorized into 4 groups (large, medium, small and tiny).




Figure 11: Evolution of F_1 -score with respect to the sizes of the training sets on real cryo-ET dataset #2, *Chlamydomonas reinhardtii* (3 classes). The F_1 -scores are obtained by comparing the membrane-bound ribosomes (*mb-ribo* particles) found by DeepFinder to expert annotations. In a similar fashion to Figure 10, this figure provides an estimate of the quantity of training data required to achieve a competitive result. It appears that this quantity is 1400 ribosomes (9 tomograms), which is a typical size for a cryo-ET dataset. On first glance, this estimate seems to contradict the estimates in Figure 12: the numbers do not coincide (the curve labeled "Large" estimates that quantity at 208 particles). Note that *SHREC'19* is a synthetic dataset, composed of 12 classes. Here, we are dealing with a real dataset consisting of 3 classes (membrane, membrane-bound ribosome and cytosolic ribosome). It appears that having a larger number of classes enables the use of smaller training sets. On the other hand, the case of real data is more difficult, notably because of the presence of "label noise" (errors due to the annotation pipeline) and other sources of signal corruption such as the missing wedge, the CTF and the low SNR.




Figure 12: Analysis of localization of *mb-ribo* and *ct-ribo* particles in a *Chlamydomonas reinhardtii* cell (Dataset #2). (A) 3D voxel-wise classification of experimental tomograms obtained by DeepFinder. Visual inspection confirms that most identified membrane-bound ribosomes (in blue) are close to the cell membranes (in gray), while the cytosolic ribosomes are not. (B) Histograms of distances for *mb-ribo* particles (left) and *ct-ribo* particles (right) with respect to the nearest cell membrane computed from the test set. The histogram mode maximum for *mb-ribo* particles is located at $136.8\text{\AA}$, which corresponds to the ribosome radius.




Figure 13: Quantitative analysis of overlap with expert annotations (Dataset #2). We varied the thresholds of template matching (left) and DeepFinder (right) to compute the *Recall*, *Precision* and F_1 -score curves. The threshold parameter for template matching is the constrained correlation coefficient, and for DeepFinder is the cluster size, which corresponds to the macromolecule volume (in voxels). Template matching and DeepFinder have good *Recall* values but template matching has a lower *Precision* than DeepFinder. This suggests that template matching can be recommended to select many candidates, but a time-consuming post-classification is required to improve *Precision*. DeepFinder has much higher *Precision* values, which confirms the results from the synthetic dataset (SHREC'19 challenge).




Figure 14: Analysis of consensus decisions and overlap sets (Dataset #2). (A) The central diagram represents the overlap between the *mb-ribo* sets S_{DF} (found by DeepFinder) and S_E (annotated by expert). Thus, $S_E \cap S_{DF}$ is the subset of *mb-ribo* particles found by both DeepFinder and the experts, $S_{DF} \setminus S_E$ is the subset of *mb-ribos* found by DeepFinder only (and missed by the experts), and $S_E \setminus S_{DF}$ is the subset of *mb-ribo* particles found by the experts only (and missed by DeepFinder). The origin of red arrows pointing to the subtomogram averages A_{DF} , A_{DF}^{ref} , A_E , A_E^{ref} indicate the particle subsets used to compute the averages. A ribosome density is clearly visible in A_{DF} , therefore one can safely assume that the FP rate in $S_{DF} \setminus S_E$ is low. (B) The subtomogram averages A_{DF}^o and A_E^o have been computed using subtomograms sampled from random positions. These averages serve to estimate a lower bound for the FSC curve. The correlation values equal or below this bound are considered "noise" values, and are caused by alignment bias⁶⁷. (C) The averages A_{DF}^{ref} and A_E^{ref} have both led to a higher resolution than A_{DF} and A_E , implying that the *mb-ribo* particles in the set $S_{DF} \setminus S_E$ and in the set $S_E \setminus S_{DF}$ are more heterogenous than the *mb-ribo* particles in the set $S_{DF} \cap S_E$. Also, A_{DF} and A_E have led to a higher resolution than A_{DF}^o and A_E^o , meaning that the impact of alignment bias is not significant.




Figure 15: Analysis of localization consensus between DeepFinder and experts. (A) Two tomogram slice ROIs depicting *Chlamydomonas reinhardtii* cells. (B) Membrane-bound ribosomes mapped into the ROIs. In blue, the ribosomes found by DeepFinder but missed by the experts ($S_{DF} \setminus S_E$). In yellow, the ribosomes found by the experts but missed by DeepFinder ($S_E \setminus S_{DF}$). The ribosomes found by both DeepFinder and the experts ($S_{DF} \cap S_E$) are identified in green. As expected, members of $S_{DF} \cap S_E$ constitute the majority of identified ribosomes. Members of $S_{DF} \setminus S_E$ have a tendency of being at locations where the membrane has less contrast (B, left) or where neighboring ribosomes are close (B, right). Members of $S_E \setminus S_{DF}$, which were obtained with the expert pipeline (template matching and CPCA clustering), may also be located at positions where membrane contrast is low (B, left). Nevertheless, it appears that this pipeline has a tendency of confusing membrane-bound and cytosolic ribosomes. The proximity of ice-contamination (B, right) also seems to be a factor responsible for miss-classifications.




Figure 16: The generalization potential of DeepFinder on P19 cells. DeepFinder has been trained on the *Chlamydomonas* (algae) dataset and applied on a tomogram (EMD-10439, <https://www.ebi.ac.uk/pdbe/entry/emdb/EMD-10439>) of P19 cells (mouse). Although the ribosome has a different structure for the two species for a given voxel size (13.68Å), the structures are similar enough for DeepFinder to identify and localize *mb-ribo* particles in a P19 cell. (A) Tomographic slice with both the superimposed segmented cell membrane (gray) and *mb-ribo* particles (blue). (B) Average density of *mb-ribo* particles. (C) Histograms of distances of *mb-ribo* particles with respect to the nearest cell membrane. The histogram mode maximum is located at 136.8Å, which corresponds to the ribosome radius.




Figure 17: **Evolution of loss and F_1 -score with respect to training iterations (Dataset #2).** (A) The loss, which quantifies the segmentation quality, is computed for the training set, as well as for the validation set. Comparing both curves allows assessment of the generalization capabilities of our 3D CNN method. The curves for both sets should ideally overlap, otherwise it indicates overfitting (the network memorizes trained samples instead of learning discriminating features). One epoch equals to 100 training iterations. (B) The F_1 -score, which quantifies the localization performance, computed on the test set. The time axis has been obtained using a Tesla K80 GPU. The curve indicates that competitive particle picking results are obtained after 20 epochs, or 4.3 hours with the required GPU.

Appendix A

Description of Template Matching method

We provide herein a short description of the Template Matching method (see PyTom implementation: <http://strubi.chem.uu.nl/pytom/doc/pytom/tutorial.html>) used in our experiments and SHREC'20 challenge³⁵, and explained in detail in (Hrabe et al., 2012).

Template matching correlates a structural model of the target molecule (template) with a tomogram. The voxels corresponding to high cross-correlation coefficients are potential locations of the molecule of interest. In (Hrabe et al., 2012), the authors utilize the Local Constrained Cross-Correlation (LCCC) for localizing candidate positions and apply Fourier Transform for fast computation. The computation of LCCC involves a binary mask or a weighting mask. In PyTom, non-symmetrical masks are supported and the Euler angle space is uniformly sampled to improve computational efficiency.

As described in the SHREC'20 challenge³⁵, the templates were modulated in the frequency domain using a standard CTF curve at 3 μm defocus. Frequencies beyond the first CTF-zero were set to 0. Spherical template masks with Gaussian smoothed edges based on the thresholded electron density were used for normalization for the cross-correlation value. We selected the best candidates with the highest cross-correlation score for each class. Then, overlapping candidates were filtered out. Two particles were considered as overlapping if the distance between their centers was smaller than the sum of their radii.

(Hrabe 2012) Hrabe T., Chen Y., Pfeffer S., Cuellar L.K., Mangold A.-V., Förster F. PyTom: a Python-based toolbox for localization of macromolecules in cryo-electron tomograms and subtomogram analysis. *J. Struct. Biol.*, 178(2):177–188, 2012.