

DRPNET - AUTOMATED PARTICLE PICKING IN CRYO-ELECTRON MICROGRAPHS USING DEEP REGRESSION

A PREPRINT

Nguyen P. Nguyen

Dept. of Electrical Engineering and Computer Science
University of Missouri
Columbia, MO, USA
npntz3@mail.missouri.edu

Ilker Ersoy

MU Informatics Institute
University of Missouri
Columbia, MO, USA
ersoyi@health.missouri.edu

Jacob Gotberg

Research Computing Support Services
University of Missouri
Columbia, MO, USA
gotbergj@missouri.edu

Filiz Bunyak

Dept. of Electrical Engineering and Computer Science
University of Missouri
Columbia, MO, USA
bunyak@missouri.edu

Tommi White

Dept. of Biochemistry and Electron Microscopy Core
University of Missouri
Columbia, MO, USA
whiteto@missouri.edu

April 23, 2019

ABSTRACT

Selection of individual protein particles in cryo-electron micrographs is an important step in single particle analysis. In this study, we developed a deep learning-based method to automatically detect particle centers from cryoEM micrographs. This is a challenging task because of the low signal-to-noise ratio of cryoEM micrographs and the size, shape, and grayscale-level differences of particles. We propose a Deep Regression Picker Network (DRPnet) containing a cascade of two convolutional neural networks (CNNs) for automated detection of particles in cryo-electron micrographs. Particles are detected by the first CNN and further refined (classified) to reduce false particle detections by the second CNN. The first network, a fully convolutional regression network (FCRN), maps the particle image to a continuous distance map that acts like a probability density function of particle centers. This approach, entitled “DRPnet”, is simple but very effective in recognizing different grayscale patterns corresponding to 2D views of 3D particles. Our experiments showed that DRPnet’s first CNN pretrained with one dataset can be used to detect particles from a different datasets without retraining. The performance of this network can be further improved by re-training the network using specific particle datasets. The second network, a classification convolutional neural network, is used to refine detection results by identifying false detections. The proposed fully automated “deep regression” system, DRPnet, pretrained with TRPV1 (EMPIAR-10005) [1], and tested on β -galactosidase (EMPIAR-10017) [2] and β -galactosidase (EMPIAR-10061) [3], was then compared to RELION’s interactive particle picking. Preliminary experiments resulted in comparable or better particle picking performance with drastically reduced user interactions and improved processing time.

Keywords convolutional neural networks · regression · deep learning · electron microscopy · image segmentation · particle picking · cryoEM · autopicking · single particle reconstruction · single particle analysis · 3D reconstruction

1 Introduction

Although high resolution 3D protein structure determination via single particle analysis (also known as single particle reconstruction) using cryo-electron microscopy (cryoEM) is becoming more widely used, it still remains a challenging technique because the resulting micrographs, especially those containing high resolution information, are noisy, usually with very low contrast intensity. In these cryoEM experiments, a purified, homogeneous protein is vitreously frozen in a thin film of solution to form a glass-like ice, which is then imaged under cryogenic temperatures (-170 °C) in a transmission electron microscope (TEM) [4]. Many cryoEM micrographs are collected, with protein “particles” imaged in different orientations in 2D. Many 2D particle views (also known as “projections”) are collected and used to reconstruct an atomic 3D model [5] by iterative alignment, classification and averaging. The resulting micrographs are noisy; the low signal is due to a number of causes - microscope-related aberrations, low doses of electron exposure applied to the radiation sensitive protein specimen [6], movement of the specimen upon imaging [7], and due to the process of high resolution image formation in the TEM [8] [9]. Thus, the protein particles of interest are sometimes challenging to identify in these 2D cryoEM micrographs, especially certain orientations of the particles. Several software solutions have been developed to address this problem. These solutions require a large number of particles to estimate accurately the relative angular orientations of these protein particles in 3D, which is used to create 3D reconstructions of averaged protein particle structures.

One solution, RELION [10], uses the Bayesian approach and expectation maximization algorithm to solve for protein structures. RELION’s workflow allows users to manually identify and select (“picking”) particles in 2D cryoEM micrographs; then aligns, classifies, and averages those particles to make 2D class averages. Suitable 2D class averages are manually selected as templates, which are then used to perform automated particle selection (“autopicking”) by assessing image blocks’ correlation to the template particles [2]. EMAN2 [11, 12] is another popular solution to create 3D reconstructions from cryoEM micrographs. EMAN offers users three methods to pick particles from manual samples: local search, region search, and a simple deep neural network for pattern recognition based on THEANO [13]. One recent software called cryoSPARC [14] has a very promising approach for picking particles whereby the algorithm searches for a global optimal model from the regularized likelihood function instead of using an expectation maximization algorithm.

Although these software solutions have different methods for the particle picking process, the common point is that template particles (a set of good particles selected by experts from noisy cryo-EM micrographs) needs to exist. This manual picking task is time consuming, and often an expert is not available. After obtaining good 2D class average templates, these software packages rely on two main methods to automatically select particles: template matching by cross-correlation or pattern recognition by a simple deep neural network. Template matching is sensitive to noise and may obtain strong bias [2], whereas neural networks are normally less sensitive and usually perform better. Unfortunately, simple neural networks are not robust enough to handle all the complicated noise patterns or the illumination/transmission variations in cryoEM micrographs. Also, necessary particle orientations cannot be usually obtained with the stringent constraints required (such as high picking threshold) due to the low contrast nature of cryoEM data. If we implement low picking thresholds, it is likely that these programs tend to pick not only particles, but also noise patterns known as false positives [15]. Unfortunately, one has to compromise to have either picking accuracy or false positive avoidance. Development of machine learning, specifically deep learning, holds much promise to improve and automate biomedical image analysis [16, 17, 18].

For microscopic image analysis, [19] many machine learning approaches have been utilized from support vector machines (SVM) to convolutional neural networks (CNN). A recent study of Shin *et al.* [20] has employed deep learning models to learn semantics in MRI scans, and to extract features to detect different organs. Xie [21] proposed a novel deep neural network for robust nucleus localization, where, instead of using a pixel-wise classifier or a regressor, they combined CNN with a nonlinear voting transformation. Inspired by the success of these examples, we sought to apply these deep learning methods to cryo-EM image analysis, specifically particle picking, which is the most tedious step in most cryoEM image analysis workflows.

In this paper, we introduce a Deep Regression Picker Network (DRPnet) to predict a probability map of particle centers and then classify those candidates to pick particles. The paper is organized as follows: in Section 2, we review recent advanced methods used to detect particles in microscopy images. Then, we explain the methods we used in this study and how DRPnet works in Section 3. In the final section, we present our experimental particle detection results for EMPIAR-deposited datasets, and 3D reconstructions obtained using these particles. Our major finding is that DRPnet finds particles with greater angular orientation coverage, and the certainty of the angular assignment is improved, as compared to those generated with RELION[10] autopicking. This leads to improvements in resolution as seen in results obtained from 3D reconstruction using RELION software.

2 Related Work

DeepPicker [22] is a fully automated particle picking program based on deep neural networks. Before DeepPicker, other packages such as RELION and EMAN2 used semi-automated algorithms and workflows to collect particles. DeepPicker slides a window, with the same size as a specified particle (“box”), across a micrograph with a default step size of four pixels to collect candidate patches. All image patches are normalized and fed into a convolutional neural network to determine whether they belong to a qualified class of particles or not. In the fully automated mode, DeepPicker has a pretrained network to pick particles for the first time as positive training samples. The negative samples are randomly selected far from the positive samples using a spatial distance threshold. All of those samples are then used to train a fresh convolutional network to pick the final particles. However, in our experiment with β -galactosidase dataset (EMPIAR-10017), the pretrained mode did not obtain a satisfactory result. Next, we worked through DeepPicker’s training mode. In training mode, DeepPicker can adopt particles from RELION (using a manually picked .star file) to train the DeepPicker network. This second approach provided suitable results with higher particle picking accuracy and better capture of different particle orientations. In summary, DeepPicker worked quite well in the semi-automated mode. However, in fully automated mode, because of its simple four layer network configuration, DeepPicker was not powerful enough to acquire appropriate results with low contrast particles in EMPIAR-10061.

Zhu *et al.* [15] introduced a convolutional neural network (CNN) composed of eight layers called DeepEM to pick particles. DeepEM has three convolutional layers, one fully connected layer, and it requires data augmentation for both learning process and recognition process. The experiment was performed on a keyhole limpet haemocyanin (KLH) dataset and showed improved detection performance avoiding unwanted particles and non-particles. This framework, however, was only tested on one dataset (KLH; collected with a Tecnai F30 Twin on a Gatan Ultrascan 4000 CCD camera). If one required a fully automated detection framework, this simple network needs to be more versatile to adapt to a variety of differently shaped particles.

Xiao and Yang [23] employed Fast-RCNN [24] to select particles from micrographs. Instead of using a default region proposal function for Fast-RCNN, they made use of a sliding window to provide candidate regions for the detection CNN. This approach is equivalent to regular CNN architectures as the one used in DeepPicker. The strength of this method is that the network was trained with three classes instead of two, in order to discriminate false positive samples.

Heimowitz *et al.* [25] introduced a new picker called APPLE picker based on cross-correlation. It automatically selects reference windows to compare and eliminate the need to prepare manual templates. Compared to RELION’s manual template generation, APPLE picker works faster, accurately, and automatically. However, APPLE picker did not work well with high concentrations and hollow cylindrical particles (KLH, GroEL).

crYOLO [26] utilized the first generation of YOLO network [27] to detect particles. YOLO is based on a fixed-grid regression method. Compared to other region-based methods, this fixed-grid helps YOLO to work faster in real-time [28]. However, YOLO have detection problems with small objects and objects with unusual aspect ratios, and also suffers from high localization errors [28]. crYOLO performed detection very fast, up to six micrographs per second with image dimensions of 1024×1024 pixels [26]. crYOLO attained higher accuracy (recall, precision, and AUC) compared to the original YOLO model when working with small objects as particles in those test datasets. Despite modifications to the original YOLO architecture, the experiments do not mention how the crYOLO architecture performs using particles of other sizes and aspect ratios.

WARP package [29] uses a picking function based on BoxNet [30] for cryo-EM data processing. This deep neural network was trained with multiple EMPIAR raw and simulated data. As a result, it can pick many types of particles and can mask out the artifact regions. WARP supports network re-training with custom data.

TOPAZ [31] is another particle picking program using deep learning. Similar to DeepPicker [22], it examines micrograph patches and uses non-maxima suppression to select only the patches with highest scores as the most likely particle instances. One major difference between these approaches is that TOPAZ was trained with positive and unlabeled samples, instead of positive and negative samples.

Since a robust and flexible deep learning tool had not been developed to date for cryoEM particle picking, we set forth to develop a deep learning-based particle picking system that would work on multiple types of data (vitreous ice and negative stain), with various particle sizes and aspect ratios, and collected using different detectors (direct detection and CCD) with different defocus ranges. Below, we present our work on “Deep Regression” particle picking (Deep Regression Picker Network; DRPnet, Figure 1). The proposed system works on multiple types of data (tested on four protein datasets), from various detectors, with improved speeds (testing time is approximately ~ 6 sec/micrograph on Nvidia GTX 1080 GPU with 8GB memory and SSD storage) compared to manual template selection (~ 30 min to select 1000 particles) required by RELION. The outputs of DRPnet were compared with the outputs of the RELION autopicking and produced final maps with similar resolution.

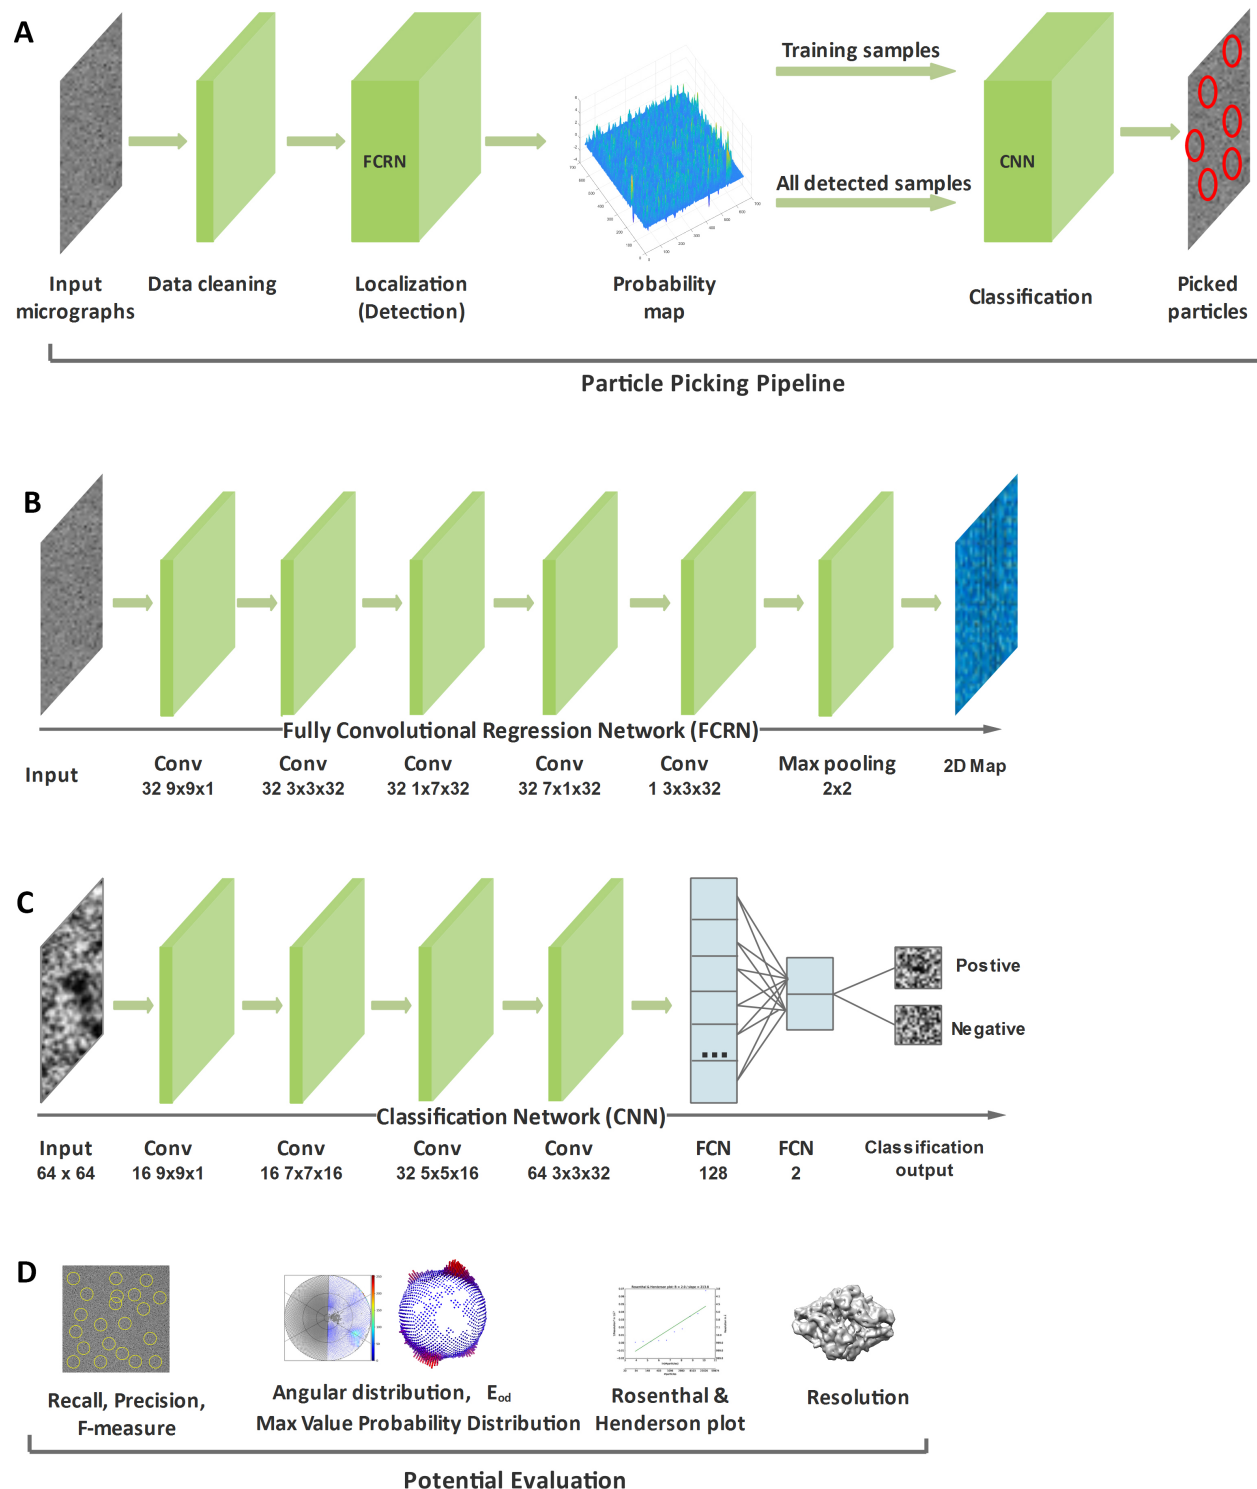


Figure 1: The cascade architecture of the Deep Regression Picker Network (DRPnet). (A) Overview of particle picking pipeline of DRPnet, from input micrographs to picked particles. (B) Architecture of Fully Convolutional Regression Network (FCRN) for initial particle detection. (C) Architecture of Convolutional Neural Network (CNN) for refining classification. (D) Potential Evaluation metrics.

3 Material and Methods

3.1 Development of DRPNet Particle Picking Pipeline

Particle picking involves localization (detection) and classification processes. In classical computer vision, object detection and classification processes often rely on carefully hand-crafted image features and descriptors such as HOG (Histogram of Oriented Gradients) [32], SIFT (Scale-Invariant Feature Transform) [33], or SURF (Speeded-Up Robust Features) [34] etc. that are extracted from image patches. These descriptors are then fed to supervised machine learning models such as SVM (support vector machine), random forests, and neural networks to cluster or to classify those feature vectors. Given an input image, a scanning window slides over the image and independently classifies each image patch based on features extracted from the patch. The success of these systems heavily depends on developer selected features that may not capture complexities of the underlying visual patterns or adapt to new patterns.

On the other side, deep learning approaches rely on data itself to identify discriminative features to perform the given task. Convolutional neural networks (CNN), a class of deep learning methods heavily used for image analysis, extract features over multiple convolutional layers with different sizes and different numbers of filters whose coefficients are learned during training. Earlier layers of convolutional networks learn deep features, while later fully connected layers perform the classification or regression tasks. Deep detection approaches can be coarsely grouped as single-stage and two-stage. Single-stage detectors such as YOLO [35] and SSD [36] rely on regular dense sampling of objects, scales, and aspect ratios and perform detection and classification in one step. Two-stage detectors such as FasterRCNN [37] and Mask-RCNN [38] first produce a sparse set of candidates (region proposals). These candidates are then classified into a number of classes, foreground versus background in our case. Faster-RCNN and Mask-RCNN have shown great performance in detection and classification tasks on image datasets such as Pascal VOC [39] and MS COCO [40]. However, these models were originally designed to work with color (RGB, HSV, CYMK) images and larger objects with rich color, texture, and shape features. Unlike color images and objects in these datasets, electron micrographs are grayscale, low contrast, and may contain noise and artifacts (support-film edges and surface ice contamination). We required more versatile network than Faster-RCNN's Region Proposal Network (RPN) [37] to perform automated particle detection in cryoEM micrographs. In order to overcome the problem with RPN, we propose the Deep Regression Picker Network (DRPnet, Figure 1), which is based on the blob detection concept [41]. After the initial training with one dataset, the proposed DRPnet can work with multiple particles having various sizes and shapes from different datasets without need for retraining. We also realized that using scanning window and a single network for both localization and classification limits the ability to recognize patterns and/or discriminate different types of false positive patterns. For this reason, our particle picking pipeline (Figure 1A) includes two steps: localization (detection) relying on a fully convolutional regression network (FCRN) (Figure 1B) and classification relying on a classification convolutional neural network (Figure 1C). Although we used a supervised CNN model to classify true versus false positives, training data for this classification network were collected in an unsupervised manner. To identify true versus false positives in a semi-supervised manner, we used FCRN's measure of confidence where high level corresponded to a true positive and low level corresponded to a true negative. The initial network was trained using the EMPIAR-10005 dataset [1]. This scheme allows DRPnet to operate on other datasets automatically from detection through classification.

Our particle picking pipeline also included a micrograph preprocessing step to enhance contrast and to correct transmission/illumination artifacts in order to ensure optimal performance by both networks. Picked particles were then fed to RELION (version 3.0) software for 3D reconstruction. Processing steps involved in the proposed automatic particle picking pipeline are described in the following subsections.

3.2 Data Preprocessing

3.2.1 Data Cleaning: Illumination/Transmission Correction

Illumination (beam centering) and/or electron transmission (transmitted electron signal) can misguide DRPnet's detection because particle picking significantly relies on intensity of grayscale images. To obtain the best DRPnet detection performance, it was necessary to reduce the influence of illumination/transmission variations. We compute local averages of illumination by applying a low-pass filter with a very large kernel. The corrected intensity value of each pixel is obtained by subtracting local and adding global intensity averages, as in Equation 1:

$$I_{new} = I - I_{\sigma} + \frac{1}{n} \sum (I_{\sigma}) \quad (1)$$

where I is the original image, n is the number of pixels in the image, and I_{σ} is image smoothed with a Gaussian filter of sigma σ , set heuristically according to particle size. Very low contrast and extremely high noise characteristics, make electron micrographs more challenging compared to typical color or grayscale images.

3.2.2 Distance Transform

Distance transform of a binary image is a grayscale map where each pixel's value represents the distance to the closest boundary. It is calculated at each pixel by the distance to the nearest background pixel. In this study, we applied distance transform to the binary mask of particles to produce a representation of a probability map of the particle centers (Figure 2B).

3.3 Deep learning

3.3.1 Fully Convolutional Regression Network (FCRN)

To address the challenges of automated particle picking, we have treated it as a blob detection problem where each particle is roughly a convex blob with texture different than the background (Figure 2A). Additionally, dimensions and shapes of particles are relatively similar. Given a rough size estimate, our goal is to locate particle centers. For the purposes of this study, segmentation of particle regions from image background is unnecessary. Our proposed model is a Fully Convolutional Regression Network (FCRN), trained to predict particle centers by producing a likelihood map where local maxima correspond to the particle centers (Figure 2). Given the ground truth particle centers (as provided by the datasets in EMPIAR, Figure 2B) and a particle size estimate, first, binary particle versus background masks were produced for the electron micrographs. Distance transforms of those binary masks are then used as training labels for the proposed FCRN model. Distance transform for the background pixels are set to zero. Distance transform for the foreground regions (particles) increase from particle boundaries towards particle centers, reaching their local maxima at the particle centers. Therefore, a continuous regression model (FCRN), rather than discrete classification models, is used to localize particle centers. Because of its continuous nature, this model enables localization of individual particles within dense clusters, and ensures better scale adaptation, which allows detection of different sized particles without need for retraining the network. This is an important feature missing in the existing algorithms reported to date. Furthermore, since the proposed network is built without fully connected layers, any size micrographs can be used as input to the network. The proposed FCRN has seven layers (Figure 1B), one input, five convolutional, and one max-pooling layers [42]. With its simple structure, our approach has smaller computational cost and training data needs, compared to more complex deep learning models such as Fast-RCNN.

The proposed network was implemented using Matlab deep learning toolbox. The network was trained with image patches from raw micrographs and associated training labels, centered and cropped on particle regions. These image patches are of the same size as the FCRN input layer (Figure 2B). The number of patches extracted from each image is equal to the number of particles in those images. The network is trained using Adam optimizer [43]. Since our network does not have fully connected layers, it does not require fixed sizes of input images. In the testing stage, we directly applied the trained FCRN model to different-sized input images (Figure 2D) to predict associated particle center probability maps (Figure 2E).

3.3.2 Classification Network

As shown in the working pipeline in Figure 1A, the FCRN probability map output indicates locations of all detected particles. The second convolutional neural network (CNN) will classify these particles for the final keep-or-delete selection process. We extracted two subsets from the particles detected by the first network, and trained a second two-class classification network using stochastic gradient descent optimization [42]. Positive training samples were particles detected with high confidence, further refined by removing those with low standard deviation of intensity. Negative training samples were particles with low probability in the detection map as well as particles which had a low standard deviation of intensity. Negative particles included patches of background, carbon edges, ice contaminated regions, and weak patterns. Using the second convolutional network, these particles are classified into positive particles (Figure 2D) and negative particles (Figure 2E). This second classification network learns texture patterns of true particles versus false positives caused by various artifacts, and provides a refinement to the particle detection process in an efficient unsupervised manner. This network has four convolutional layers, with each followed by a max-pooling layer [42] for feature extraction and two fully connected layers [42] to decide whether to pick or discard a particle patch using the learned features. Figure 1C illustrates this CNN architecture.

3.4 Performance Metrics

In this research, we applied multiple metrics to measure our DRPnet-based particle picking performance (Figure 1D). The first metric uses recall, precision, and F-measure [44] to evaluate how accurately our model predicts, which can only be calculated if a ground truth exists. After picking particles, we use them to reconstruct 3D maps of protein structures to measure resolution, which serves as our second metric. Finally, we analyze angular distributions of

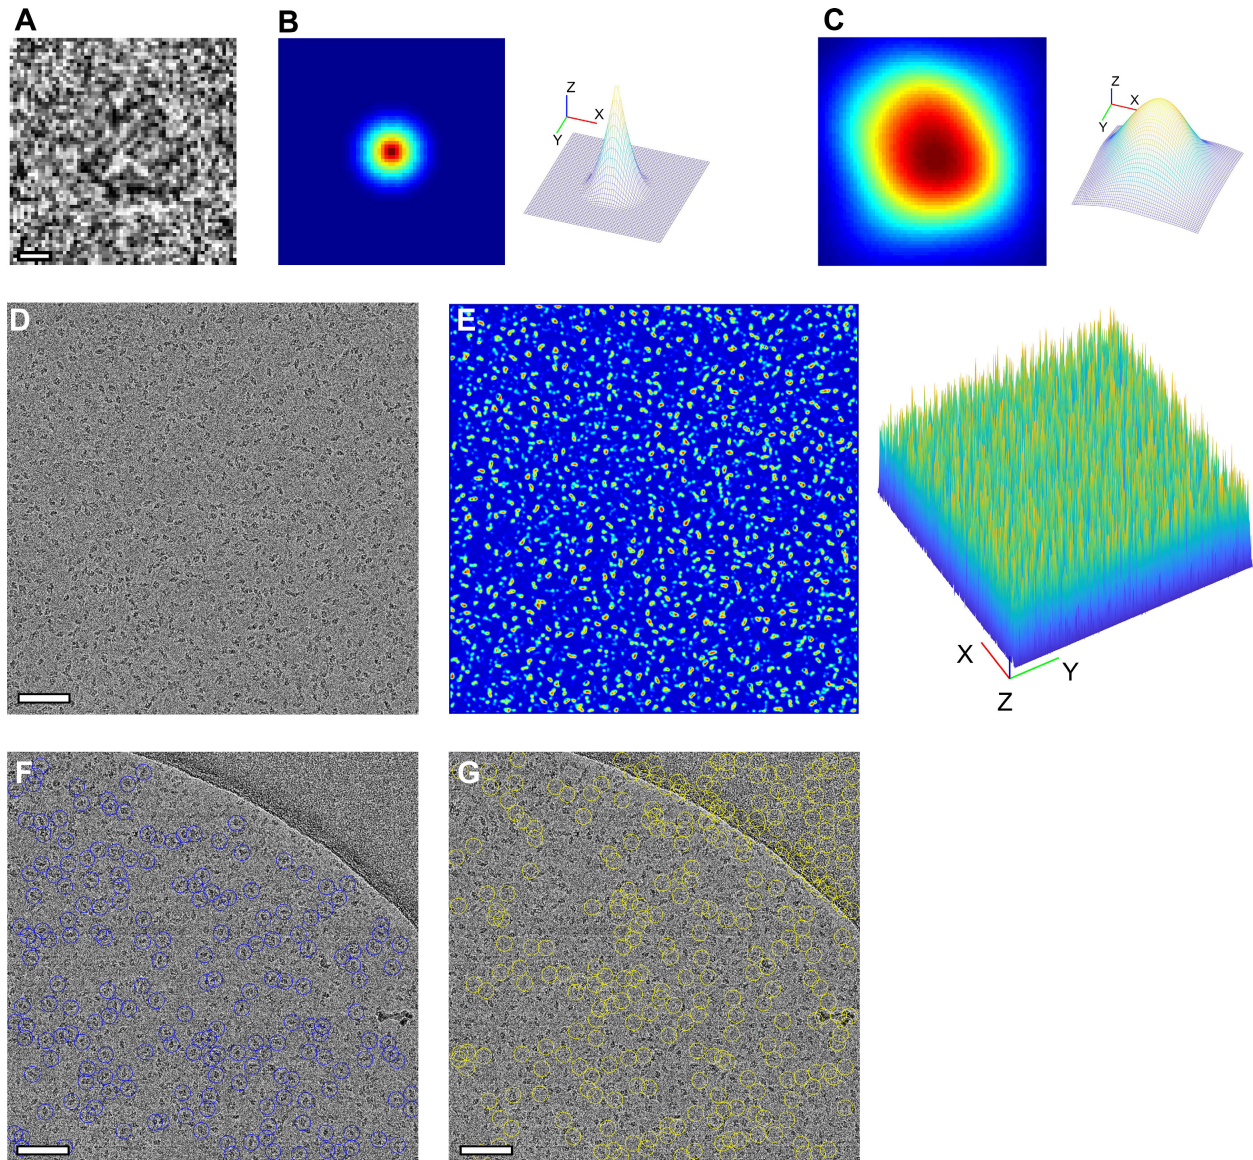


Figure 2: Training and test examples of DRPnet on cryoEM micrographs. other scale bars have length of) (A) Input particle from cryoEM micrograph of TPRV1 (scale is 50 Å) . (B) Distance transform of particle used as Ground truth and its 3D visualization. (C) Predicted map of individual particle and its 3D visualization. (D) Example of original input cryoEM micrograph (scale is 500 Å). (E) Predicted map of entire micrograph by DRPnet, and 3D visualization of predicted map. Positive (F) and negative (G) samples for training the classification network (scale is 500 Å).

those 3D reconstructions and calculate the efficiency E_{od} [45] to understand the effect of orientations of the final 3D reconstruction resolution from those particles picked by either DRPnet or RELION. We also visualize the effect of number of particles on structures' resolution by Rosenthal-Henderson Plots [46] [47].

3.4.1 Recall, Precision & F-measure

The overall detection performance of all particles are evaluated using three metrics: recall, precision, and F-measure [44]. Initially, we collected the detection results in confusion matrix which has the columns to indicate the actual fact or the ground truth, and the rows to indicate prediction results. True Positive (TP) and True Negative (TN) are two cases that prediction and ground truth match together (Table 1). Otherwise, when prediction and ground truth are different, we will get false positive (FP) if we detect more than the ground truth, or we will have false negative (FN) if we miss the ground truth cases. From these basis terms, Recall is defined as the number of true positives divided by the number of true positives plus the number of false negatives (Equation 2). Recall indicates the ability to catch all the ground truth cases. Precision is defined as the number of true positives divided by the number of true positives plus the number of false positives. Precision indicates how the prediction is relevant to the ground truth (Equation 3). In order to detect all the ground truth particles, one may decide to return all the images patches. In which case, recall would be high but precision would be low because only a small portion of the returned image patches would correspond to true particles. Conversely, if one only selects the high confident predictions, precision would be high. However, this would result in some ground truth cases to be missed, resulting in low recall scores. F-measure, computed as two times of Product of Recall and Precision, and divided by sum of Recall and Precision, to assess our models' prediction (Equation 4).

Table 1: Prediction Outcomes

	Actual Positive	Actual Negative
Predict Positive	TP	FP
Predict Negative	TN	FN

$$Recall = \frac{TP}{TP + FN} \quad (2)$$

$$Precision = \frac{TP}{TP + FP} \quad (3)$$

$$F - measure = \frac{2 * (Recall * Precision)}{Recall + Precision} \quad (4)$$

3.4.2 Resolution

The second evaluation metric we applied to evaluate the performance of picking is the resolution of 3D reconstruction models built from those picked particles. However, the optical resolution is not applicable in this case. Instead of that, we rely on the concept of signal to noise ratio, and measure the internal consistency of each dataset containing picked particles. We take advantages of some characteristics of Fourier transform to apply to our data: correlation between to volumes of 3D model can be represented by a product in Fourier space, and then we split that product into shells by radial frequency to compute Fourier Shell Correlation (FSC; Equation 5) [48, 49, 10, 50].

$$FSC(k, \Delta k) = \frac{\text{Real} \left\{ \sum_{(k, \Delta k)} F_1(K) F_2(K) \right\}}{\left\{ \sum_{(k, \Delta k)} |F_1(K)|^2 |F_2(K)|^2 \right\}^{1/2}} \quad (5)$$

K : spatial frequency vector

$k = |K|$, is the size of spatial frequency

ΔK : ring width or shell thickness

$F_1(K), F_2(K)$: Fourier transforms of two half set reconstruction

As a result, we construct a plot of Fourier Shell Correlation vs. spatial frequency. In addition, the relationship between FSC and Spectral Signal to Noise Ratio (SSNR) is explained by [51] as following:

$$SSNR = \frac{FSC}{1 - FSC}, \quad FSC = \frac{SSNR}{SSNR + 1} \quad (6)$$

FSC is a biased estimate of SSNR. For a large number of images: $\text{Variance}(SSNR) = \text{Variance}(FSC)$, and the bias is negligible. When FSC is calculated for a data set which is split into halves, the relationship is:

$$SSNR = \frac{2FSC}{1 - FSC} \quad (7)$$

With the relationship between FSC and SSNR, we use FSC to measure the internal consistency of our data. For a specific level of FSC, we take the plot to identify an equivalent spatial frequency. That frequency has a unit of 1/Angstroms which is the inverse of the second measured metric, resolution. In order to divide our data into two volumes to calculate FSC, we can select odd particles for the first volume and the even particles for the second volume after we construct the 3D model. This way is very simple but two halves of our data are not independent. Another way is splitting particles from the beginning and build two 3D models. This method is often referred as “Gold standard Fourier Shell Correlation”. RELION uses this standard to compute resolution.

A cut-off level is necessary to identify structures’ resolution. A reasonable criterion includes only Fourier information that is above the noise level. Typical cut-off level for Gold standard split is 0.143 [52] [10] (Figure 6C, 6F and 6I), and for odd-even split is 0.5 [53, 10]

$$FSC = \frac{\sum (S + N_1)(S + N_2)}{\sum (S^2 + 2SN + N^2)} \quad (8)$$

where S denotes the signal and $N = N_1 = N_2$ are noise in half datasets.

3.4.3 Rosenthal-Henderson Plots

As originally proposed by Rosenthal-Henderson [46], plotting of the inverse-squared resolution as a function of the number of particles allowed us to compare numbers of particles picked with DRPnet or RELION’s and the effect of particle numbers on 3D-reconstruction resolution [47].

3.4.4 Angular distribution

Besides two metrics above, we analyzed the angular distributions recorded in RELION .star files of 3D Refinement jobs to understand the difference in orientation of particles picked by DRPnet and RELION’s Autopick. Those distributions showed the effect of orientation on resolution and clarified the difference of final resolution of structures by DRPnet and RELION. We visualized the angular distribution by 2D scatter plots of rot-tilt angles [54], and visualized the bild files of 3D refinement jobs by 3D histograms in Chimera [55]. Furthermore, we employed efficiency number E_{od} [45] to confirm the angular distributions’ effect. By using corresponding point spread function, Naydenova and Russo [45] computed E_{od} to assess how the angular distribution contributes to the reconstruction results. An E_{od} value of 0.8–1 indicates a good orientation distribution and a uniform Fourier space coverage. Otherwise, if E_{od} is lower than 0.5, the orientation distribution is not effective to 3D reconstruction. Finally we plotted the “Max Value Probability Distribution”, an output from RELION’s 3D Refinement, where values close to one indicate that the probability distribution over all orientations and classes have converged to near-delta functions. We use this metric to provide confidence that RELION was accurately assigning orientations to the particles.

3.5 Overview of Methods

Taking cryoEM datasets available in EMPIAR [56] or locally generated negative stain micrographs, we utilized both DRPnet-particle picking or RELION (v3.0) autopicking to generate datasets with similar numbers of particles. For each EMPIAR data utilized, RELION requires generation of a manually picked particle set, usually around 1000 particles total. Particles used in our RELION comparison results were generated by randomly selecting a subset (1,000) of particles deposited in EMPIAR that contributed to the final, refined high resolution 3D reconstructions previously published. These randomly selected particles can be considered the “manually picked” particles. These randomly-selected, manually picked particles were then processed through the standard RELION pipeline (template generation by 2D classification). DRPnet particles were generated as described in Sections 3.2-3.4. Both DRPnet particles and

the RELION's autopicked particles were then processed through RELION's 2D classification to identify good 2D class averages (see Figure 2 in chapter 6 of Methods in Enzymology [57]). Both DRPnet and RELION-autopicked 2D class average templates were further 3D classified (allowing 5 possible classes) using the corresponding deposited, low-resolution (60 Å) filtered 3D reconstruction as the initial model. All good 3D classes were selected and refined to generate the final 3D reconstruction. RELION used gold standard as cut off level to calculate those structures' resolutions.

4 Results and Discussion

4.1 Datasets

To demonstrate our model's ability of working with multiple types of particles, we prepared a pretrained DRPnet model from one dataset and tested it on other datasets. We downloaded three cryoEM datasets from EMPIAR including TRPV1 (EMPIAR-10005) [1], β -galactosidase (EMPIAR-10017) [2], β -galactosidase (EMPIAR-10061) [3]. CryoEM datasets were acquired using an 300 kV accelerating voltage on FEI Titan Krios cryo-transmission electron microscopes or on a FEI Tecnai F30 Polara, whereas the negatively-stained data at 200 kV on a FEI Tecnai F30 Twin transmission electron microscope. We selected TRPV1 to train our DRPnet model. Liao *et al.* [1] collected TRPV1 cryoEM micrographs using a Gatan K2 direct electron detector with a pixel size of 1.2156 Å.

Our second test dataset is EMPIAR-10017 (β -galactosidase). Scheres [2] captured this set using FEI Falcon II camera, and a pixel size of 1.77 Å. EMPIAR-10017 contains 84 micrographs and 40,863 manually selected particles.

The third test set has 661 micrographs selected from EMPIAR-10061 (β -galactosidase) dataset. Bartesaghi, et al. [3] collected this dataset with Gatan K2 summit camera on the end of a Gatan Imaging Filter (GIF) energy filter with a pixel size of 0.3185 Å. For a summary of these datasets, please see Table 2.

Table 2: Experimental datasets from EMPIAR used for DRPnet testing

Dataset	Voltage (kV)	Camera	Type	Defocus Range (μm)	# Total Images	# Test Images	# Picked Particles
Empiar-10005 (TRPV1)	300	Gatan K2	Vitreous	1.5 - 3.5	871	661 (subset)	60,000
Empiar-10017 (β -gal)	300	FEI Falcon 2	Vitreous	1.0 - 3.0	84	84 (full set)	50,000
Empiar-10061 (β -gal)	300	Gatan K2 Summit GIF K2	Vitreous	0.6 - 3.0	1538	661 (subset)	60,000

4.2 Training on EMPIAR 10005 (TRPV1)

To train the first network of DRPnet (FCRN) for particle detection, we used a subset of TRPV1 data (EMPIAR-10005), including 50 TRPV1 micrographs and 9751 ground truth particles extracted from deposit coordinates used for high resolution 3D reconstruction. These micrographs have dimensions of 3710x3710 pixels and a pixel size of 1.2156 Å. We scaled them down by 3x to fit the scale of deposit particles' coordinates, and used a box size of 64x64 pixels to extract particle patches. Those deposit particles' coordinates were used to generate binary masks; and then the distance transform of those masks were used as ground truth to train the FCRN network. The second network of DRPnet for classification was trained after we tested the first network to detect positive samples and negative samples.

4.3 Testing on EMPIAR-deposited datasets

4.3.1 Test on Empiar-10005 (TRPV1)

In order to validate our pretrained DRPnet model with TRPV1, we tested it on a set of 661 TRPV1 micrographs, collected on a Gatan K2 direct detector. TRPV1 micrographs have low contrast, and some do not contain any particles. We set the first network (FCRN) at normal detection level ($\sigma = 9$, threshold = $2*5$) to detect particles. To discriminate true positive and false positive detections, we select classification training samples at positive level ($\sigma = 9$, threshold = $3*5$) and negative level ($\sigma = 9$, threshold = $1*5$) to train the second network. As a result, DRPnet picked 61,282 particles. For RELION's autopick, we selected $\sigma = 1$ and threshold = 0.475 to pick 61,599 particles. We extracted particles with box size 200x200 pixels without scaling to perform the subsequent jobs. Taking both sets of

picked particles forward through the RELION 3.0 pipeline, we selected 23,147 good particles picked by DRPnet and 22,830 good particles picked by RELION autopick.

Table 3: Test results (values in parentheses were obtained without imposing symmetry - C1)

Dataset	Program	# Picked Particles	# Selected Particles	# Missing Views (in total 12,288 views)	E _{od}	Resolution (Å)
Empiar-10005 (TRPV1)	DRPnet	61,282	23,147	10,164 (7,807)	0.57 (0.47)	3.92 (6.57)
	RELION	61,599	22,830	10,389 (8,591)	0.51 (0.32)	3.92 (6.75)
Empiar-10061 (β -gal)	DRPnet	60,267	49,582	9,196 (4,019)	0.82 (0.76)	2.78 (3.02)
	RELION	61,597	49,570	9,288 (4,320)	0.80 (0.74)	2.81 (3.06)
Empiar-10017 (β -gal)	DRPnet	49,604	41,914	9185 (2,250)	0.69 (0.70)	4.21 (4.72)
	RELION	49,855	40,669	9,250 (2,412)	0.69 (0.70)	4.26 (5.06)

Table 4: Recall, Precision and F-measure of Empiar-10017 (β -galactosidase)

	# Picked Particles	Recall	Precision	F-measure
DRPnet	49,604	87.73%	71.10%	78.54%
RELION	49,855	73.36%	59.88%	65.94%

In this experiment, the resolution of DRPnet structure was 3.92 Å, and equal to the resolution of RELION autopick structure which was also 3.92 Å. We studied the angular distribution to further analyze the results of resolutions. The histograms of Rot-Tilt angles (Figure 3G and 3H) and Table 3 show that the particles of RELION autopick missed more views than DRPnet's particles (gray region). This result also agreed with the efficiency numbers [45] E_{od} in Table 3. DRPnet had E_{od} of 0.57 while RELION had E_{od} of 0.51. We imposed symmetry C4 for TRPV1 structure to get these results of E_{od} and final resolution. Without imposing symmetry (C1), the values in parentheses in Table 3 show that DRPnet missed less views, and had better E_{od} and resolution than RELION autopick. Besides the orientation angles, we evaluated the Max Value Probability Distribution to see how confident RELION was of the picked particles' orientations in the 3D reconstruction, with a low Max Probability Distribution value indicating low confidence of orientation. As shown in Figure 3I, the histogram of Max Value Probability Distribution shows that DRPnet picked more particles with higher values (blue area under the curve) than RELION (red area under the curve). Those probability values of DRPnet have a mean of 0.230, and a standard deviation of 0.173 while RELION has a mean of 0.237 and a standard deviation of 0.161. Rosenthal & Henderson plot in Figure 6A shows that when the number of EMPIAR-10005 particles increase from 600 to 20,000 particles, DRPnet obtained better structure resolution than RELION's autopick, especially if that structure was not imposed symmetry (C1).

Next, we tested on two other datasets to verify that DRPnet could not only select particles with diversified shapes, but also improved orientation ranges resulting in higher resolution structures.

4.3.2 Test on Empiar-10061 (β -galactosidase)

In the second test, we picked particles from a selected 661 micrographs from EMPIAR-10061. This dataset was collected with little defocus and at high magnification, which resulted in extremely low particle contrast. These micrographs have a size of 7676x7420 pixels, and the original pixel size is 0.3187 Å. We set DRPnet to work at detection level (sigma = 1, threshold = 1*0.01) and scale factor of 1/8, and then trained classification network with positive samples as a subset of detected particles with high standard deviation (s > 0.25) and negative samples as a subset of detected particles with high standard deviation (s < 0.25). Accordingly, DRPnet picked 60,267 particles from those micrographs. For RELION's autopick, we selected sigma = 1.2 and threshold = 0.15 to pick 61,597 particles. All subsequent jobs including extraction, classification and refinement used a binning factor of 2, a box size of 384x384 pixels and a pixel size of 0.637 Å. In 3D structure refinement, we kept 49,592 good particles of DRPnet and 49,570 good particles of RELION autopick. The results showed that DRPnet structure had a resolution of 2.78 Å, while RELION structure had a resolution of 2.81 Å. Figure 4G and 4H illustrate the angular distribution of particles picked from dataset Empiar-10061 (β -galactosidase). The histograms of Rot-Tilt angles show that the particles of DRPnet (Figure 4G) missed less views than particles of RELION autopick (gray region). This result also agreed with the efficiency numbers E_{od} in Table 3. In our experiment, DRPnet had E_{od} of 0.82 while RELION had E_{od} of 0.80 with two-fold symmetry (D2). Without imposing symmetry (C1), the values in parentheses in Table 3 still show that DRPnet missed less views, and had better E_{od} and resolution than RELION autopick. Besides the angles, we evaluated the Max Value Probability

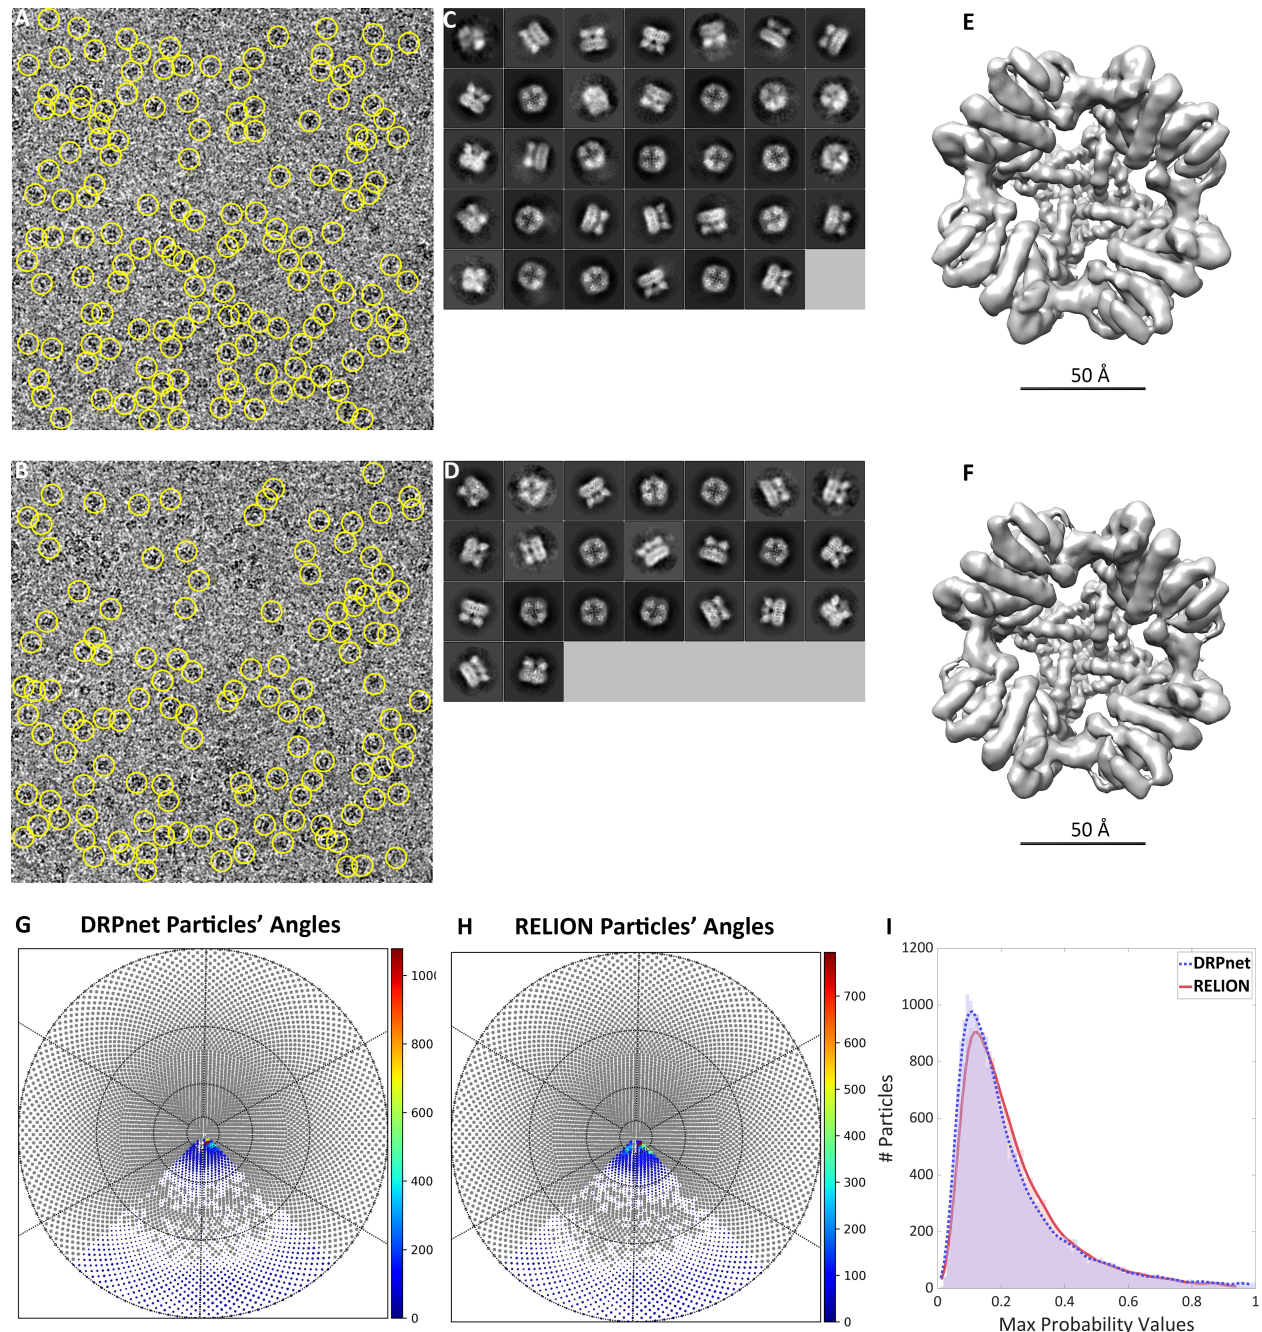


Figure 3: Test results of Empiar-10005 (TRPV1): Representative cryoEM micrograph shows picked particle output (yellow circles) by DRPnet (A) and RELION (B). All picked particles were classified using RELION to select quality 2D class averages as displayed for DRPnet (C) and RELION (D). Then those particles were used for 3D reconstruction, and 3D maps were generated from DRPnet's result (E) and RELION's result (F), are superimposed, and displayed in gray isosurface respectively. The particles' angular distributions in 3D reconstruction process are displayed in Rot-Tilt angle scatter plots by results of DRPnet (G) and RELION (H). A gray "X" symbol indicates a missing view. The last sub-figure overlays the histogram of Max Value Probability Distribution (I) of particles picked by either DRPnet (blue dashed line) or RELION (red solid line).

Distribution to see how those picked particles fit best to the reconstruction process of RELION. The histogram of Max Value Probability Distribution (Figure 4I) shows that DRPnet had more higher values than RELION. Those probability values of DRPnet have a mean of 0.339, and a standard deviation of 0.201 while RELION has a mean of 0.324 and a standard deviation of 0.205. Rosenthal & Henderson plot in Figure 6B shows that when the number of EMPIAR-10061 particles increase from 1,200 to 38,000 particles, DRPnet obtained better resolution than RELION's autopick in both cases with and without symmetry imposed.

4.3.3 Test on Empiar-10017 (β -galactosidase) data

After EMPIAR-10005 (TRPV1) and EMPIAR-10061 (β -galactosidase) test, we conducted the third test on entire EMPIAR-10017 (β -galactosidase) dataset. This dataset has an image size of 4096x4096 pixels and a pixel size of 1.77 Å. We set DRPnet to detect particles at picking level ($\sigma = 7$, threshold = $2*7$), and selected training samples for classification at positive level ($\sigma = 16$, threshold = $3*7.5$) and negative level ($\sigma = 12.5$, threshold = $2*5$) to pick 49,604 particles. On the other side, RELION's autopick picked 49,855 particles. Those particles were extracted with a box size of 200x200 pixels at original scale for 3D map reconstruction later. With a provided ground truth (particles manually picked by an expert [2]), we could evaluate detection accuracy for this dataset. Table 3 shows that our DRPnet has recall of 87.73 %, and precision of 77.10 %, and F-measure of 78.54 %. Compared to RELION autopicking, DRPnet's results are better when RELION used manually selected templates to pick particles. RELION achieved a recall of 73.36 %, precision of 59.88 %, and F-measure of 65.94 %. With the total number of picked particles around 50,000 particles from each algorithm, we run 2D classification to filter good particles. According to visual inspection, we kept only good average templates after 2D and 3D classification for 3D map reconstruction. The number of good particles corresponding to DRPnet's pick and RELION's autopick are 41,914 and 40,669, respectively. Our experiment showed that DRPnet structure had a resolution of 4.21 Å while RELION structure had a resolution of 4.26 Å. Figures 5G and 5H illustrate the angular distribution of particles picked from dataset Empiar-10017 (β -galactosidase). The histograms of Rot-Tilt angles show that the particles of DRPnet (Figure 4G) miss less views than particles of RELION autopick (gray region). In table 3, with symmetry imposing (D2), DRPnet misses 9,185 views while RELION's autopick misses 9,250 views. Without symmetry imposing (C1), the number of missing views between them is further increased to 162 views. Both DRPnet and RELION's autopick gained the same number of E_{od} of 0.69 with symmetry D2, and 0.70 without symmetry C1. Besides the view angles, the Max Value Probability Distribution shows us how those picked particles fit best to the 3D reconstruction process. The histogram of Max Value Probability Distribution (Figure 5I) indicates that DRPnet has more higher values than RELION. Those probability values of DRPnet have a mean of 0.27, and a standard deviation of 0.151 while RELION has a mean of 0.24 and a standard deviation of 0.136. Rosenthal & Henderson plot in Figure 6C shows that when the number of EMPIAR-10017 particles increase from 1,200 to 38,000 particles, DRPnet obtained better structure resolution than RELION's autopick, especially in case of no symmetry (C1).

5 Conclusions

In summary, we propose a Deep Regression Picker Network (DRPnet) and successfully demonstrate the ability of picking particles on a dataset of cryo-EM micrographs which is different from the training data. Our deep learning based network is simple and effective for automatically picking particles from 2D cryo-EM micrographs, including those with low contrast or having a large particle box size.

DRPnet has a two-part cascade of convolutional neural networks for detection and classification. It provides flexibility to retrain either or both networks if desired. Compared to other deep learning tools for particle picking in cryoEM micrographs reported to date, this cascade strategy is unique. After the initial training with TRPV1, we found it was unnecessary to retrain this network when applied to other cryoEM micrographs containing differently shaped particles (e.g. β -galactosidase) in other datasets (EMPIAR-10061). Without the need to retrain, automating 2D template preparation to reconstruct 3D protein structures by single particle analysis is feasible. If desired, one can retrain the classification CNN – which may be useful for very low contrast images or images that have artifacts, such as shapes similar to particles of interest.

Additionally, utilizing a pretrained network to pick particles significantly reduces time to obtain a particle dataset. DRPnet is very efficient for selecting a set of particles. For comparison, picking ~1000 particles manually in our laboratory from the above EMPIAR datasets took about 30 minutes per dataset. There is also additional time spent running 2D classification to select good 2D class average templates, as well as optimizing autopicking parameters. With DRPnet, it takes about 30 seconds per micrograph to pick the particles, which results in significant time savings.

When investigating traditional computer science identification metrics of recall, precision, and F-measure, we did not feel these indicators were informative of performance. Many times, existing ground truth sets are lacking due to 1) low contrast particles not contained in the “experts” ground truth, or 2) the particles composing final 3D reconstruction not




Figure 4: Test results of Empiar-10061 (β -galactosidase): The first two images are examples of picked particles by DRPnet (A) and RELION (B). All picked particles were classified using RELION to select quality 2D class averages as displayed for DRPnet (C) and RELION (D). Then those good particles were used for 3D reconstruction, and the isosurface superimposed renderings of the maps were generated from DRPnet's result (E) and RELION's result (F) respectively. The particles' angular distributions in reconstruction process are displayed in Rot-Tilt angle scatter plots by results of DRPnet (G) and RELION (H). A gray "X" symbol indicates a missing view. The last sub-figure overlays the histogram of Max Value Probability Distribution (I) of particles picked by DRPnet (blue dashed line) and RELION (red solid line).

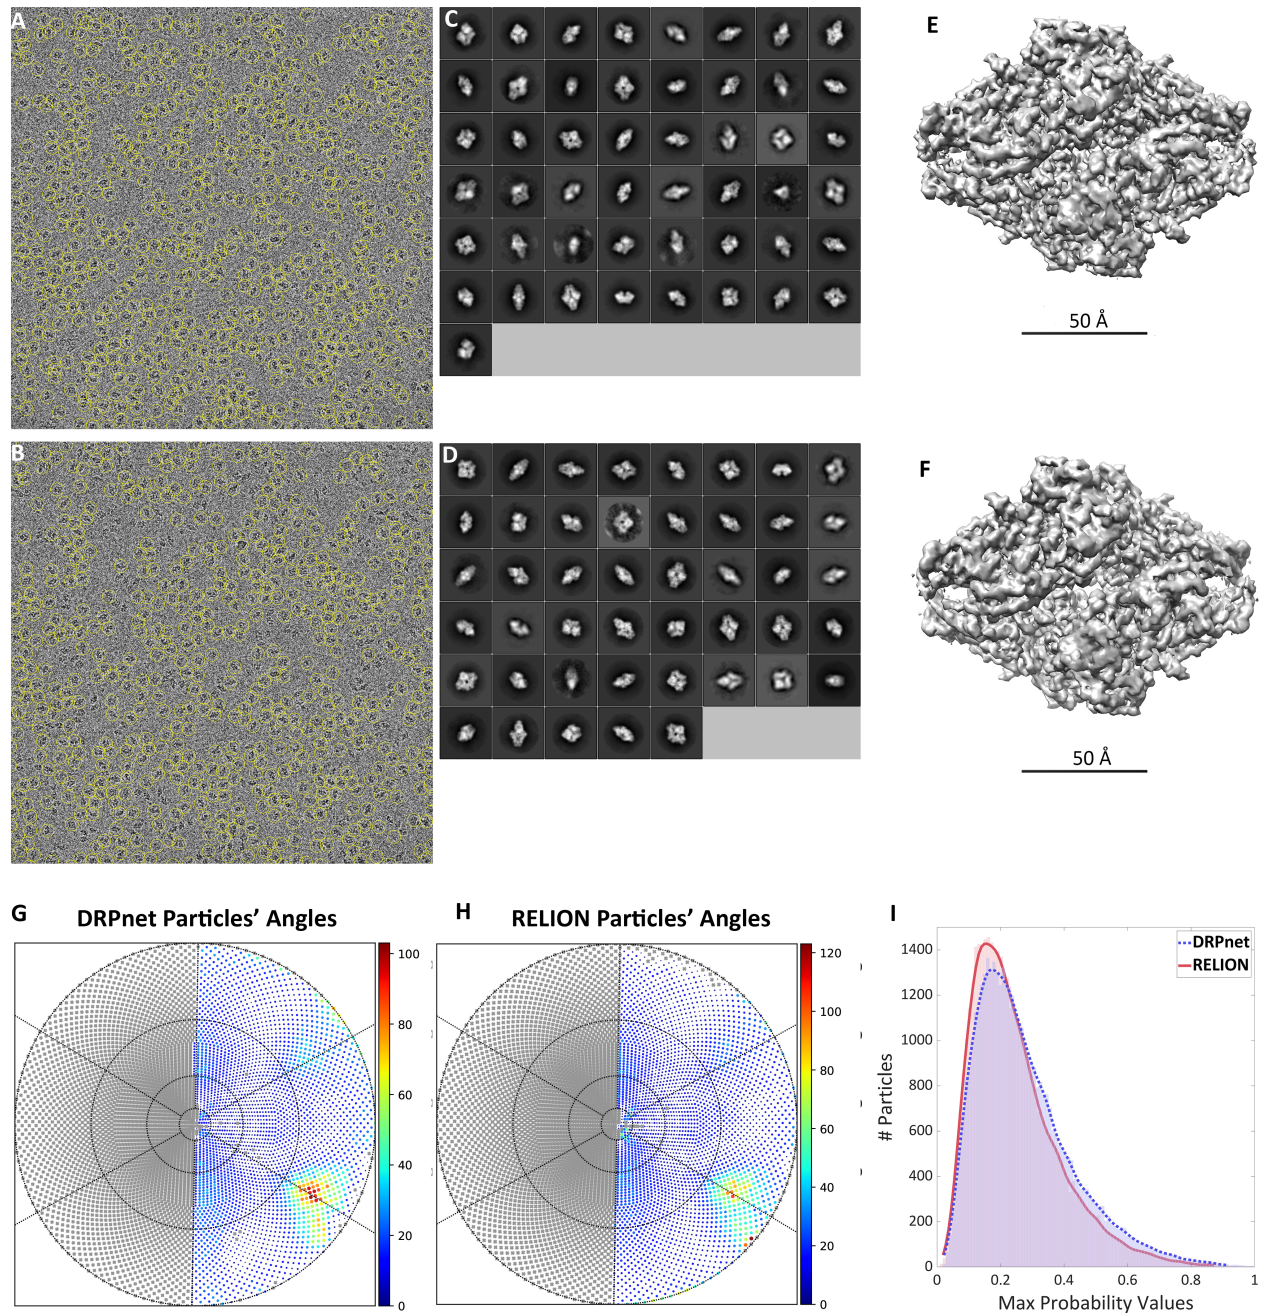


Figure 5: Test results of Empiar-10017 (β -galactosidase): Two examples of picked particles by DRPnet (A) and RELION (B). All picked particles were classified using RELION to select good 2D class averages as shown for DRPnet (C) and RELION (D). Then those good particles were used for 3D reconstruction, and the isosurface superimposed renderings of the maps were generated from DRPnet's result (E) and RELION's result (F), respectively. The particles' angular distributions in reconstruction process are displayed in Rot-Tilt angle scatter plots by results of DRPnet (G) and RELION (H). A gray "X" symbol indicates a missing view. The last sub-figure overlays the histogram of Max Value Probability Distribution (I) of particles picked by DRPnet (blue dashed line) and RELION (red solid line).

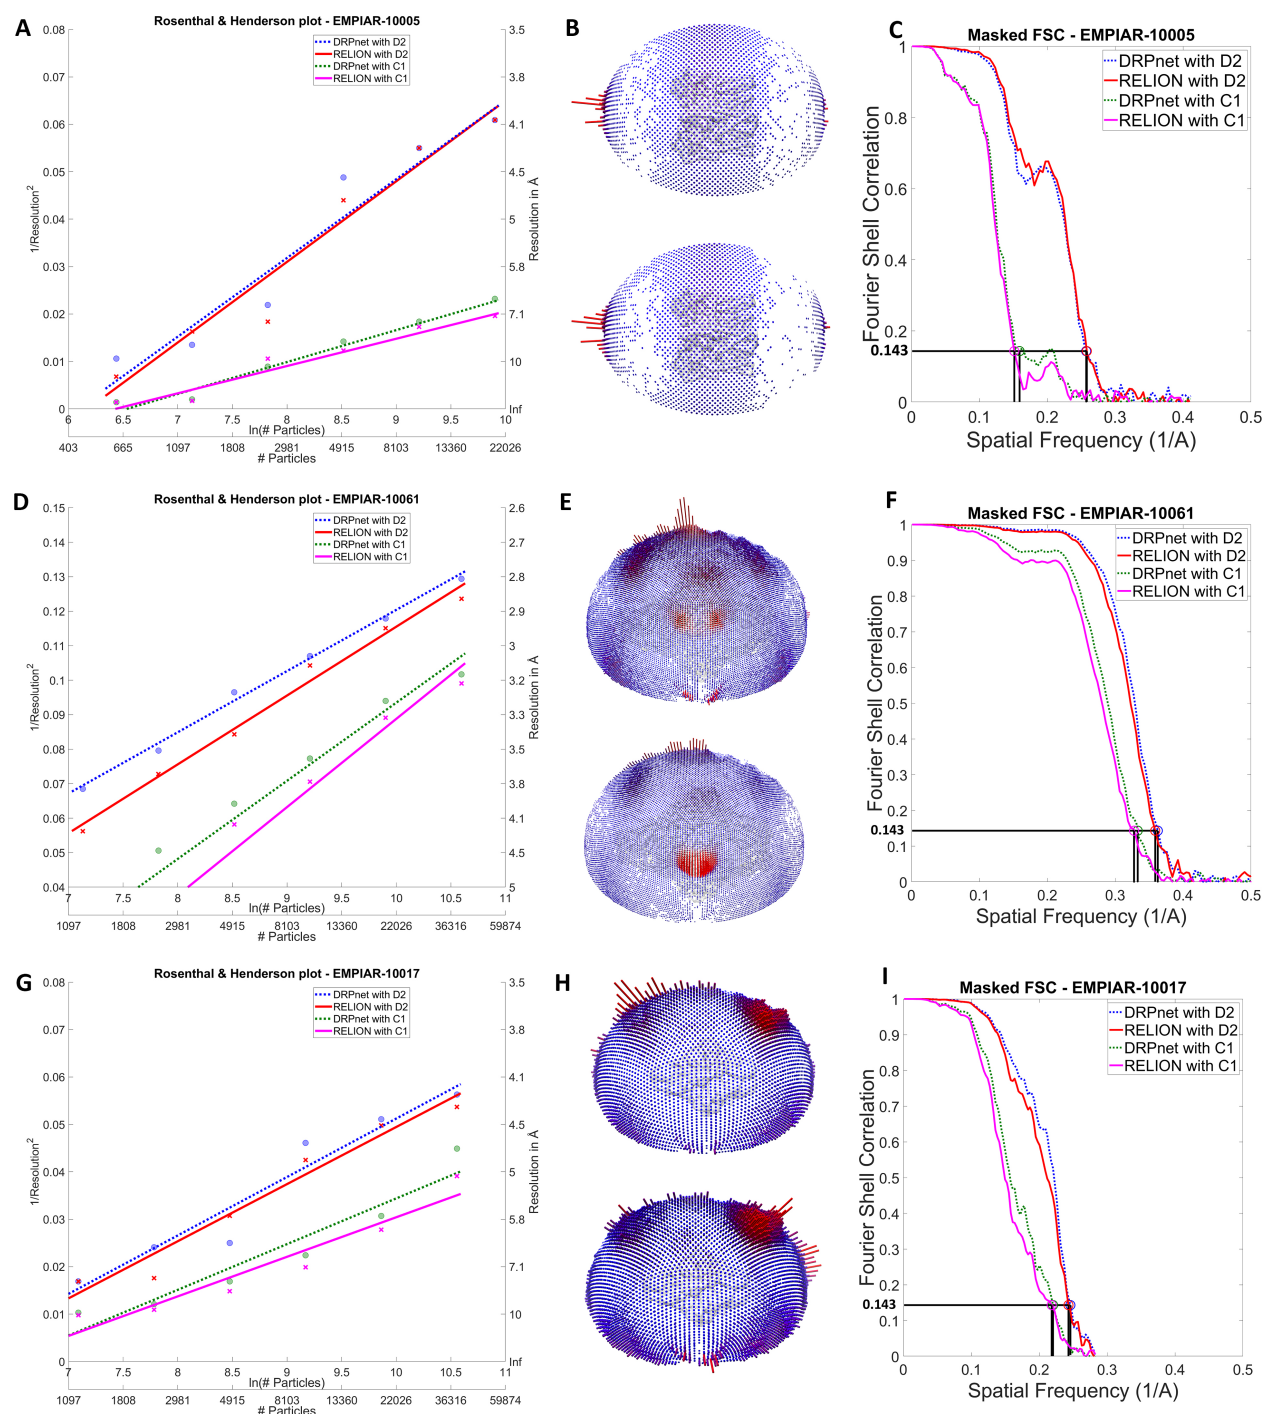


Figure 6: 3D reconstruction outcomes: the right column contains Fourier Shell Correlation curves, the middle column contains visualization of bild files for the angle distributions of good particles picked from 3 datasets by DRPnet's pick (upper) and RELION's autopick (lower), and the left column contains Rosenthal & Henderson plots show how resolution of 3D structures changed when the number of picked particles changed. The first row shows EMPIAR-10005 (TRPV1)'s results: (A) Rosenthal & Henderson plot, (B) angle distribution by DRPnet's pick (upper) and RELION (lower), (C) Masked FCS curve; The second has EMPIAR-10061 (β -galactosidase) results: (D) Rosenthal & Henderson plot, (E) angle distribution by DRPnet's pick (upper) and RELION (lower), (F) Masked FCS curve; The last row shows EMPIAR-10017 (β -galactosidase)'s results: (G) Rosenthal & Henderson plot, (H) angle distribution by DRPnet's pick (upper) and RELION (lower), (I) Masked FCS curve.

representing the total particles used for alignment, classification, and averaging. Specifically, in case of EMPIAR-10017 (β -galactosidase), although the ground truth was the most comprehensive, it still did not cover all particle patterns, thus orientations. This may be due to fluctuations in intensity level, which may effect the ground truth's particle patterns. Cryo-electron microscopy images captured with low contrast (low defocus or minimal electron radiation exposure, therefore enhancing high resolution terms) are most likely have high resolution 3D protein structural information, and conversely, those captured with high contrast (high defocus or high electron radiation exposure, enhancing low resolution terms) contain low resolution information. Because images in EMPIAR-10017 (β -galactosidase) have different contrast levels, a true positive particle in a low contrast image can be a false positive pattern in high contrast image. For the above reasons, the detection stage F-measure of 78.54% was not high, but it was the best result we could achieve with EMPIAR-10017 (β -galactosidase) data.

Therefore, we sought additional metrics to guide our analysis. Metrics we found useful include reporting resolution using the gold standard Fourier Shell Correlation (correlation threshold of 0.143) [46], angular distributions plots, efficiency (E_{od}) value [45], plots of maximum value of probability distributions, and inspecting the 3D isosurface.

Gold-standard Fourier Shell Correlation or 3D reconstruction resolution reveals the internal consistency in 3D reconstruction output maps [46]. Consistently, we obtained slightly higher resolutions with DRPnet over Relion's autopick, and DRPnet excelled with low contrast data sets (those collected on Gatan direct detectors with low defocus values). Plotting the angular distributions allows one to visually inspect the orientations (angles and rotations) of the picked particles and also their abundance. With DRPnet, there is more coverage (less white background) in all datasets. This finding is also corroborated by the efficiency metric, with DRPnet consistently showing improved E_{od} values, indicating a more robust sampling of particle angular orientations in the final 3D reconstruction. When assigning a particle's orientation during 3D refinement, histogram plots of the Maximum Probability Distribution Values provide a reference of how confidently a particle's orientation was assigned (a value of 1 indicates that the probability value has converged, and a value of 0 indicates great uncertainty) [58]. The shape of this distribution is informative; we note that with DRPnet-picked particles for EMPIAR-10061 & EMPIAR-10005, there is a greater number of particles having a higher maximum probability distribution value, with the shape of the histogram distribution shifted to the right. For our DRPnet analysis of EMPIAR-100017, the two curves are very similar; with RELION's autopicked particle distribution values having slight improvement over DRPnet's. Finally, visual inspection of the 3D reconstruction isosurface with USCF Chimera [59] or generating difference densities revealed little to no improvement in secondary structure and side chain densities with DRPnet versus RELION autopicked 3D reconstructions. 3

We noted minor improvements in 3D reconstruction resolution with DRPnet as compared to RELION's picking strategy, and believe this is due to picked particles having better angular orientation coverage. This was true for both TRPV1 (EMPIAR-10005) and β -galactosidase (EMPIAR-10061), which are low contrast (low defocus, high resolution) datasets collected from Gatan direct detectors. Based on the concept of blob detection, DRPnet can pick diversified patterns which usually results in wider particle orientation angle coverage, allowing an improvement in resolution. This is different than the manually selected templates used in RELION's autopicking function. It is likely especially for low contrast (low defocus, high resolution) datasets that bias from the user is minimized, preferred high contrast particle views are picked as well as other less obvious views with less contrast with the resulting particle set having more diverse angular orientations. We also note that in these aforementioned results, we used a subset of the total data. When using an entire EMPIAR-10017 (β -galactosidase) dataset to test, containing numerous particles (around 50,000) to perform 3D reconstructions, DRPnet and RELION reached the same resolution. If limited by low contrast data or not enough particles, DRPnet may be particularly useful.

Here, we present deep regression-based particle picking algorithms composed of a dual cascade of convolutional neural networks, entitled "DRPnet". The power of this strategy is that 1) the picking algorithm does not have to be, but can be, retrained, 2) efficiency of automated particle picking compared to manual template generation, and 3) usefulness for low contrast datasets (low defocus, high resolution) or data that is not complete (needs better angular sampling). We also provide functional assessment metrics for cryoEM micrograph particle picking beyond traditional recall, precision, and F-measure. In the future, we anticipate developing a powerful classification network for our workflow, trained with multiple types of particles (heterogeneous populations of particles, for example). Having a more powerful classification would also reduce the running time of our picking program while increasing its accuracy (reducing false positive particles such as noisy and/or low contrast particles). The second issue we also need to improve is more accurate localization of the particles within the image patches. Tackling these issues will improve the DRPnet algorithm and also will provide a refined dataset to perform high resolution 3D reconstructions.

References

- [1] M. Liao, E. Cao, D. Julius, and Y. Cheng. Structure of the TRPV1 ion channel determined by electron cryo-microscopy. *Nature*, 504(7478):107–12, 2013.
- [2] Sjoers H. W. Scheres. Semi-automated selection of cryo-EM particles in RELION-1.3. *Journal of Structural Biology*, 189(2):114–122, 2015.
- [3] Alberto Bartesaghi, Alan Merk, Soojay Banerjee, Doreen Matthies, Xiongwu Wu, Jacqueline L. S. Milne, and Sriram Subramaniam. 2.2 Å resolution cryo-EM structure of β -galactosidase in complex with a cell-permeant inhibitor. *Science*, 348(6239):1147–1151, 2015.
- [4] Marc Adrian, Jacques Dubochet, Jean Lepault, and Alasdair W. McDowell. Cryo-electron microscopy of viruses. *Nature*, 308, 1984.
- [5] Pawel Penczek, Michael Radermacher, and Joachim Frank. Three-dimensional reconstruction of single particles embedded in ice. *Ultramicroscopy*, 40(1):21, 1992.
- [6] Wah Chiu. Electron microscopy of frozen, hydrated biological specimens. *Annual Review of Biophysics and Biophysical Chemistry*, 15:21, 1986.
- [7] Axel F. Brilot, James Z. Chen, Anchi Cheng, Junhua Pan, Stephen C. Harrison, Clinton S. Potter, Bridget Carragher, Richard Henderson, and Nikolaus Grigorieff. Beam-induced motion of vitrified specimen on holey carbon film. *Journal of Structural Biology*, 177(3):630 – 637, 2012.
- [8] Monya Baker. Cryo-electron microscopy shapes up. *Nature*, 561(7724):565-567, 2018.
- [9] David Bernard Williams and C. Barry Carter. *Transmission Electron Microscopy - A Textbook for Materials Science*. Springer, 2009.
- [10] S. H. Scheres and S. Chen. Prevention of overfitting in cryo-EM structure determination. *Nat Methods*, 9(9):853–4, 2012.
- [11] J. M. Bell, M. Chen, P. R. Baldwin, and S. J. Ludtke. High resolution single particle refinement in EMAN2.1. *Methods*, 100:25–34, 2016.
- [12] S. J. Ludtke. Single-particle refinement and variability analysis in EMAN2.1. *Methods Enzymol*, 579:159–89, 2016.
- [13] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.02688, May 2016.
- [14] A. Punjani, J. L. Rubinstein, D. J. Fleet, and M. A. Brubaker. cryoSPARC: algorithms for rapid unsupervised cryo-em structure determination. *Nat Methods*, 14(3):290–296, 2017.
- [15] Y. Zhu, Q. Ouyang, and Y. Mao. A deep convolutional neural network approach to single-particle recognition in cryo-electron microscopy. *BMC Bioinformatics*, 18(1):348, 2017.
- [16] S. Kevin Zhou. *Medical Image Recognition, Segmentation and Parsing*. Academic Press, 2016.
- [17] S. Kevin Zhou, Hayit Greenspan, and Dinggang Shen. *Deep Learning for Medical Image Analysis*. Academic Press, 2017.
- [18] Guorong Wu, Dinggang Shen, and Mert R. Sabuncu. *Machine Learning and Medical Imaging*. Academic Press, 2016.
- [19] F. Xing and L. Yang. Chapter 4 - machine learning and its application in microscopic image analysis. In Guorong Wu, Dinggang Shen, and Mert R. Sabuncu, editors, *Machine Learning and Medical Imaging*, pages 97 – 127. Academic Press, 2016.
- [20] H.-C. Shin, Matthew Orton, David J Collins, and M.O. Leach. *Organ Detection Using Deep Learning*, book section 7, pages 123–153. Academic Press, 2016.
- [21] Yuanpu Xie, Fuyong Xing, and Lin Yang. Chapter 7 - deep voting and structured regression for microscopy image analysis. In S. Kevin Zhou, Hayit Greenspan, and Dinggang Shen, editors, *Deep Learning for Medical Image Analysis*, pages 155 – 175. Academic Press, 2017.
- [22] F. Wang, H. Gong, G. Liu, M. Li, C. Yan, T. Xia, X. Li, and J. Zeng. DeepPicker: A deep learning approach for fully automated particle picking in cryo-EM. *Journal of Structural Biology*, 195(3):325–36, 2016.
- [23] Yifan Xiao and Guangwen Yang. A fast method for particle picking in cryo-electron micrographs based on fast R-CNN. In *AIP conference proceedings*, volume 1836, page 020080, 2017.

- [24] Ross Girshick. Fast r-cnn. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV), ICCV '15*, pages 1440–1448, Washington, DC, USA, 2015. IEEE Computer Society.
- [25] A. Heimowitz, J. Anden, and A. Singer. Apple picker: Automatic particle picking, a low-effort cryo-EM framework. *J Struct Biol*, 204(2):215–227, 2018.
- [26] Thorsten Wagner, Felipe Merino, Markus Stabrin, Toshio Moriya, Christos Gatsogiannis, and Stefan Raunser. SPHIRE-crYOLO: A fast and well-centering automated particle picker for cryo-EM. *bioRxiv*, 2018.
- [27] Joseph Redmon and Ali Farhadi. YOLO9000: Better, faster, stronger. *arXiv*, 2016.
- [28] Zhong-Qiu Zhao, Peng Zheng, Shou-tao Xu, and Xindong Wu. Object detection with deep learning: A review. *arXiv:1807.05511v1*, 2017.
- [29] Dimitry Tegunov and Patrick Cramer. Real-time cryo-EM data pre-processing with warp. June 2018.
- [30] Lin Yang, Yizhe Zhang, Zhuo Zhao, Hao Zheng, Peixian Liang, Michael T. C. Ying, Anil T. Ahuja, and Danny Ziyi Chen. Boxnet: Deep learning based biomedical image segmentation using boxes only annotation. *CoRR*, abs/1806.00593, 2018.
- [31] T. Bepler, A. Morin, A. J. Noble, J. Brasch, L. Shapiro, and B. Berger. Positive-unlabeled convolutional neural networks for particle picking in cryo-electron micrographs. *Res Comput Mol Biol*, 10812:245–247, Apr 2018.
- [32] P.V.C. Hough. A method and means for recognizing complex patterns, 1962.
- [33] David G. Lowe. Distinctive image features from scale-invariant keypoints (sift). *International Journal of Computer Vision*, 60(2):28, 2004.
- [34] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Computer Vision and Image Understanding*, 110(3):14, 2008.
- [35] Joseph Redmon and Ali Farhadi. YOLOv3: An incremental improvement. *arXiv*, 2018.
- [36] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: Single shot multibox detector. *arXiv*, 2016.
- [37] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *International Conference on Neural Information Processing Systems*, pages 91–99.
- [38] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask RCNN. *arXiv*, 2017.
- [39] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, June 2010.
- [40] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014.
- [41] Richard Szeliski. *Computer Vision: Algorithms and Applications*. Springer-Verlag, Berlin, Heidelberg, 1st edition, 2010.
- [42] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. The MIT Press, 2016.
- [43] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- [44] Jiawei Han, Micheline Kamber, and Jian Pei. *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 3rd edition, 2011.
- [45] K. Naydenova and C. J. Russo. Measuring the effects of particle orientation to improve the efficiency of electron cryomicroscopy. *Nat Commun*, 8(1):629, 2017.
- [46] Peter B. Rosenthal and Richard Henderson. Optimal determination of particle orientation, absolute hand, and contrast loss in single-particle electron cryomicroscopy. *Journal of Molecular Biology*, 333(4):721–745, 2003.
- [47] Jasenko Zivanov, Takanori Nakane, Björn O Forsberg, Dari Kimanius, Wim JH Hagen, Erik Lindahl, and Sjors HW Scheres. New tools for automated high-resolution cryo-em structure determination in relion-3. *eLife*, 7:e42166, nov 2018.
- [48] Joachim Frank. *Three-Dimensional Electron Microscopy of Macromolecular Assemblies: Visualization of Biological Molecules in Their Native State*. OXFORD UNIVERSITY PRESS, 2006.
- [49] Pawel A. Penczek. *Resolution Measures in Molecular Electron Microscopy*, pages 73–100. Methods in Enzymology. 2010.
- [50] G. Harauz and M. van Heel. Exact filters for general geometry three dimensional reconstruction. *OPTIK*, 73(4):146–156, 1986.

- [51] Michael Unser, Benes L. Trus, and Alasdair C. Steven. A new resolution criterion based on spectral signal-to-noise ratios. *Ultramicroscopy*, 23:4, 1987.
- [52] R. Henderson, S. Chen, J. Z. Chen, N. Grigorieff, L. A. Passmore, L. Ciccarelli, J. L. Rubinstein, R. A. Crowther, P. L. Stewart, and P. B. Rosenthal. Tilt-pair analysis of images from a range of different specimens in single-particle electron cryomicroscopy. *J Mol Biol*, 413(5):1028–46, 2011.
- [53] B. Bottcher, S.A. Wynne, and R.A. Crowther. Determination of the fold of the core protein of hepatitis b virus by electron cryomicroscopy. *Nature*, 386:4, 1997.
- [54] Robert Langlois, Jesper Pallesen, Jordan T. Ash, Danny N. Ho, John L. Rubinstein, and Joachim Frank. Automated particle picking for low-contrast macromolecules in cryo-electron microscopy. *Journal of Structural Biology*, In Press, 2014.
- [55] E. F. Pettersen, T. D. Goddard, C. C. Huang, G. S. Couch, D. M. Greenblatt, E. C. Meng, and T. E. Ferrin. Ucsf chimera—a visualization system for exploratory research and analysis. *J Comput Chem*, 25(13):1605–12, 2004. Pettersen, Eric F Goddard, Thomas D Huang, Conrad C Couch, Gregory S Greenblatt, Daniel M Meng, Elaine C Ferrin, Thomas E eng P41 RR001081/RR/NCRR NIH HHS/ P41-RR01081/RR/NCRR NIH HHS/ Research Support, U.S. Gov't, P.H.S. *J Comput Chem*. 2004 Oct;25(13):1605-12. doi: 10.1002/jcc.20084.
- [56] Ardan Patwardhan, José-Maria Carazo, Bridget Carragher, Richard Henderson, J Bernard Heymann, Emma Hill, Grant J Jensen, Ingvar Lagerstedt, Catherine L Lawson, Steven J Ludtke, David Mastronarde, William J Moore, Alan Roseman, Peter Rosenthal, Carlos-Oscar S Sorzano, Eduardo Sanz-García, Sjors H W Scheres, Sriram Subramaniam, John Westbrook, Martyn Winn, Jason R Swedlow, and Gerard J Kleywegt. Data management challenges in three-dimensional EM. *Nature Structural & Molecular Biology*, 19(12):1203–1207, December 2012.
- [57] Sjors H. W. Scheres. Chapter six - processing of structurally heterogeneous cryo-EM data in RELION. In R. A. Crowther, editor, *The Resolution Revolution: Recent Advances In cryoEM*, volume 579 of *Methods in Enzymology*, pages 125 – 157. Academic Press, 2016.
- [58] Melody G Campbell, David Veisler, Anchi Cheng, Clinton S Potter, and Bridget Carragher. 2.8 Å resolution reconstruction of the *Thermoplasma acidophilum* 20S proteasome using cryo-electron microscopy. *eLife*, 4:e06380, mar 2015.
- [59] Eric F. Pettersen, Thomas D. Goddard, Conrad C. Huang, Gregory S. Couch, Daniel M. Greenblatt, Elaine C. Meng, and Thomas E. Ferrin. Ucsf chimera - a visualization system for exploratory research and analysis. *Journal of Computational Chemistry*, 25(13):1605–1612, 2004.