

DeepPicker: A deep learning approach for fully automated particle picking in cryo-EM



Feng Wang^{c,1}, Huichao Gong^{b,1}, Gaochao Liu^{a,d,1}, Meijing Li^{a,d}, Chuangye Yan^{a,e}, Tian Xia^{c,*}, Xueming Li^{a,d,e,*}, Jianyang Zeng^{b,*}

^a School of Life Sciences, Tsinghua University, Beijing 100084, China

^b Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing 100084, China

^c School of Electronic Information and Communications, Huazhong University of Science and Technology, Wuhan 430074, China

^d Beijing Advanced Innovation Center for Structure Biology, Tsinghua University, Beijing 100084, China

^e Tsinghua-Peking Joint Center for Life Sciences, Tsinghua University, Beijing 100084, China

ARTICLE INFO

Article history:

Received 28 April 2016

Received in revised form 8 July 2016

Accepted 11 July 2016

Available online 14 July 2016

Keywords:

Cryo-EM

Particle picking

Automation

Deep learning

ABSTRACT

Particle picking is a time-consuming step in single-particle analysis and often requires significant interventions from users, which has become a bottleneck for future automated electron cryo-microscopy (cryo-EM). Here we report a deep learning framework, called DeepPicker, to address this problem and fill the current gaps toward a fully automated cryo-EM pipeline. DeepPicker employs a novel cross-molecule training strategy to capture common features of particles from previously-analyzed micrographs, and thus does not require any human intervention during particle picking. Tests on the recently-published cryo-EM data of three complexes have demonstrated that our deep learning based scheme can successfully accomplish the human-level particle picking process and identify a sufficient number of particles that are comparable to those picked manually by human experts. These results indicate that DeepPicker can provide a practically useful tool to significantly reduce the time and manual effort spent in single-particle analysis and thus greatly facilitate high-resolution cryo-EM structure determination. DeepPicker is released as an open-source program, which can be downloaded from <https://github.com/nejyeh/DeepPicker-python>.

© 2016 Elsevier Inc. All rights reserved.

1. Introduction

Recent advance of the single-particle electron cryo-microscopy (cryo-EM) technique has been revolutionizing the structural biology field (Bai et al., 2015; Cheng, 2015; Cheng et al., 2015), and enabled protein complex structure determination at near-atomic resolution (Liao et al., 2013; Yan et al., 2015; Zhao et al., 2015; Wu et al., 2015). However, the high-resolution cryo-EM studies of molecular complexes, especially those asymmetric ones, often require the selection of a tremendous number of (e.g., hundreds of thousands of) high-quality particles from micrographs. This particle picking step is a labor-intensive step in single-particle data analysis and is a major obstacle for automated cryo-EM pipeline. In the past, particles from cryo-EM micrographs are often selected manually. Such a manual picking process is usually a laborious,

tedious and time-consuming task which inevitably requires a considerable amount of human effort to obtain a sufficient number of good-quality particles to ensure high-resolution 3D reconstruction. In addition, manual particle selection is normally subjective and can easily introduce bias and inconsistency due to change in human judgement over time.

To relieve the bottleneck in single-particle data analysis, numerous computational approaches have been proposed to facilitate the particle picking process (Nicholson and Glaeser, 2001; Zhu et al., 2004; Ogura and Sato, 2001; Langlois et al., 2011; Chen and Grigorieff, 2007; Hall and Patwardhan, 2004; Huang and Penczek, 2004; Adiga et al., 2004; Voss et al., 2009; Sorzano et al., 2009; Arbez et al., 2011; Zhao et al., 2013; Norouzi et al., 2013). These methods can be basically divided into three categories, including generative (Chen and Grigorieff, 2007; Hall and Patwardhan, 2004; Huang and Penczek, 2004), unsupervised (Adiga et al., 2004; Voss et al., 2009) and discriminative approaches (Langlois et al., 2011; Sorzano et al., 2009; Arbez et al., 2011). The generative approaches (Chen and Grigorieff, 2007; Hall and Patwardhan, 2004; Huang and Penczek, 2004) usually measure

* Corresponding authors.

E-mail addresses: isutian@gmail.com (T. Xia), lixueming@mail.tsinghua.edu.cn (X. Li), zengjy321@tsinghua.edu.cn (J. Zeng).

¹ These authors contributed equally to this work.

the similarity to a reference to identify particle candidates from micrographs. A typical generative method employs a template-matching technique (Hall and Patwardhan, 2004; Huang and Penczek, 2004) with a cross-correlation similarity measure to accomplish particle selection. The unsupervised approaches distinguish the images of particle-like objects from background or noise in micrographs via an unsupervised learning manner (i.e., without any labeled training data) (Adiga et al., 2004; Voss et al., 2009). The discriminative methods first train a classifier based on a labeled dataset of positive and negative examples, and then apply this trained classifier to detect and recognize particle images from micrographs (Langlois et al., 2011; Sorzano et al., 2009; Arbez et al., 2011). Although these computational approaches have greatly reduced time and effort spent on single-particle data analysis, there still remain gaps to achieve a fully automated pipeline for efficient particle picking. For instance, most of the generative methods require the user to prepare an initial set of high-quality reference particles used as templates to search for similar particle candidates from micrographs, while the discriminative approaches usually demand the user to manually pick a number of positive and negative samples to train the classifier. Thus, these approaches generally depend upon a certain level of human intervention to provide a portion of manually-picked particles to initialize the particle selection process. Although the unsupervised approaches do not heavily rely on hand-labeled data, they rarely fully exploit the intrinsic and unique characters of particles to facilitate automated particle picking. Therefore, the unsupervised approaches are often combined with the template-matching or classification based approaches to achieve decent picking results (Zhao et al., 2013; Norousi et al., 2013).

In recent years, deep learning has become an increasing popular tool in the machine learning field probably due to the availability of large-scale training data, the advance of powerful computing platforms and the development of efficient learning algorithms (LeCun et al., 2015; Hinton et al., 2012; Sutskever et al., 2014). Tests on several well-known benchmark datasets have demonstrated that deep learning can achieve better performance especially in large-scale data analysis than traditional machine learning approaches (LeCun et al., 2015; Hinton et al., 2012; Krizhevsky et al., 2012; Goldberg, 2015). So far deep learning has been successfully applied to a wide range of data science fields, such as computer vision (Krizhevsky et al., 2012; Lawrence et al., 1997; Farabet et al., 2013; Tompson et al., 2014), natural language processing (Hinton et al., 2012; Mikolov et al., 2011; Xu et al., 2015) and computational biology (Zhou and Troyanskaya, 2015; Ma et al., 2015). More recently, the Google team has shown that the deep learning framework is capable of achieving an impressive and amazing level of artificial intelligence that can closely mimic the problem solving skills of human experts (Silver et al., 2016; Mnih et al., 2015). Despite these successful stories in a variety of applications, it still remains unknown whether the deep learning technique can also be effectively used to address the current problems in particle picking and achieve a fully automated particle selection procedure in single-particle cryo-EM data analysis. Here, we aim to answer this question and propose a deep learning framework, called DeepPicker, to fill the aforementioned gaps in automated particle selection and liberate structural biologists from the dreary manual picking process to focus on more interesting work.

Although our deep learning model can be trained with a semi-automated manner as in the conventional classification based approaches, it can be equipped with a new training scheme, which fully exploits the known particles of other molecular complexes that are different from the target one and whose structures have been previously determined by cryo-EM. This new training strategy does not require any manual effort in particle picking for the

current target molecular complex and thus is considered *fully automated*. Our deep learning model with such a new cross-molecule training scheme is innovative and can effectively capture the common abstract representation of latent features from the known particles of the previously-determined molecular structures. It can closely mimic human intelligence and effectively use the extracted cross-molecule features to initialize the particle picking process of the current target complex. To our knowledge, our approach is the *first* successful attempt to fully exploit the cross-molecule data to achieve *full* automation in particle picking without any human interference.

We have implemented our deep learning method for particle picking and tested it on the real cryo-EM data of three molecules that have been published in the past three years, including TRPV1 (Liao et al., 2013) (from EMPIAR (<https://www.ebi.ac.uk/pdbe/emdb/empiar/>)), human γ -secretase (Sun et al., 2015) and yeast spliceosome (Yan et al., 2015). Our tests have demonstrated that the proposed deep learning framework with either semi-automated or fully automated training scheme can accurately detect and select a sufficient number of particles that are comparable to those picked manually by human experts. Our new automated picking approach can significantly reduce time and labor spent in single-particle data analysis and thus greatly relieve a bottleneck in the automated cryo-EM structure determination pipeline.

2. Methods

2.1. The fully automated particle picking pipeline

We have implemented a deep learning framework for automated particle picking in single-particle cryo-EM structure determination (Fig. 1). The automated particle picking pipeline consists of two modules, i.e., model training and particle picking (Fig. 1). In the model training module, a set of labeled positive and negative samples is used to train a convolutional neural network (CNN) model (see Sections 2.3 and 2.4), while in the particle picking module, the trained CNN classifier is then used to select particle images from input micrographs. The particle picking module is further divided into five steps: scoring, cleaning, filtering, sorting and iterating. In the scoring step, a sliding window (i.e., a square box) of a fixed size is used to scan each micrograph from the top left corner to the bottom right corner with a constant step size. The box size of the sliding window is chosen to be slightly larger than the particle size, which can be easily estimated and defined as a parameter. During the above scanning process, each patch within the sliding window is extracted and fed to the trained CNN classifier as input data. The prediction score between 0 and 1 output by the CNN model, which represents the probability of being a particle at the current position, is then assigned to the center of the corresponding window. After that, we obtain a “scored map”, which describes the distribution of the likelihood scores of particles over the whole micrograph. As ice can easily introduce false positives during the picking process, we also employ a cleaning step to discard these false particles from the candidate list. In this cleaning step, we first connect any two neighboring pixels if their prediction scores are both above a threshold, and then examine the size of each connected domain (i.e., the portion of all connected pixels). If the size of a connected domain is larger than a cutoff value, it is regarded as a potential false positive probably due to ice, and thus removed from the candidate list. In the filtering step, we aim to refine the current set of particle candidates and also identify the center coordinates of the final remaining particles from the scored map. We first introduce a concept of peak window, the size of which is related to the minimum distance between cen-

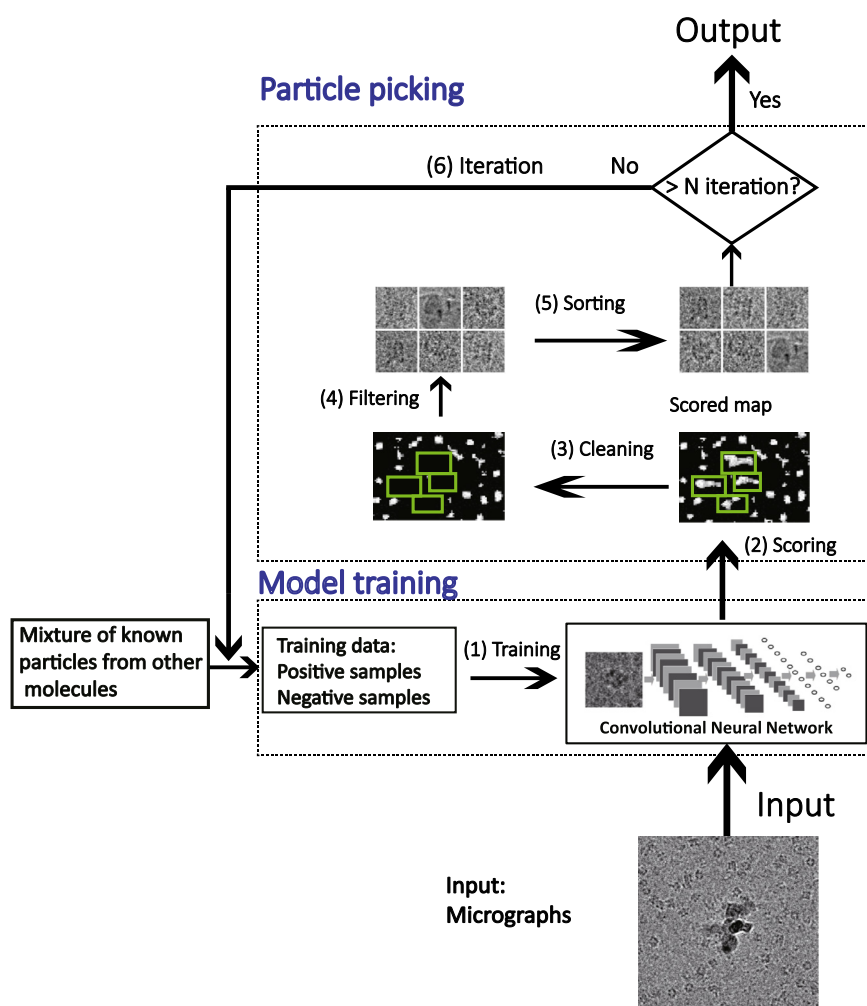


Fig. 1. Schematic overview of our automated particle picking pipeline. The particle selection procedure consists of six steps: (1) training, (2) scoring, (3) cleaning, (4) filtering, (5) sorting and (6) iteration. In our fully automated picking pipeline, the initial training data are obtained from the known particles of other molecular complexes whose structures have been previously solved via cryo-EM.

ters of two possible particles. Then the position with the maximum prediction score in each peak window is chosen and output as the center of a particle. We also remove bad particle candidates when they are far away from the mean case of particle candidates in the same micrograph. More specially, the following two steps are performed to accomplish this task. First, for each particle candidate in a micrograph, we count the number of extreme pixels, in which the pixel value is three standard deviations away from the mean pixel value of the whole micrograph. Next, we remove those bad particle candidates, in which the number of extreme pixels is larger than three standard deviations away from the mean number of extreme pixels over all particle candidates over the whole micrograph. In the sorting step, we sort the remaining particle candidates according to their prediction scores. In the iterating step, we use the particles picked by the previous CNN classifier which was trained over the known particles of other molecules to further refine the CNN model. After a certain number of iterations, the algorithm outputs the top list of the highest-rated particles. The detailed setting of the parameters described for the above operations is provided in Section 2.6.

In general, the particles of different molecular complexes display distinct sizes and shapes. Thus, how to effectively exploit the cross-molecule features to pick particles from micrographs of the current complex is a major challenge in our work. As described above, this problem has been tackled in our framework mainly

using two strategies. First, the known particles from multiple types of other molecular complexes are combined together to train the deep learning classifier and enable it to capture common features of particles. Second, the particles of the target molecule identified from the deep learning classifier trained by a cross-molecule manner in the first iteration are further used to refine the model to better describe the features of the current target complex. In our automated particle picking scheme, the user is not required to manually pick any particle from micrographs of the current target molecule. Thus, it is considered a fully automated procedure without requiring any manual intervention during the particle selection process.

2.2. Datasets and data preprocessing

The data used in this paper were divided into two non-overlapping parts: training and test datasets. The training data consisted of 100, 400, 100, 84 and 400 micrographs of γ -secretase (Sun et al., 2015), spliceosome (Yan et al., 2015), TRPV1 (Liao et al., 2013), β -galactosidase (Bartesaghi et al., 2014) and N-ethylmaleimide sensitive factor complex (Zhao et al., 2015), respectively, while the test data contained a separate set of micrographs from γ -secretase, spliceosome and TRPV1, each having 100 micrographs. All the results described in the paper are referred to those on the test datasets.

The defocus value of each micrograph was calculated using CTFIND4 (Rohou and Grigorieff, 2015). For each micrograph, we first used a Gaussian filter as a low pass filter to remove white noise with high frequency components. Then the binning strategy (Li et al., 2013) was used to convert each original micrograph to an image ranging between 1000 and 2000 pixels. All the coordinates of reference particles were identified manually by human experts (see also Section 2.5) except for β -galactosidase, in which the coordinates of reference particles were obtained from EMPIAR (entry ID: EMPIAR-10017). In addition, all the coordinates of the reference particles were further aligned using FREALIGN (Grigorieff, 2007).

Note that we did not test our particle picking method on the classic KKH dataset (Zhu et al., 2004), which was collected in 2004. This is mainly because, as the technique of electron cryo-microscopy detector has been significantly revolutionized since 2012 (Cheng, 2015), we believe that the micrographs collected nowadays are quite different from the old datasets. Thus, we mainly tested the new datasets that were published quite recently (i.e., after 2012). In addition, we mainly focused on automated particle picking for small or medium molecules, and excluded the particles of large complexes or virus from our training or test data. This is mainly because automated particle selection of these large complexes or virus is usually a relatively easier task. For example, these large particles can be relatively easily detected using the edge detection techniques (Canny, 1986) or other image processing methods (Cheng et al., 2015).

The γ -secretase and spliceosome datasets were obtained from Dr. Yigong Shi's lab at Tsinghua University, which were acquired by an FEI Titan Krios electron microscope operating at 300 kV with a magnification of 22,500 \times . These micrographs were first taken by a Gatan K2 Summit using a super-resolution mode at 0.66 Å/pixel, and then binned to a final pixel size of 1.32 Å. The TRPV1 dataset was downloaded from EMPIAR (entry ID: EMPIAR-10005), in which the micrographs were recorded by a Gatan K2 Summit equipped with a FEI POLARA 300 operating at 300 kV in a super-resolution counting mode, with a final calibrated super resolution pixel size of 1.2156 Å. The β -galactosidase dataset was downloaded from EMPIAR (entry ID: EMPIAR-10017), which was recorded by a FEI Falcon-II camera. More detailed information of the β -galactosidase data collection can be found in Bartesaghi et al. (2014). The NSF (N-ethylmaleimide Sensitive Factor) fusion complex (20S particle) dataset was kindly offered by Dr. Minglei Zhao from Dr. Axel T. Brunger's Lab, and more details about its data collection can be found in the original paper (Zhao et al., 2015).

2.3. The convolutional neural network model

Our deep learning framework employs a convolutional neural network (CNN) as the classifier to discriminate correct particles (i.e., positive samples) from background random noise images (i.e., negative samples). Here, we briefly describe the concept of a CNN. The reader is referred to Cun et al. (1990), Lecun et al. (1998), LeCun et al. (2015), Krizhevsky et al. (2012), Lawrence et al. (1997) for more details of this deep learning model.

A convolutional neural network is a specific type of deep learning (Krizhevsky et al., 2012; Lecun et al., 1998; Zeiler and Fergus, 2013) with a multi-layer structure, which consists of an input layer, an output layer and several hidden layers (between input and output layers). Each layer is comprised of a number of units, also called artificial neurons. In a CNN, the output layer is a softmax layer (Bishop, 2006) in which the number of artificial neurons is equivalent to that of the classification classes, and the value of each unit represents the chance of being the corresponding class. In a softmax layer, the output of a CNN is ensured to be a probability value between 0 and 1. More specifically, a softmax function is defined as

$$P(Y_i) = \frac{e^{W_i X + b_i}}{\sum_j e^{W_j X + b_j}}$$

where Y_i represents the output, i and j stand for the indexes of the artificial neurons in the last layer (which also represent the indexes of the output classes), X represents the input to the softmax layer, W_i represents the corresponding weight parameter, and b_i stands for the bias of the i th unit in the output layer.

The hidden layers mainly include convolutional, pooling and fully-connected layers. The convolutional and fully-connected layers are also called the learning layers, since their weights are optimized during the training procedure. The dimension of the weight parameters in a convolutional layer is $N \times M \times K \times K$, where N stands for the number of kernels in the convolutional layer, M stands for the number of square matrices in the kernel (e.g., a grayscale image has one square matrix), and K represents the dimension of a square matrix which is typically less than 10. The convolutional layer takes the convolution operation between the input X and the weights W and then feeds the output into a nonlinear function. In particular, the convolution operation is defined as

$$\text{convolution}(X)_{ijn} = F_{\text{nonlinear}} \left(\sum_{m=0}^{M-1} \sum_{k=0}^{K-1} \sum_{l=0}^{K-1} W_{mkl}^n X_{m,k+i-1,l+j-1} \right)$$

where i and j denote the location indexes of the output, n is the index of the output kernel, m is the index of the input matrix, k and l stand for the location indexes of the kernel, and $F_{\text{nonlinear}}$ stands for a nonlinear function, e.g., $\tanh(x)$, $\text{sigmoid}(x) = \frac{1}{1+e^{-x}}$ and the rectified linear function (ReLU) $\text{ReLU}(x) = \max(0, x)$. The ReLU function has been popularly used in a CNN architecture, as it can significantly speed up the training process (Krizhevsky et al., 2012; Nair and Hinton, 2010).

The pooling layer first divides the input image into non-overlapping windows, and then takes a max or mean operation over each window. The pooling operation can reduce the number of parameters and thus the complexity of the learning model. In addition, it can effectively relieve the over-fitting problem (Jarrett et al., 2009; Scherer et al., 2010).

The architecture of our CNN model consists of four convolutional layers, each of which uses an ReLU function and is followed by a max pooling layer. Then these convolutional and max pooling layers are stacked with two fully-connected layers. The input data of our CNN framework is the image within sliding window (see Section 2.1), and the results of the last fully-connected layer are fed into a softmax layer with two units, which represent the positive and negative labels, respectively. The “dropout” technique (Srivastava et al., 2014), in which the unit states in a layer are set to zero with probability p , whose default value is set to 0.5, is applied in the first fully-connected layer to address the overfitting problem (Krizhevsky et al., 2012; Srivastava et al., 2014).

In our deep learning framework, the first convolutional layer filters the input image of size $1 \times 64 \times 64$ with 8 kernels of size $1 \times 9 \times 9$ with a stride of 1. The second convolutional layer has 16 kernels of size $8 \times 5 \times 5$ with a stride of 1. The third convolutional layer has 32 kernels of size $16 \times 3 \times 3$ with a stride of 1. The fourth convolutional layer has 64 kernels of size $32 \times 2 \times 2$ with a stride of 1. The first fully-connected layer has 256 neurons and the second (i.e., the last output layer) has two neurons.

2.4. Training

The stochastic gradient descent algorithm (Bottou, 2010) is used to train our deep learning model. The training data to our deep learning classifier include a set of positive samples and an equivalent number of negative samples, which are selected from

those regions at least 0.6 times the size of sliding window away from positive samples. The positive samples are either from the known particles of other molecules whose structures have been previously determined via cryo-EM (i.e., in a fully automated manner) or the manually picked particles of the target molecule (i.e., in a semi-automated manner). As particles may have different sizes, we normalize both positive and negative samples into images of 64 pixels \times 64 pixels. This normalization scheme is particularly useful when we use cross-molecule data to train our deep learning model, as the particles of distinct molecular complexes generally have different sizes. During the training process, we also use a separate validation dataset, which is independent from training data and test data, to determine some super-parameters of our deep learning model, e.g., the number of iterations to reach the convergence. The size of such a validation dataset is chosen to be around 1/9 of the training dataset.

In our fully automated picking scheme, we first use a mixture of 10,000 known particles of other molecules as training data that were different from the target complex, each contributing to a roughly equivalent number of training examples. Then in the iterating step (Fig. 1), the top 10,000 particles predicted by the previously trained deep learning classifier are used as new training data to further refine the deep learning model. For all the test results shown in the paper, we only ran one iteration of the training process.

2.5. Performance evaluation

We evaluated the performance of our particle picking approach mainly by comparing the automated picking results to the corresponding reference particles picked manually by human experts. These reference particles were first picked by one cryo-EM expert and then further verified by two additional experts. The coordinates of all manually-picked particles were further aligned using FREALIGN (Grigorieff, 2007). An automatically-picked particle was said to agree with a manually picked particle, if the distance between their centers was less than a threshold, which was set to 0.2 times the size of sliding window. We mainly used recall and precision to evaluate the accuracy of the particles selected by our automated approach. Let TP denote the number of particles that were picked by our deep learning approach and also agreed with reference particles picked by human experts. Let FP denote the number of particles that were picked by our approach but did not agree with any reference particle. Let FN denote the number of the reference particles that were picked by human experts but not by our automated approach. Then recall and precision are defined as $recall = TP / (FN + TP)$, and $precision = TP / (FP + TP)$, respectively.

In addition to recall and precision, we also measured the deviations of the centers of the particles picked by our automated approach from those of the reference particles selected by human experts. In particular, we assessed their average distances, normalized by the size of sliding window.

We also compared the 2D clustering and class averaging results of the particles picked by our approach to those from the manually picked images in the reference set. The 2D clustering and class averaging operations were performed using the corresponding commands in RELION 1.3. Such a 2D analysis process was conducted in 2–3 iterations. In each iteration, those obvious image artifacts, invalid particles or empty objects from the 2D averaging results were removed from the candidate list. The same data analysis procedure was done for both manually and automatically picked particles. Here the 2D averaging results were mainly used to examine whether the particles picked by DeepPicker were comparable to those picked by human experts. Note that the recall and

precision values were accessed before the 2D alignment filtering operation.

2.6. Parameter setting

In this section, we explain the parameter setting of our deep learning model for automated particle picking. In the data preprocessing stage, the standard deviation of the Gaussian filter was set to 0.1. The standard deviation we used here is in Fourier space, so it is equivalent to a 10 Å low-pass filter. In the model training stage (i.e., step 1 in Fig. 1), the learning rate was set to 0.01. In the scoring step of the particle picking stage (i.e., step 2 in Fig. 1), the size of sliding window was set to 180 pixels for TRPV1 and γ -secretase, and 320 pixels for spliceosome, and the step size was set to 4 pixels. In the cleaning step (i.e., step 3 in Fig. 1), two neighboring pixels in the scored map with prediction scores above 0.5 were connected. In addition, we removed those connected domains whose sizes were four times larger than the average size. In the filtering step (i.e., step 4 in Fig. 1), the size of peak window was set to 0.8 times the size of sliding window, in other words, the minimum distance between centers of two potential particles was set to 0.4 times the size of sliding window. In the iterating step (i.e., step 6 in Fig. 1), we chose 10,000 samples from the top sorted particles picked by the previously-trained classifier to further refine the CNN model.

2.7. Implementation

DeepPicker was originally implemented based on Torch7 (<http://torch.ch>), an open source deep learning library and the Lua programming language. While the paper was under the reviewer, we have reimplemented it in Python programming language based on the TensorFlow library (Abadi et al., 2015). It can be run either with or without graphic processing units (GPUs). In our tests, a NVIDIA Quadro K4000 GPU was used for training the model and picking the particles from micrographs. The time required for training the model with 10,000 samples was less than an hour, and picking the particles from a micrograph of size 3710 pixels \times 3838 pixels took about 1.5 min, when the step size was set to 4. In general, the picking time can vary with respect to different parameter settings. In particular, the time requirement was mainly related to the values of box size and step size, which principally determined the time of inferring each particle candidates and the number of particle candidates, respectively.

3. Results

3.1. Performance evaluation on fully automated particle picking

3.1.1. Accuracy of particle picking

We have tested and evaluated the performance of our fully automated particle picking approach on the published datasets of γ -secretase, spliceosome and TRPV1 (see Section 2.2 for details of data preprocessing), in which the known particles of the other two molecules that were different from the target complex were used as training data (Fig. 2). In addition, another two datasets, including β -galactosidase (Bartesaghi et al., 2014) and NSF (N-ethylmaleimide Sensitive Factor) fusion complex (20S particle) (Zhao et al., 2015), were incorporated into training data. We compared the fully automated picking results with the coordinates of the reference particles identified manually by human experts and measured both precision and recall scores (see Section 2.5). As in the current single-particle analysis pipeline, the fraction of true particles that have been picked is generally more important for

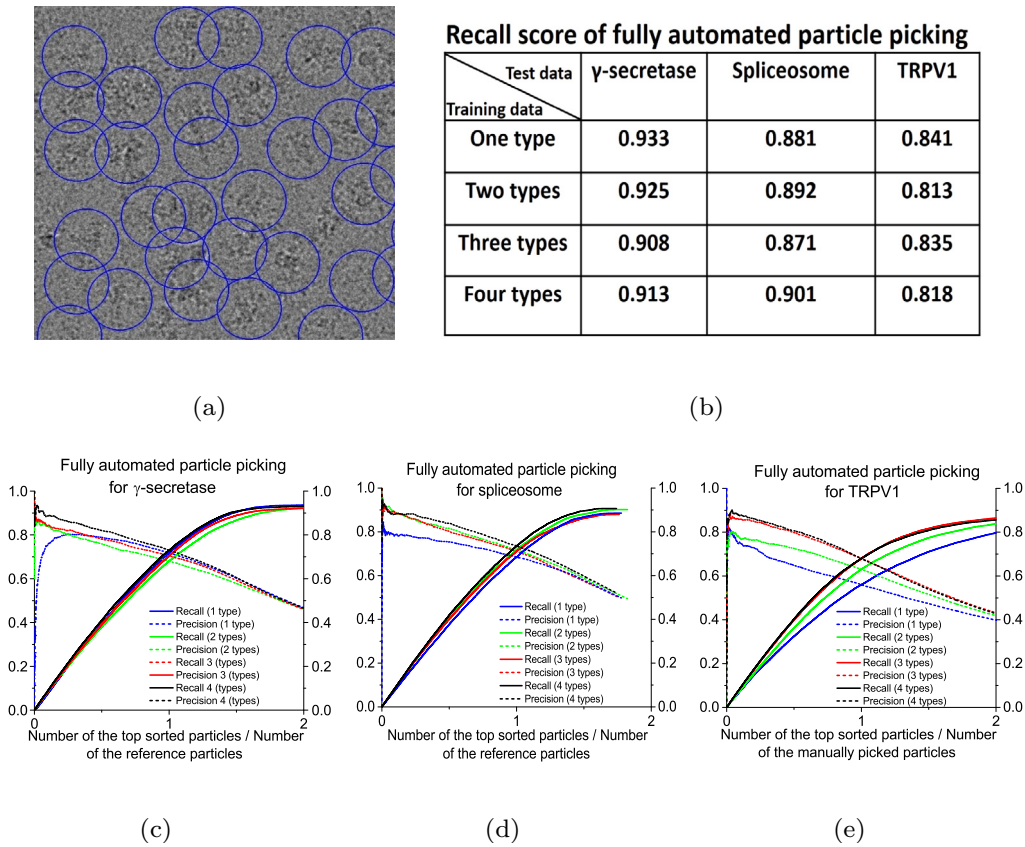


Fig. 2. Results on fully automated particle picking. (a) An example of our fully automated particle picking results for spliceosome, using a mixture of 5000 TRPV1 and 5000 γ -secretase particles picked manually by human experts as training data. (b) The summary of recall scores for γ -secretase, TRPV1 and spliceosome using different combinations of training data. The recall scores were computed on those picked particles with prediction scores above 0.5. (c)–(e) The trends of precision and recall vs. the number of picked particles for γ -secretase, spliceosome and TRPV1, respectively. The horizontal axis represents the ratio of the number of the top sorted particles picked by our fully automated method vs. the number of manually picked particles in the reference set. The blue, green, red and black curves represent the results of using 10,000 manually picked particles as training data contributed roughly equally by one, two, three and four types of other molecules (which were different from the target molecule), respectively. The real and dashed curves represent the recall and precision scores, respectively. More details about different combinations of known particles from different molecules in training data can be found in Table 1.

3D model reconstruction (Langlois et al., 2014), we focused more on the recall metric.

A typical example of the fully automated picking results is shown in Fig. 2(a), which has demonstrated that most of the particles picked by our fully automated approach from a micrograph were reasonable. All different tests (Fig. 2(b)) show that our fully automated picking method can achieve high recall scores (above 0.81), which implies that most of the reference particles manually picked by human experts can also be identified by our deep learning framework with a cross-molecule training strategy. In addition, when the number of the top scoring particles picked by our fully automated approach was close to that of the reference particles, both recall and precision scores were relatively high (Fig. 2(c)–(e)). These results imply that a major fraction of the particles identified by our fully automated scheme were consistent with those manually picked by human experts. The comparisons among results by using different training datasets (see Section 2.4) show that our CNN classifier trained by multiple datasets achieved more robust performance than that using only a single dataset as training data (Fig. 2(c)–(e)). Such a result was expected, as the hidden features derived from multiple datasets were supposed to be more general. In practice, we can include as many known datasets as possible into training data to boost the classification performance of our deep learning model.

In the cross-molecule training strategy, we also tested different combinations of the known particles from other molecules as training data. In particular, we looked into the test results on the

deep learning classifier trained by the known particles of one, two, three and four types of molecules. Here we only show the typical results of each case (Fig. 2). More details about different combinations of training data are provided in Table 1. The tests of other combinations show similar results for individual cases.

All the above results indicate that, despite different contrast levels of micrographs, different shapes and sizes of the known particles from other molecules used as training data, our fully automated approach can effectively identify accurate particles that were comparable to the manual picking results and thus achieve a

Table 1
Different combinations of training data for the test results shown in Fig. 2.

For γ -secretase:	
1 type	TRPV1
2 types	TRPV1 and spliceosome
3 types	TRPV1, spliceosome and NSF fusion complex
4 types	TRPV1, spliceosome, NSF fusion complex and β -galactosidase
For spliceosome:	
1 type	γ -Secretase
2 types	γ -Secretase and TRPV1
3 types	γ -Secretase, TRPV1 and NSF fusion complex
4 types	γ -Secretase, TRPV1, NSF fusion complex and β -galactosidase
For TRPV1:	
1 type	Spliceosome
2 types	Spliceosome and γ -secretase
3 types	Spliceosome, γ -secretase and NSF fusion complex
4 types	Spliceosome, γ -secretase, NSF fusion complex and β -galactosidase

near human level of particle selection. Therefore, our fully automated particle picking method can be practically useful for cryo-EM structure determination. As it does not require any human intervention, in practice, it can help the user save a huge amount of time and effort in single-particle data analysis.

3.1.2. The 2D clustering and class averaging results

The 2D clustering and class averaging operations are commonly used right after the particle picking step in the single-particle data

analysis pipeline to further remove false positives and refine the list of the selected particles for 3D map reconstruction (Nogales and Scheres, 2015; Scheres, 2012). Thus, an additional effective method for evaluating the practicability of an automated particle picking approach is to further examine the 2D clustering and class averaging results of the identified particles. We have performed such an investigation and refined the original list of the picked particles by filtering those obviously bad-quality particles (see Section 2.5) during the 2D clustering and averaging operations as in




Fig. 3. The comparisons between the 2D clustering and class averaging results derived from the particles picked manually by human experts and fully automatically by our deep learning framework. (a), (c) and (e) The 2D clustering and class averaging results of the reference particles picked manually by human experts for γ -secretase, spliceosome and TRPV1, respectively. (b), (d) and (f) The 2D clustering and class averaging results of the particles picked fully automatically by our CNN classifier for γ -secretase, spliceosome and TRPV1, respectively. A mixture of 10,000 known particles of the other two molecules (each contributing to 5000 particles) among γ -secretase, spliceosome and TRPV1, which were different from the target molecule, were used as training data. Top 9 classes are shown here.

a standard single-particle data analysis pipeline. The comparison results have demonstrated that the 2D clustering and class averaging results of our fully automated picking method were comparable to those of the reference particles manually selected by human experts (Fig. 3). For example, the side-view particles of γ -secretase that were identified and used in the original paper (Sun et al., 2015) to reconstruct the near-atomic resolution 3D map were almost all present in our automated results (Fig. 3(a) and (b)). This implies that the particles automatically picked by our approach can represent a majority of good-quality 2D images that were practically useful for downstream data analysis in cryo-EM structure determination. Thus, our fully automated particle picking approach can provide a good starting point for 3D map reconstruction in cryo-EM.

3.1.3. Identification of particle centers

Identifying the accurate centers of particles is crucial for downstream single-particle data analysis, such as 2D clustering and 3D map reconstruction. In the previous 2D clustering and averaging results, we have shown that most of the particles picked by our fully automated approach were well centralized: almost all the averaging particle images selected by our program were fully

covered by the mask circle (Fig. 3). To further investigate the accuracy of our fully automated approach in identifying the centers of particles, we also measured the average distance with respect to the box size of sliding window between the particles picked fully automatically by our CNN classifier and manually by human experts (Fig. 4). We observed that the particles picked by our fully automated scheme deviated less than 10% of the box size from the reference images. These results imply that our fully automated approach can accurately detect the centers of particles that were quite close to those manually identified by human experts.

3.2. Semi-automated particle picking with an alternative training strategy

In our fully automated picking scheme, the known particles from other molecular complexes that are different from the current target molecule are used to train the deep learning classifier. An alternative training scheme is to let the user manually select a small number of particles as positive samples to train the CNN model and initialize the particle selection process (Fig. 5(a)). Such a scheme requires a certain level of human intervention and thus is considered semi-automated. As in the previous evaluation of fully

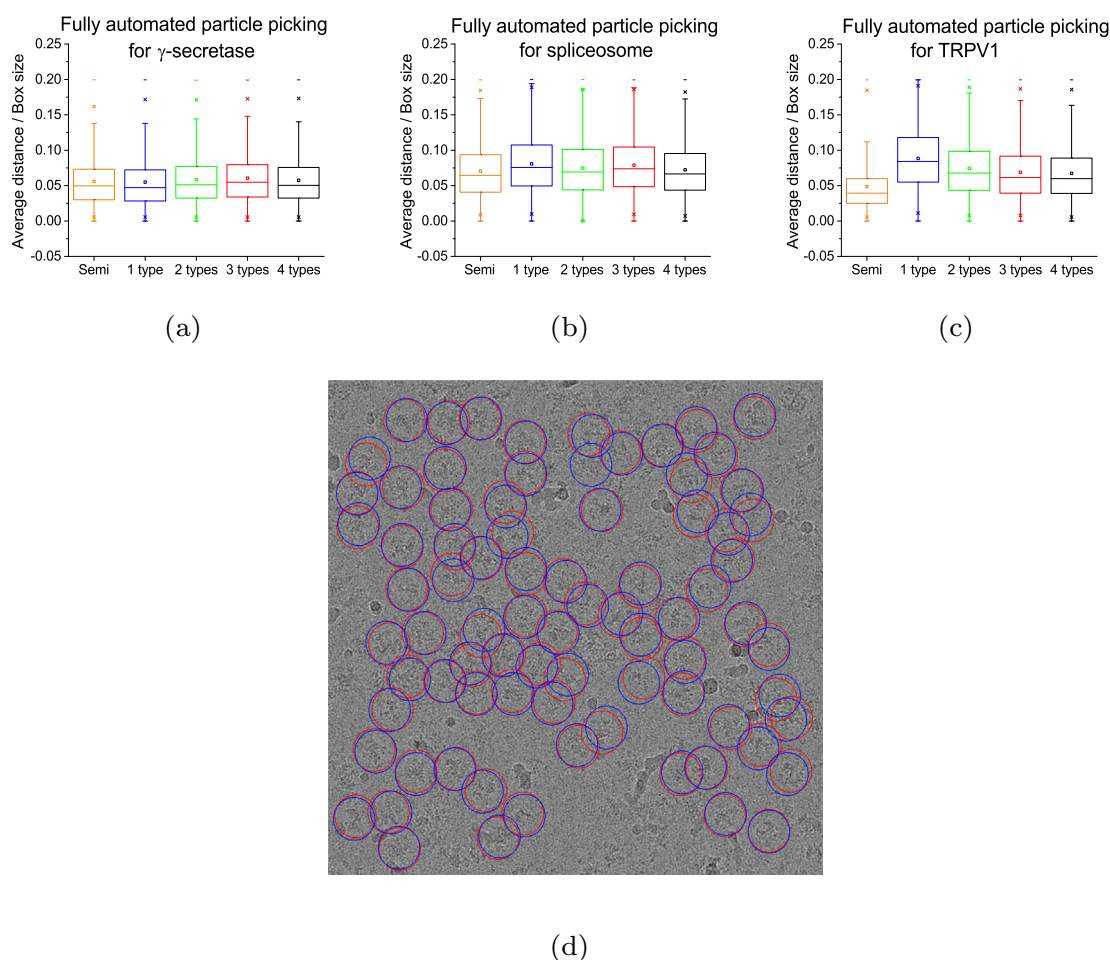


Fig. 4. Results on identifying the centers of particles. (a)–(c) The distance between our automatically picked particles and their manually picked counterparts in the reference set, normalized by the box size of sliding window, for γ -secretase, spliceosome and TRPV1, respectively. The orange plots represent the results of the CNN classifier trained by 10,000 particles of the same molecule (such a scheme is also called semi-automated picking, see also Section 3.2). The blue, green, red and black plots represent the results of the CNN classifier trained by a mixture of 10,000 known particles roughly equally contributed by one, two, three and four types of other molecules that were different from the target complex, respectively. (d) An example of the comparison between the spliceosome particles picked manually by human experts and fully automatically by our CNN classifier, which was trained by a mixture of 5000 γ -secretase and 5000 TRPV1 particles. The red and blue circles represent the particles selected by human experts and by our fully automated approach, respectively.

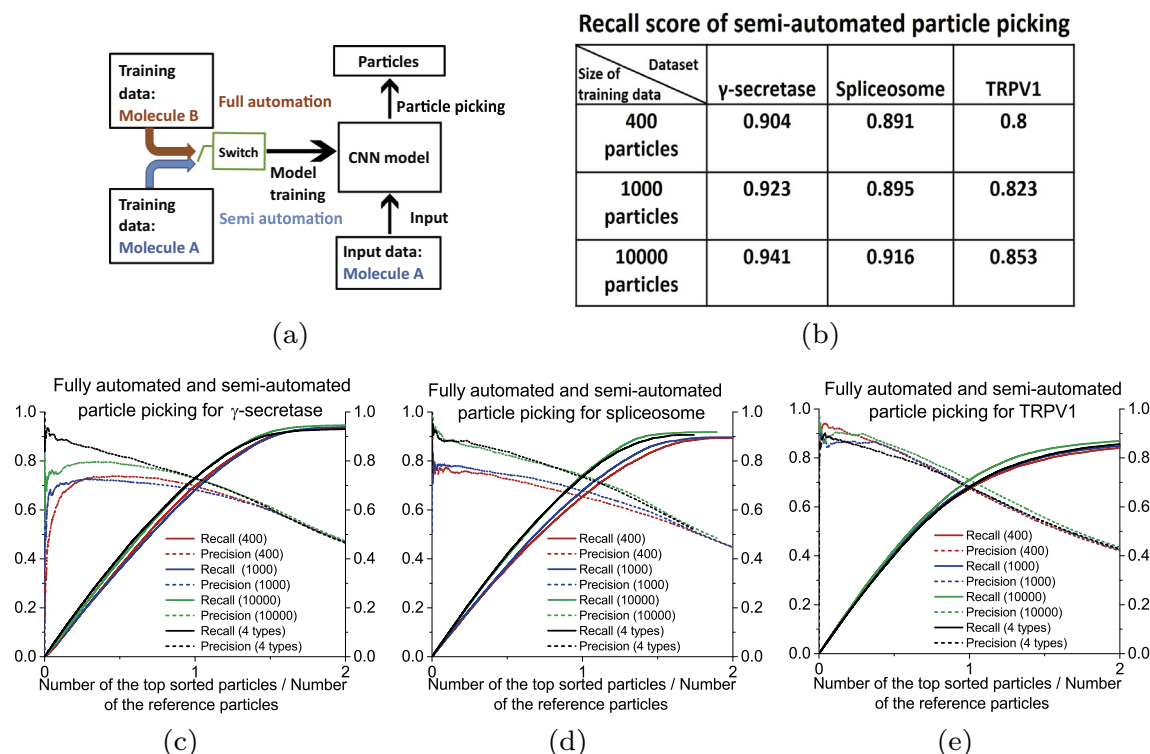


Fig. 5. Results on semi-automated particle picking of our deep learning framework with an alternative training scheme. (a) A simple illustration on the alternative choice between fully automated and semi-automated training schemes. (b) The summary of recall scores with prediction scores above 0.5 for the semi-automated particle picking scheme under our deep learning framework for three different datasets with various sizes of training data. (c)–(e) The trends of precision and recall vs. the number of the picked particles for γ -secretase, spliceosome and TRPV1, respectively. The horizontal axis represents the ratio between the number of the top sorted particles picked by both semi-automated and fully automated schemes vs. the number of manually picked particles in the reference set. The red, blue and green lines represent the semi-automated results of the CNN classifier trained by 400, 1000 and 10,000 particles (which were manually picked from micrographs of the target molecule), respectively. The black lines represent the fully automated results in which a mixture of the manually picked particles of four other molecules that were different from the target molecule were used as training data, each contributing to 2500 particles. The real and dashed lines represent the recall and precision scores, respectively.

automated particle picking, we also performed the similar tests for the semi-automated picking scheme on the same datasets of γ -secretase, spliceosome and TRPV1.

Overall, various tests on these three datasets have demonstrated that the semi-automated picking method resulting from the alternative training scheme can detect and identify a high percentage of correct particles that were consistent with the manual picking results by human experts (Fig. 5(b)). These results were similar to those in the previous tests of fully automated picking (Fig. 2(b)). The semi-automated picking results with 10,000 particles as training data were only slightly better than the fully automated picking results (Fig. 5(c)–(e)). We also observed similar trends of precision and recall curves with respect to the number of picked particles for both fully automated and semi-automated schemes (Fig. 5(b)–(e)). From the other perspective, our fully automated picking scheme can achieve performance close to that of semi-automated picking, and both schemes can yield outcomes that were comparable to those manual picking results by human experts.

Another noticeable observation in the test results of semi-automated particle picking is that, when the number of manually picked particles used in training data varied from 400 to 10,000, the performance of our semi-automated picking method did not significantly change (Fig. 5(b)–(e)). This result indicates that in practice the user only needs to select a small number (e.g., 400) of particles, which is sufficient enough to train our deep learning model for accurately picking new particles from the remaining micrographs. Thus, it can only demand a minimum of human intervention and can also provide

a practically useful tool for the current cryo-EM structure determination pipeline.

4. Discussion

The difficulty of picking particles from a micrograph is generally associated with the contrast level of the image, which is generally related to its defocus level. To examine the influence of the contrast levels on the automated particle selection results, we further examined the recall scores of our fully automated scheme at different defocus levels (Fig. 6). The tests on the micrographs of γ -secretase, spliceosome and TRPV1 with defocus ranging from 1 μ m to 3.5 μ m have shown that the performance of our fully automated particle picking method was relatively robust at different defocus levels (Fig. 6(a)–(c)). This implies that our automated method can identify a sufficient number of correct particles at various conditions of data collection and thus may have a wide range of applications in single-particle cryo-EM data analysis.

Recently, it has been pointed out that picking particles from background or random noise with extremely low thresholds might lead to the overfitting problem and fall into a potentially dangerous pitfall (Henderson, 2013; van Heel, 2013; Subramaniam, 2013), termed “Einstein-from-noise” (Shatsky et al., 2009), in which a photograph of Einstein was used as a template and pure noise images were aligned to this reference to reconstruct the Einstein image. It would be interesting to know whether our approach also suffers from this overfitting problem, when using the cross-molecule particles as training data and picking particles from

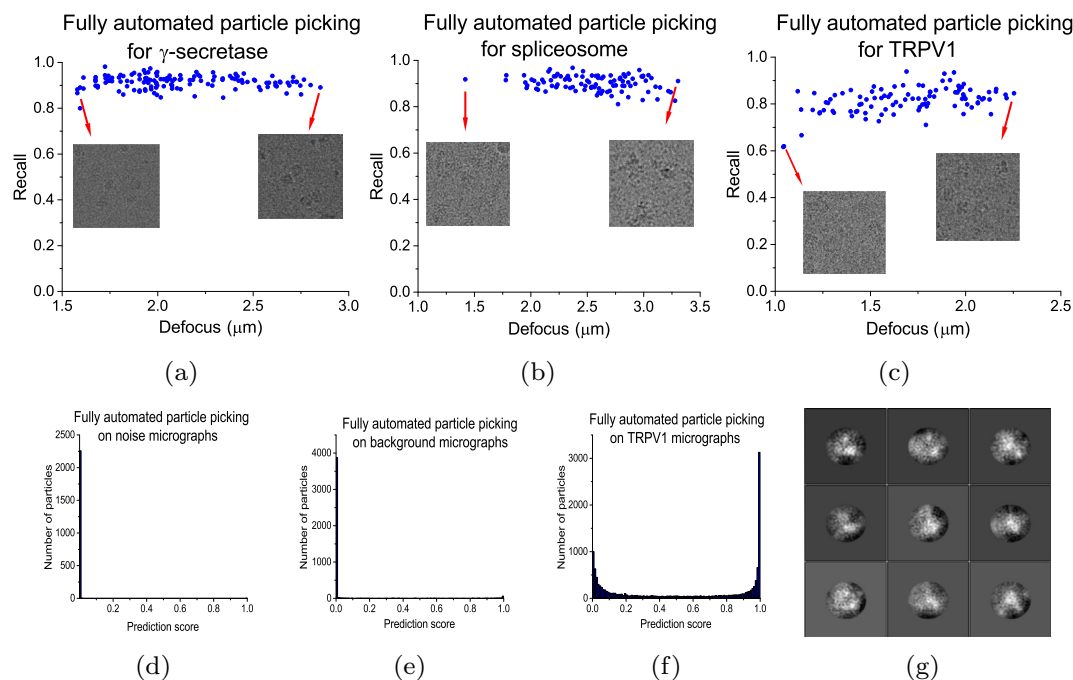


Fig. 6. The results of fully automated particle picking from micrographs at different defocus levels and from background or random noise. (a)–(c) The recall results of automated picking at different defocus levels for γ -secretase, spliceosome and TRPV1, respectively. The CNN classifier was trained by a mixture of 10,000 particles of other four molecules that were different from the target one, each contributing to 2500 particles. The inset panels in each subfigure show typical image examples at the minimum and maximum defocus levels, respectively. (d)–(f) The distributions of the prediction scores in our fully automated particle picking approach on pure noise (i.e., without adding any sample), pure background (i.e., without adding any protein in the sample) and TRPV1 micrographs, respectively. (g) The 2D class averaging results of the TRPV1 particles picked by DeepPicker using the prediction scores ranging between 0.1 and 0.5 in (f). Top 9 classes are shown here.

random noise. To investigate this issue, we tested our CNN model, which was trained by a mixture of 5000 γ -secretase and 5000 spliceosome particles, on pure *background* micrographs which were collected using the normal specimen preparation but without putting any protein on the glow-discharged holey carbon grids, and pure *noise* micrographs, which were measured from an empty cryo-EM instrument without putting any sample inside. In this additional test on background and noise images, we found that it was difficult to pick particles with the default threshold value (0.5) as used in the previous tests of γ -secretase, spliceosome or TRPV1 (see Section 3). We then set the threshold to zero and investigated the distribution of the prediction scores of the selected particles. We found that for both noise micrographs and background micrographs, almost all the picked particles were associated with a low prediction score (less than 0.05) (Fig. 6(d) and (e)), which were greatly different from our previous results, e.g., the test of TRPV1 (Fig. 6(f)). This result indicates that our scoring function was able to separate true particles from background or noise. Next, we performed the 2D clustering and class averaging processed on the particles picked by DeepPicker using the low prediction scores ranging from 0.1 to 0.5 in Fig. 6(f) and checked whether these picked particles can result in the overfitted images. We found that the class averaging results (Fig. 6(g)) were quite different from the training data in our model, which indicates that very unlikely our automated particle picking process with the cross-molecule training scheme would suffer from the previously-mentioned pitfall.

In addition to its own virtue in full automation, our deep learning framework with the cross-molecule training strategy has several additional advantages. First, unlike the subjective manual picking, the picking results of our deep learning classifier are objective and mainly determined by the in-house “computational” experts learned from training data. Second, as we apply a cross-molecule training method, in principle, we can incorporate as many available known particles from previously-solved complex

structures as possible into training data, which will thus increase the robustness of the training process and further improve the automated picking results. Third, in our fully automated picking scheme under the deep learning framework, the cross-molecule training process can be finished before collecting micrograph data of the current target molecule. Such an offline training manner is much more efficient than the online training methods (i.e., the classifiers are trained after data collection) in most of existing classification based frameworks. Moreover, the deep learning framework can take full advantage of current available powerful computing platforms, e.g., graphics processing unit (GPU), to speed up the learning process. Thus, our fully automated picking approach can be generally much more efficient and more robust than previous semi-automated methods.

In summary, we have proposed a deep learning framework for automated particle picking. The powerful deep learning model can be applied to pick particles either in a semi-automated manner or through a fully automated fashion. In the fully automated particle picking strategy, we innovatively use the known particles of the previously solved structures to train the deep learning model, which can accurately describe the common and cross-molecule features of particles. Both our fully automated and semi-automated particle picking schemes have been tested on the published cryo-EM datasets of several molecules, including γ -secretase, spliceosome and TRPV1. The comparisons between the particles picked by either our fully automated or semi-automated method and those manually identified by human experts have demonstrated that our deep learning framework can achieve near human-level performance in particle picking. Thus, our fully automated or semi-automated particle picking approach can provide a reliable and practically useful tool to liberate structural biologists from the time-consuming and laborious task of manual particle picking in cryo-EM data analysis. Our current version of DeepPicker is the first attempt to apply deep learning to address the

automated particle picking problem in cryo-EM. Several future directions may be pursued to improve the current particle picking results. First, some specific deep learning models, such as auto-encoder (Vincent et al., 2010), can be used for image denoising, which may help detect the correct particle objects from background or noise. Second, the cross-molecule knowledge about the common features of particles may be better exploited using a more sophisticated deep learning model and can also be combined with the 2D class averaging results to improve automated particle selection. Third, the current particle picking procedure may be combined with the downstream 3D map reconstruction steps in a fully automated and iterative manner to further improve the accuracy of particle selection.

Author contributions

XL and JZ initiated the project. JZ and XL designed the functions of the program. JZ, TX, FW and HG designed the deep learning algorithm. FW and HG implemented the program. FW, HG, GL and ML performed the testing experiment and analysis. CY acquired part of the data and provided corresponding reference images. JZ, HG and FW drafted the manuscript. JZ, XL and TX revised the manuscript.

Acknowledgements

We thank Prof. Yigong Shi, Prof. Nieng Yan, Dr. Minglei Zhao and Mr. Jianping Wu for generously sharing their cryo-EM data for testing our particle picking approach. We are grateful to Prof. Hongwei Wang, Prof. Jian Peng and Mr. Jinglin Yang for their insightful discussions about our work. We are also grateful to Ms. Rudan Chen for her help on our draft preparation. We thank Mr. Fangping Wan for testing the released version of our program. Part of the test data are from EMPIAR (<https://www.ebi.ac.uk/pdbe/emdb/empiar/>). This work was supported by the National Basic Research Program of China (Grant 2011CBA00300, 2011CBA00301 to JZ), the National Natural Science Foundation of China (Grant 61033001, 61361136003 and 61472205 to JZ, 31570730 to XL, and 61571202 to TX), China Youth 1000-Talent Program by the State Council of China (to JZ, XL and TX), Beijing Advanced Innovation Center for Structural Biology (to JZ and XL), and Tsinghua-Peking Joint Center for Life Sciences (to XL). We acknowledge the China National Center for Protein Sciences Beijing for providing the facility support and the support of NVIDIA Corporation with the donation of the Titan X GPU used for this research.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X., 2015. TensorFlow: Large-scale machine learning on heterogeneous systems, software available from tensorflow.org. URL <http://tensorflow.org/>.
- Adiga, P.U., Malladi, R., Baxter, W., Glaeser, R.M., 2004. A binary segmentation approach for boxing ribosome particles in cryo EM micrographs. *J. Struct. Biol.* 145 (12), 142–151. <http://dx.doi.org/10.1016/j.jsb.2003.10.026>.
- Arbelez, P., Han, B.-G., Typke, D., Lim, J., Glaeser, R.M., Malik, J., 2011. Experimental evaluation of support vector machine-based and correlation-based approaches to automatic particle selection. *J. Struct. Biol.* 175 (3), 319–328. <http://dx.doi.org/10.1016/j.jsb.2011.05.017>.
- Bai, X., McMullan, G., Scheres, S.H., 2015. How cryo-EM is revolutionizing structural biology. *Trends Biochem. Sci.* 40 (1), 49–57. <http://dx.doi.org/10.1016/j.tibs.2014.10.005>.
- Bartesaghi, A., Matthies, D., Banerjee, S., Merk, A., Subramaniam, S., 2014. Structure of β -galactosidase at 3.2-Å resolution obtained by cryo-electron microscopy. *Proc. Nat. Acad. Sci.* 111 (32), 11709–11714. <http://dx.doi.org/10.1073/pnas.1402809111>.
- Bishop, C.M., 2006. *Pattern Recognition and Machine Learning* (Information Science and Statistics). Springer-Verlag New York, Inc., Secaucus, NJ, USA.
- Bottou, L., 2010. Large-scale machine learning with stochastic gradient descent. In: Lechevallier, Y., Saporta, G. (Eds.), *Proceedings of COMPSTAT'2010*. Physica-Verlag HD, pp. 177–186. http://dx.doi.org/10.1007/978-3-7908-2604-3_16.
- Canny, J., 1986. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.* PAMI-8 (6), 679–698. <http://dx.doi.org/10.1109/TPAMI.1986.4767851>.
- Chen, J.Z., Grigorieff, N., 2007. SIGNATURE: A single-particle selection system for molecular electron microscopy. *J. Struct. Biol.* 157 (1), 168–173. <http://dx.doi.org/10.1016/j.jsb.2006.06.001>.
- Cheng, Y., 2015. Single-particle cryo-EM at crystallographic resolution. *Cell* 161 (3), 450–457. <http://dx.doi.org/10.1016/j.cell.2015.03.049>.
- Cheng, Y., Grigorieff, N., Penczek, P.A., Walz, T., 2015. A primer to single-particle cryo-electron microscopy. *Cell* 161 (3), 438–449. <http://dx.doi.org/10.1016/j.cell.2015.03.050>.
- Cheng, M.-M., Mitra, N.J., Huang, X., Torr, P.H.S., Hu, S.-M., Zhang, G.-X., Mitra, N.J., Huang, X., Hu, S.-M., 2015. Global contrast based salient region detection. *IEEE Trans. Pattern Anal. Mach. Intell.* 37 (XX), 569–582. <http://dx.doi.org/10.1109/TPAMI.2014.2345401>, URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6871397>.
- Cun, Y.L., Boser, B., Denker, J.S., Howard, R.E., Habbard, W., Jackel, L.D., Henderson, D., 1990. *Advances in Neural Information Processing Systems*, vol. 2. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, Ch. Handwritten Digit Recognition with a Back-propagation Network, pp. 396–404.
- Farabet, C., Couprie, C., Najman, L., LeCun, Y., 2013. Learning hierarchical features for scene labeling. *IEEE Trans. Pattern Anal. Mach. Intell.* 35 (8), 1915–1929. <http://dx.doi.org/10.1109/TPAMI.2012.231>.
- Goldberg, Y., 2015. A primer on neural network models for natural language processing, CoRR abs/1510.00726.
- Grigorieff, N., 2007. FREALIGN: High-resolution refinement of single particle structures. *J. Struct. Biol.* 157 (1), 117–125. <http://dx.doi.org/10.1016/j.jsb.2006.05.004>.
- Hall, R.J., Patwardhan, A., 2004. A two step approach for semi-automated particle selection from low contrast cryo-electron micrographs. *J. Struct. Biol.* 145 (12), 19–28. <http://dx.doi.org/10.1016/j.jsb.2003.10.024>.
- Henderson, R., 2013. Avoiding the pitfalls of single particle cryo-electron microscopy: Einstein from noise. *Proc. Nat. Acad. Sci.* 110 (45), 18037–18041. <http://dx.doi.org/10.1073/pnas.1314449110>.
- Hinton, G., Deng, L., Yu, D., Dahl, G.E., Mohamed, A.-r., Jaitly, N., Senior, A., Vanhoucke, V., Nguyen, P., Sainath, T.N., Kingsbury, B., 2012. Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups. *Signal Process. Mag.*, IEEE 29 (6), 82–97. <http://dx.doi.org/10.1109/msp.2012.2205597>.
- Huang, Z., Penczek, P.A., 2004. Application of template matching technique to particle detection in electron micrographs. *J. Struct. Biol.* 145 (12), 29–40. <http://dx.doi.org/10.1016/j.jsb.2003.11.004>.
- Jarrett, K., Kavukcuoglu, K., Ranzato, M., LeCun, Y., 2009. What is the best multi-stage architecture for object recognition? Proceedings of the IEEE International Conference on Computer Vision, 2146–2153. <http://dx.doi.org/10.1109/ICCV.2009.5459469>.
- Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. ImageNet classification with deep convolutional neural networks. In: Pereira, F., Burges, C., Bottou, L., Weinberger, K. (Eds.), *Advances in Neural Information Processing Systems*, vol. 25. Curran Associates, Inc., pp. 1097–1105.
- Langlois, R., Pallesen, J., Frank, J., 2011. Reference-free particle selection enhanced with semi-supervised machine learning for cryo-electron microscopy. *J. Struct. Biol.* 175 (3), 353–361. <http://dx.doi.org/10.1016/j.jsb.2011.06.004>.
- Langlois, R., Pallesen, J., Ash, J.T., Ho, D.N., Rubinstein, J.L., Frank, J., 2014. Automated particle picking for low-contrast macromolecules in cryo-electron microscopy. *J. Struct. Biol.* 186 (1), 1–7. <http://dx.doi.org/10.1016/j.jsb.2014.03.001>.
- Lawrence, S., Giles, C.L., Tsoi, A.C., Back, A.D., 1997. Face recognition: A convolutional neural network approach. *IEEE Trans. Neural Networks* 8, 98–113. <http://dx.doi.org/10.1109/72.554195>.
- LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., 1998. Gradient-based learning applied to document recognition. In: *Proc. IEEE*, pp. 2278–2324.
- LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. *Nature* 521, 436–444. <http://dx.doi.org/10.1038/nature14539>.
- Liao, M., Cao, E., Julius, D., Cheng, Y., 2013. Structure of the TRPV1 ion channel determined by electron cryo-microscopy. *Nature* 504 (7478), 107–112. <http://dx.doi.org/10.1038/nature12822>.
- Li, X., Mooney, P., Zheng, S., Booth, C.R., Braumfeld, M.B., Gubbens, S., Agard, D.A., Cheng, Y., 2013. Electron counting and beam-induced motion correction enable near-atomic-resolution single-particle cryo-EM. *Nat. Meth.* 10 (6), 584–590. <http://dx.doi.org/10.1038/nmeth.2472>.
- Ma, J., Sheridan, R.P., Liaw, A., Dahl, G.E., Svetnik, V., 2015. Deep neural nets as a method for quantitative structure activity relationships. *J. Chem. Inf. Model.* 55 (2), 263–274. <http://dx.doi.org/10.1021/ci500747n>.
- Mikolov, T., Deoras, A., Povey, D., Burget, L., Cernocky, J., 2011. Strategies for training large scale neural network language models. In: Nahamoo, D., Picheny, M. (Eds.), *ASRU*. IEEE, pp. 196–201.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., Hassabis, D., 2015. Human-level control through deep reinforcement learning. *Nature* 518 (7540), 529–533. <http://dx.doi.org/10.1038/nature14236>.

- Nair, V., Hinton, G.E., 2010. Rectified linear units improve restricted boltzmann machines. In: Frnkranz, J., Joachims, T. (Eds.), *ICML*. Omnipress, pp. 807–814.
- Nicholson, W.V., Glaeser, R.M., 2001. Review: Automatic particle detection in electron microscopy. *J. Struct. Biol.* 133 (23), 90–101. <http://dx.doi.org/10.1006/jsbi.2001.4348>.
- Nogales, E., Scheres, S.H., 2015. Cryo-EM: A unique tool for the visualization of macromolecular complexity. *Mol. Cell* 58 (4), 677–689. <http://dx.doi.org/10.1016/j.molcel.2015.02.019>.
- Norouzi, R., Wickles, S., Leidig, C., Becker, T., Schmid, V.J., Beckmann, R., Tresch, A., 2013. Automatic post-picking using MAPPOS improves particle image detection from cryo-EM micrographs. *J. Struct. Biol.* 182 (2), 59–66. <http://dx.doi.org/10.1016/j.jsb.2013.02.008>.
- Ogura, T., Sato, C., 2001. An automatic particle pickup method using a neural network applicable to low-contrast electron micrographs. *J. Struct. Biol.* 136 (3), 227–238. <http://dx.doi.org/10.1006/jsbi.2002.4442>, URL <http://www.sciencedirect.com/science/article/pii/S1047847702944424>.
- Rohou, A., Grigorieff, N., 2015. CTFFIND4: Fast and accurate defocus estimation from electron micrographs. *J. Struct. Biol.* 192 (2), 216–221. <http://dx.doi.org/10.1016/j.jsb.2015.08.008>.
- Scherer, D., Müller, A., Behnke, S., 2010. Evaluation of pooling operations in convolutional architectures for object recognition. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 6354 LNCS (PART 3), 92–101. http://dx.doi.org/10.1007/978-3-642-15825-4_10.
- Scheres, S.H., 2012. RELION: Implementation of a bayesian approach to cryo-EM structure determination. *J. Struct. Biol.* 180 (3), 519–530. <http://dx.doi.org/10.1016/j.jsb.2012.09.006>.
- Shatsky, M., Hall, R.J., Brenner, S.E., Glaeser, R.M., 2009. A method for the alignment of heterogeneous macromolecules from electron microscopy. *J. Struct. Biol.* 166 (1), 67–78. <http://dx.doi.org/10.1016/j.jsb.2008.12.008>.
- Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., Hassabis, D., 2016. Mastering the game of Go with deep neural networks and tree search. *Nature* 529 (7587), 484–489. <http://dx.doi.org/10.1038/nature16961>.
- Sorzano, C., Recarte, E., Alcorlo, M., Bilbao-Castro, J., San-Martín, C., Marabini, R., Carazo, J., 2009. Automatic particle selection from electron micrographs using machine learning techniques. *J. Struct. Biol.* 167 (3), 252–260. <http://dx.doi.org/10.1016/j.jsb.2009.06.011>.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15 (1), 1929–1958.
- Subramaniam, S., 2013. Structure of trimeric HIV-1 envelope glycoproteins. *Proc. Nat. Acad. Sci.* 110 (45), E4172–E4174. <http://dx.doi.org/10.1073/pnas.1313802110>.
- Sun, L., Zhao, L., Yang, G., Yan, C., Zhou, R., Zhou, X., Xie, T., Zhao, Y., Wu, S., Li, X., Shi, Y., 2015. Structural basis of human γ -secretase assembly. *Proc. Nat. Acad. Sci. USA* 112 (19), 6003–6008. <http://dx.doi.org/10.1073/pnas.1506242112>.
- Sutskever, I., Vinyals, O., Le, Q.V., 2014. Sequence to sequence learning with neural networks. *CoRR abs/1409.3215*.
- Tompson, J.J., Jain, A., Lecun, Y., Bregler, C., 2014. Joint training of a convolutional network and a graphical model for human pose estimation. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., Weinberger, K. (Eds.), *Advances in Neural Information Processing Systems*, vol. 27. Curran Associates, Inc., pp. 1799–1807.
- van Heel, M., 2013. Finding trimeric HIV-1 envelope glycoproteins in random noise. *Proc. Nat. Acad. Sci. USA* 110 (45), E4175–E4177. <http://dx.doi.org/10.1073/pnas.1314353110>.
- Vincent, P., Larochelle, H., Lajoie, I., Bengio, Y., Manzagol, P.-A., 2010. Stacked denoising autoencoders: learning useful representations in a deep network with a local denoising criterion. *J. Mach. Learn. Res.* 11 (3), 3371–3408. <http://dx.doi.org/10.1111/1467-8535.00290>, arXiv:0-387-31073-8.
- Voss, N., Yoshioka, C., Radermacher, M., Potter, C., Carragher, B., 2009. DoG picker and TiltPicker: Software tools to facilitate particle selection in single particle electron microscopy. *J. Struct. Biol.* 166 (2), 205–213. <http://dx.doi.org/10.1016/j.jsb.2009.01.004>.
- Wu, J., Yan, Z., Li, Z., Yan, C., Lu, S., Dong, M., Yan, N., 2015. Structure of the voltage-gated calcium channel Cav1.1 complex. *Science* 350 (5), aad2395. <http://dx.doi.org/10.1126/science.aad2395>.
- Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhutdinov, R., Zemel, R., Bengio, Y., 2015. Show, attend and tell: Neural image caption generation with visual attention. In: Blei, D., Bach, F. (Eds.), *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, JMLR Workshop and Conference Proceedings, pp. 2048–2057.
- Yan, C., Hang, J., Wan, R., Huang, M., Wong, C.C.L., Shi, Y., 2015. Structure of a yeast spliceosome at 3.6-angstrom resolution. *Science* 349 (6253), 1182–1191. <http://dx.doi.org/10.1126/science.aac7629>.
- Zeiler, M.D., Fergus, R., 2013. Visualizing and understanding convolutional networks. *CoRR abs/1311.2901*.
- Zhao, J., Brubaker, M.A., Rubinstein, J.L., 2013. TMacS: A hybrid template matching and classification system for partially-automated particle selection. *J. Struct. Biol.* 181 (3), 234–242. <http://dx.doi.org/10.1016/j.jsb.2012.12.010>.
- Zhao, M., Wu, S., Zhou, Q., Vivona, S., Cipriano, D.J., Cheng, Y., Brunger, A.T., 2015. Mechanistic insights into the recycling machine of the SNARE complex. *Nature* 518 (7537), 61–67. <http://dx.doi.org/10.1038/nature14148>.
- Zhao, J., Benlekhir, S., Rubinstein, J.L., 2015. Electron cryomicroscopy observation of rotational states in a eukaryotic V-ATPase. *Nature* 521 (5), 241–245. <http://dx.doi.org/10.1038/nature14365>.
- Zhou, J., Troyanskaya, O.G., 2015. Predicting effects of noncoding variants with deep learning-based sequence model. *Nat. Meth.* 12 (10), 931–934. <http://dx.doi.org/10.1038/nmeth.3547>.
- Zhu, Y., Carragher, B., Glaeser, R.M., Fellmann, D., Bajaj, C., Bern, M., Mouche, F., de Haas, F., Hall, R.J., Kriegman, D.J., Ludtke, S.J., Mallick, S.P., Penczek, P.A., Roseman, A.M., Sigworth, F.J., Volkman, N., Potter, C.S., 2004. Automatic particle selection: results of a comparative study. *J. Struct. Biol.* 145 (12), 3–14. <http://dx.doi.org/10.1016/j.jsb.2003.09.033>.