DeepEMhacer: a deep learning solution for cryo-EM volume post-processing

R Sanchez-Garcia[1], J Gomez-Blanco[2], A Cuervo[1], JM Carazo[1], COS Sorzano[1*], J Vargas[2,3*]

[1] Biocomputing Unit, Centro Nacional de Biotecnología-CSIC C/ Darwin 3, 28049 Cantoblanco, Madrid, Spain

[2] Departament of Anatomy and Cell Biology, McGill University 3640 Rue University, Montréal, QC H3A 0C7, Canada

[3] Departamento de Óptica, Universidad Complutense de Madrid, Avda. Complutense s/n, 28040 Ciudad Universitaria Madrid, Spain, 28040

*Correspondence to coss@cnb.csic.es, jvargas@fis.ucm.es

Abstract

Cryo-electron microscopy (cryo-EM) maps are among the most valuable sources of information for protein structure modeling. However, due to the loss of contrast at high frequencies, they generally need to be post-processed before modeling in order to improve their interpretability. To that end, approaches based on B-factor correction are the most popular choices, yet they suffer from some limitations such as the fact that the correction is applied globally, ignoring the presence of heterogeneity in the map local quality that cryo-EM reconstructions tend to exhibit.

With the aim of overcoming these limitations, here we present DeepEMhacer, a deep learning approach designed to perform automatic post-processing of cryo-EM maps. Trained on a dataset of pairs of experimental cryo-EM maps and maps sharpened by LocScape using their respective atomic models, DeepEMhacer has automatically learned how to post-process experimental maps performing masking-like and sharpening-like operations in a single step. DeepEMhacer has been evaluated on a testing set of 20 different experimental maps, showing its ability to obtain much cleaner and detailed versions of the experimental maps, thus, improving their interpretability. Additionally, we have illustrated the benefits of DeepEMhacer with a use case in which the structure of the SARS-CoV 2 RNA polymerase is improved.

Introduction

Almost one decade after the beginning of the so-called "resolution revolution", cryogenic electron microscopy (cryo-EM) has become one of the most versatile tools in the field of structural biology. Beginning from thousands of single particle projection images, cryo-EM workflows are capable of obtaining three-dimensional (3D) reconstructions of many macromolecules at "near-atomic" resolution levels. However, the ultimate goal of cryo-EM Single Particle Analysis is not the obtention of 3D maps but the detailed atomic understanding through the derivation of atomic models.

During the atomic model building process, raw 3D maps are rarely employed, as they suffer from loss of contrast at high resolution (Rosenthal and Henderson, 2003) that makes difficult the detection and interpretability of residues and secondary structure. Fortunately, loss of contrast can be alleviated using different contrast restoration algorithms, which are usually known as sharpening methods. The first sharpening approach for cryo-EM maps was introduced by Rosenthal and Henderson (Rosenthal and Henderson, 2003) and its formulation, based on the B-factor correction, is still at the basis of the most commonly employed sharpening methods, including RELION postprocessing (Kimanius et al., 2016; Zivanov et al., 2018) or Phenix AutoSharpen (Terwilliger et al., 2018). The principle behind these algorithms consists in the correction of the raw maps by boosting the amplitude of their high frequency Fourier components. The strength of the amplitude boost at each frequency depends on the frequency itself and on a single number, the B-factor, that measures the global loss of contrast. Thus, although the different B-factor-based methods differ in the procedures employed to determine the B-factor that is applied, all of them modify the volume globally in a similar manner.

Despite being widely used, B-factor-based approaches present an important limitation: they do not consider the differences in quality that different parts of the map may present and they produce density maps that do not correspond to the scattering properties of biological macromolecules (Vilas et al., 2020). Consequently, for the case of maps that exhibit heterogeneous local resolution, some regions could be undersharpened whereas others could be oversharpened. Recently, local sharpening algorithms, that alleviate this shortcoming, have been proposed. Thus, the LocScale (Jakobi et al., 2017) algorithm uses the information contained in an atomic model to locally scale up a map. Such transformation is achieved by means of a sliding window approach in which the amplitudes of the map region that lay inside the window are scaled up to agree with the atomic model provided. Following a totally different strategy, the LocalDeblur (Ramírez-Aportela et al., 2020) algorithm employs a Wiener filtering approach that performs local deblurring with a strength proportional to an estimation of the local resolution, that has to be pre-computed. Similarly, LocSpiral (Kaur et al., 2020) employs the spiral phase transformation to factorize the volume and then, perform a local enhancement based on the normalization and thresholding of the amplitudes.

Despite their benefits, current local sharpening approaches present some drawbacks. Thus, both LocSpiral and LocalDeblur depend on masks to distinguish the macromolecule from the noise and LocalDeblur requires also from an estimation of the local resolution of the map. On the other hand, the main strength of LocScale, its ability to employ the structural information of atomic models, could also be regarded as its main weakness since the availability of atomic models limits its applicability.

With the aim of overcoming these shortcomings, in this work, we present Deep cryo-EM Map Enhancer (DeepEMhacer), a fully automatic deep learning-based approach that performs cryo-EM volume post-processing. Deep learning has revolutionized the field of Artificial Intelligence and its impact has been felt in many others including cryo-EM. Deep learning in cryo-EM was firstly applied for the problem of particle picking (Wagner et al., 2019; Wang et al., 2016; Zhu et al., 2017) and since then, it has evolved to deal with other questions such as map reconstruction (Gupta et al., 2020; Zhong et al., 2019), map segmentation (Maddhuri Venkata Subramaniya et al., 2019; Si et al., 2020) or local resolution determination (Avramov et al., 2019; Ramírez-Aportela et al., 2019). As in most of those methods, our approach relies on a convolutional neural network (CNN) that is trained on massive quantities of data. Particularly, our development, that follows a simple image super-resolution setup (Yang et al., 2019), exploits the vast amount of structural information that is contained in the Electron Microscopy Data Bank (EMDB) database (Lawson et al., 2015) in order to mimic the local sharpening effect of the LocScale algorithm. However, DeepEMhacer does not require any atomic model to function and, contrary to previous methods, it also performs automatic (tight) masking of input maps. Our results show that DeepEMhacer, that works in a fully automatic manner, is able to largely improve the interpretability of the maps contained in our benchmark, performing better than classical B-factor approaches.

Methods

DeepEMhacer is based on an end-to-end U-net architecture (Ronneberger et al., 2015) trained in a supervised manner. Thus, pairs of input maps (X) and target maps (Y), consisting in tightly masked LocScale post-processed maps, are required to that end. The following sections describe in detail the data preparation, training and evaluation processes.

Raw data collection:

DeepEMhacer has been trained and evaluated using as input a subset of cryo-EM maps obtained from the EMDB (Lawson et al., 2015) that meet the following requirements: 1) resolution better than 7 Å; 2) have one and only one atomic model associated; 3) correlation between the atomic model and the map better than 0.6 and 4) half maps available. As a result, an original list of 415 maps was compiled. However, this initial list is highly redundant and, in order to avoid biases in both the training and evaluation procedures, this list was further filtered to reduce its redundancy, (see subsection *Redundancy control*). Finally, after a visual inspection aimed at removing problematic cases that survived to the automatic filtering procedure, a total amount of 151 maps, with an average reported resolution of 3.8 Å, was selected.

Since the main objective of DeepEMhacer is to perform a sharpening-like post-processing transformation, it is important to ensure that the maps used in this study were not previously sharpened. Given the fact that most of the maps deposited in EMDB are sharpened and many are also masked, we decided to employ only the half-maps available in EMDB (condition number 4). Due to the lack of an appropriate searching tool in EMDB and a file name convention, we had to analyze all the map file names included in the database looking for the substring "half" to recover the half maps. Full maps were obtained averaging respective half maps.

As learning targets, we employed the output generated by LocScale using as input the aforementioned maps and their associated atomic models. Additionally, the output maps were tightly masked using as masks the maps simulated from the atomic models after a thresholding operation. Although it is true that the simulated maps could be directly employed as targets, we discarded this alternative for two reasons. The first one is empirical: we obtained better results when targets were produced with LocScale, probably because the input and target maps, although different, they still share some similar properties such as intensity ranges or local quality, which are not necessarily preserved when using simulated maps as targets. The other reason is that we wanted to reproduce the state-of-the-art local sharpening effect and not a new type of post-processing that could not be compatible with downstream atomic modelling tools.

Data preparation:

Due to the fact that the monomers (amino acids, nucleotides…) that compose the macromolecules have fixed size but the deposited maps vary in voxel size, both the input and the target maps were resampled to 1 Å/voxel size with the aim of facilitating the learning process. After that, the intensity of each volume was normalized using the classical cryo-EM approach by which the map noise statistics are forced to adopt a fixed mean and standard deviation (0 and 0.1 respectively). Finally, due to GPU size limitations, the maps were chunked into 64x64x64 cubes, the maximum size that our computing systems were able to efficiently manage. As a result, more than 70k volume cubes, including both signal cubes and noise-only cubes were used for training.

Redundancy control:

In order to perform the train/test/validation split used to develop and evaluate our method, it is important to consider that the universe of proteins is highly redundant and that the EMDB entries are even more redundant. Serve as an example the case of the ribosome, that supposes ~10% of the all EMDB entries. Thus, in order to avoid an over-optimistic performance estimation, we have ensured that the train, test and validation sets are mutually exclusive in the sense that their intersections are empty under a certain equivalence criterion. Particularly, we consider that two EMDB entries are equivalent if they share one sequence that belongs to the same 30% sequence identity cluster. Similarly, with the aim of eliminating potential bias in the evaluation, we have guaranteed that only one member per cluster is included in testing and validation sets. On the contrary, we have relaxed our quite strict redundancy control policy in the training set allowing up to five cluster representatives in an attempt to increase the size of this set. This decision is founded on the fact that even maps of the same exact protein may present different statistics due to the intrinsic variability of cryo-EM reconstruction workflows and thus, limiting their presence in the training set may difficult the generalization of the neural network.

As a result, a list of 110, 21 and 20 maps were used for training, validation and testing respectively. The full list of the EMDB entries used can be found in Supplementary Material.

Neural network architecture:

We have employed a 3D U-net-like neural network (Ronneberger et al., 2015) as a regression model for the estimation of post-processed maps. Our neural network consists of three downsampling blocks and three upsampling blocks with skip connections. Each block contains three convolutional layers followed by group normalization (Wu and He, 2020) and PRelu activation (He et al., 2015). The number of filters for each block is 3x32, 3x64 and 3x128 respectively. Downsampling is carried out using strided convolutions and upsampling is performed via transposed convolution. See Supplementary Material for additional details.

Neural network training:

Our neural network was trained using stochastic gradient descent with a batch size of 8 cubes. Initial learning rate was set to 10^{-3} and decreased by a factor of 0.5 when the validation loss did not improve during 5 epochs. As loss function, mean absolute error was employed. Data augmentation, consisting in random 90º rotations, gaussian blurring and patch corruption was applied to the training data.

Neural network inference:

In order to perform volume post-processing, the input volume is pre-processed as described in the *Data preparation* subsection. Then, the resized and normalized volume is chunked into overlapping cubes of size 64x64x64 with strides of 16 voxels. Each cube is individually processed by the trained neural network, yielding post-processed cubes. After that, the post-processed cubes are re-assembled into the final volume averaging the overlapping parts. Finally, the processed volume is resized to the size of the original volume, thus, showing the correct sampling rate value.

Evaluation:

With the aim of guiding the cross-validation process, we computed the correlation coefficient between the maps produced by DeepEMhacer and the maps used as learning targets (masked LocScale post-processed maps). Once the final model was selected, the quality of DeepEMhacer predictions were assessed comparing the input and processed maps against the reference maps obtained from the atomic models. Specifically, we computed the Fourier Shell Correlation coefficient (FSC) between them and we estimated the resolution using 0.5 as threshold. Due to the fact that DeepEMhacer performs a non-conventional post-processing operation, including masking and enhancement operations, in order to disentangle the two effects, the FSC was also computed after masking the maps to compare with a tight mask derived from the atomic model.

As a complementary metric, we also applied DeepRes (Ramírez-Aportela et al., 2019) over the input and processed maps. DeepRes is a deep learning-based local resolution method that, contrary to others, is sensitive to the sharpening

process and thus, it can provide an alternative estimation of the post-processing effect.

Finally, for comparison purposes, we repeated the FSC and DeepRes experiments using the Relion postprocessing program (Kimanius et al., 2016; Zivanov et al., 2018). As Relion automatic masking is very simple, in order to make the comparison more interesting, we decided to execute the postprocessing algorithm using the mask derived from the atomic models. Similarly, since the automatic determination of the B-factor can produce worse results than a manually selected one, in addition to the maps computed using an automatically determined B-factor by Relion, we also considered the sharpened map deposited in EMDB.

EMD-30178 map evaluation and atomic model modification:

DeepEMhacer was applied to the half maps deposited in EMDB entry EMD-30178. The original and post-processed maps were visually inspected using Coot (Emsley and Cowtan, 2004) and chimera (Pettersen et al., 2004), and chosen regions on the 7btf PDB were newly built or modified using Coot.
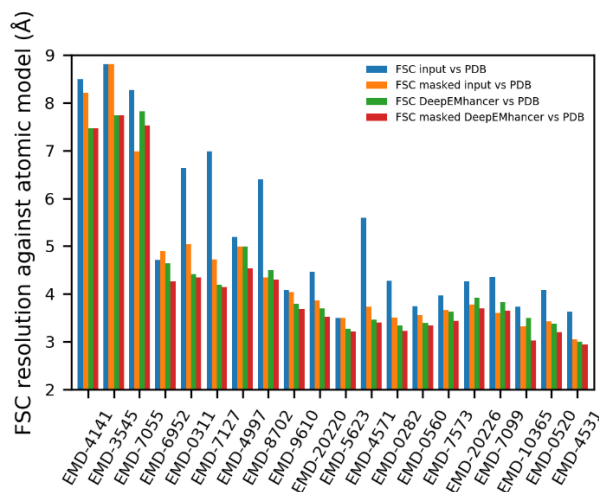

Figure 1. DeepEMhacer produces maps that are more similar to the atomic models. Resolution (determined by Fourier Shell Correlation coefficient, FSC) between the reference maps obtained from the atomic model and 1) the input maps (blue), 2) the input maps tightly masked (orange), 3) the post-processed maps (green) and 4) the post-processed maps tightly masked (red). EMDB entries are sorted by published global resolution.

Results

## Deep cryo Map Enhancer performance on the testing set.

In order to assess the quality of DeepEMhacer predictions, we first compared them against the target maps generated by LocScale. Thus, for DeepEMhacer maps, we measured a median correlation coefficient of 0.91 in contrast to 0.60 for input maps (see Supplementary Material Figure S1). Such an important increase in the correlation coefficient implies that DeepEMhacer has learned to accurately reproduce the effect of LocScale sharpening with one important advantage: no atomic models are required to employ DeepEMhacer.
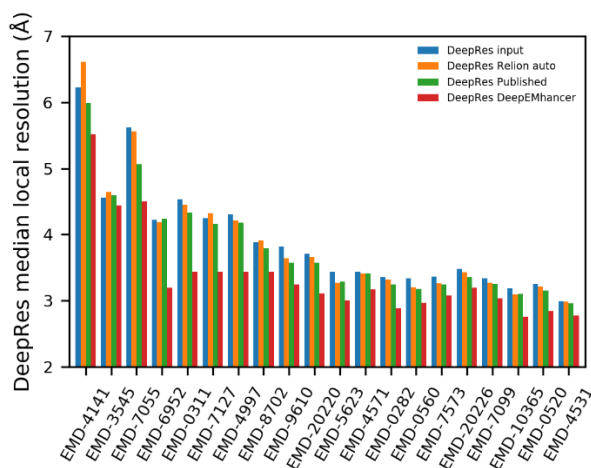
Figure 2. DeepEMhacer produces better quality maps. DeepRes median local resolution estimation for 1) the input maps (blue), 2) the post-processed maps obtained with Relion postprocessing automatic B-factor (orange), 3) the post-processed maps deposited in EMDB (green) and 4) the post-processed maps obtained with DeepEMhacer (Red). EMDB entries are sorted by published global resolution.

Although reproducing the LocScale sharpening effect was our main objective, the ultimate goal of map post-processing is to simplify the process of atomic model building. With the aim of studying if DeepEMhacer also contributes to that purpose, we next explored whether DeepEMhacer post-processed maps were more similar to the actual atomic models. To do so, we computed, for all the maps included in the testing set, the Fourier Shell Correlation coefficient (FSC) between the input and post-processed maps against the reference maps obtained from the atomic models. As it is shown in Figure 1, for all the examples included in the testing set, the application of DeepEMhacer

increased the similarity of the input maps with respect to the references (blue and green bars). Indeed, the post-processed maps exhibit a median FSC resolution value of 3.3 Å compared to 3.9 Å for the input maps. Particularly, the median improvement achieved by DeepEMhacer was ~0.6 Å (~14% in the frequency domain). Such an important improvement confirms that the maps computed by DeepEMhacer are much more similar to the target maps.
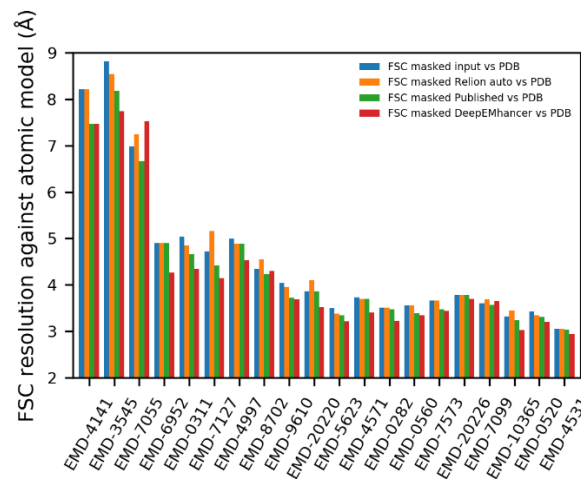

Figure 3. DeepEMhacer produces better results than B-factor-based methods. Resolution (determined by Fourier Shell Correlation coefficient, FSC) between the reference maps obtained from the atomic model and 1) the input maps (blue), 2) the post-processed maps obtained with Relion postprocessing automatic B-factor (orange), 3) the post-processed maps deposited in EMDB (green) and 4) the post-processed maps obtained with DeepEMhacer (red). EMDB entries are sorted by published global resolution.

DeepEMhacer post-processing operation performs a non-linear transformation of the experimental volume that produces a set of effects that could be broadly classified as masking/denoising and sharpening-like features enhancement. In order to disentangle the contribution of the different effects, we have also computed the FSC of the input and post-processed maps using a very tight mask derived from the atomic model. As it can be observed in Figure 1, the FSC resolution obtained for the post-processed maps tend to be better than the values computed for the input independently of the mask application (green and red bars vs orange bar), which implies that the masking effect is of high-quality, as the resolutions for the unmasked DeepEMhacer results tend to be better than the ones for the masked input maps.

Similarly, and, although it is true that the trend is not as strong as in the previous experiment, DeepEMhacer also tends to improve the resolution of the masked regions (Figure 1, orange vs red bars), which supposes an enhancement of the map features. Leaving aside some problematic examples such as EMD-7055
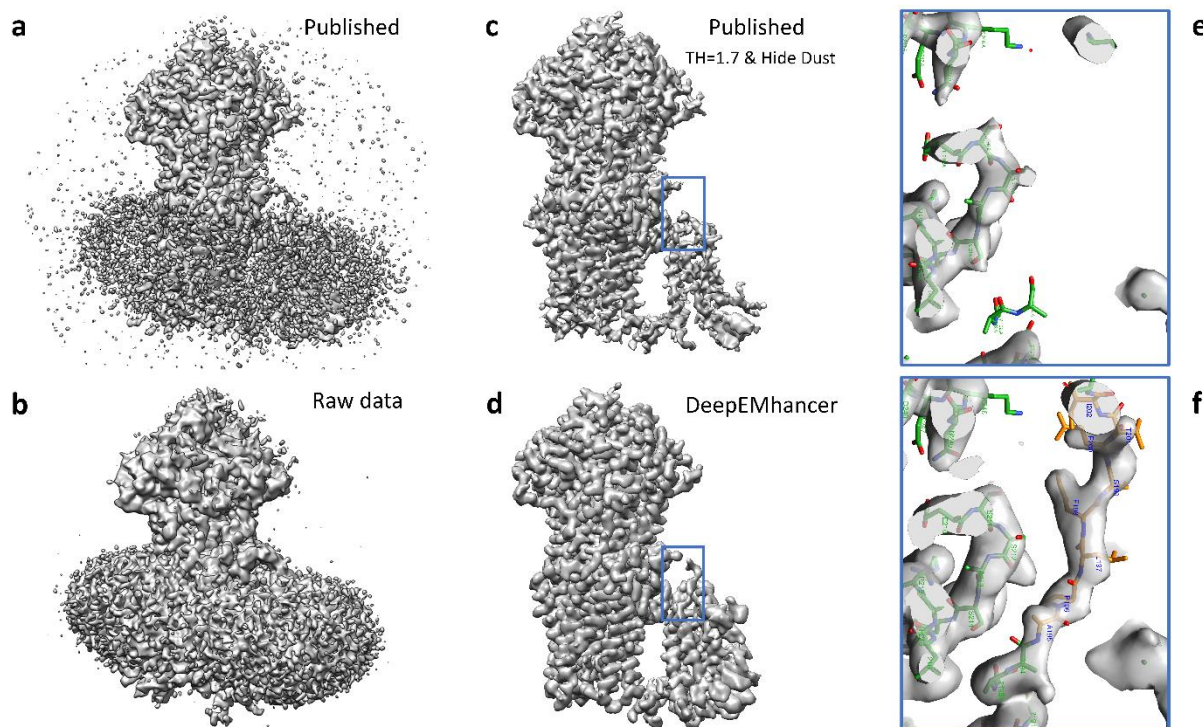
Figure 4. Testing map EMD-7099. **a**, Lateral view of the published map (B-factor sharpened, shown at recommended threshold). **b**, Lateral view of the raw data map obtained from the half maps that was used as input for DeepEMhacer. **c**, Lateral view of the published map after rising the threshold and removing the small connected components so that the signal coming from the lipids was removed. As a collateral consequence, some densities corresponding to the protein were also lost. **d**, Lateral view of the map obtained with DeepEMhacer. **e**, Zoom-in of the region marked with a blue box in c. **f**, Zoom-in of the region marked with a blue box in d, in which DeepEMhacer post-processed map, contrary to the published map, shows the densities corresponding to a missing loop in PDB 6bhu chain A. As a result, the residues A195 to I203 have been *de novo* modeled (new residues depicted in yellow, published in green).

(Tenthorey et al., 2017), that will be discussed in Supplementary Material, most of the evaluated maps exhibit a non-negligible improvement in resolution, especially notable when compared to B-factor-based results (see next section), with a median value of ~0.3 Å.

Alternatively, with the aim of obtaining a complementary measurement of improvement, we computed the DeepRes local resolution for the input and post-processed maps. As can be appreciated in Figure 2, all test cases treated with DeepEMhacer improved in terms of DeepRes local resolution, with dramatic improvements of more than 0.8 Å and a median improvement of ~0.4 Å. Again, those figures, consistent with the FSC-based measurements, point out that DeepEMhacer is improving the interpretability of the maps.

Comparison with B-factor-based methods.

With the aim of comparing DeepEMhacer with the commonly employed B-factor-based sharpening methods, we repeated the same experiments using the post-processed maps obtained with the Relion postprocessing algorithm (Kimanius et al., 2016; Zivanov et al., 2018). Before, it is important to notice that, contrary to DeepEMhacer, Relion
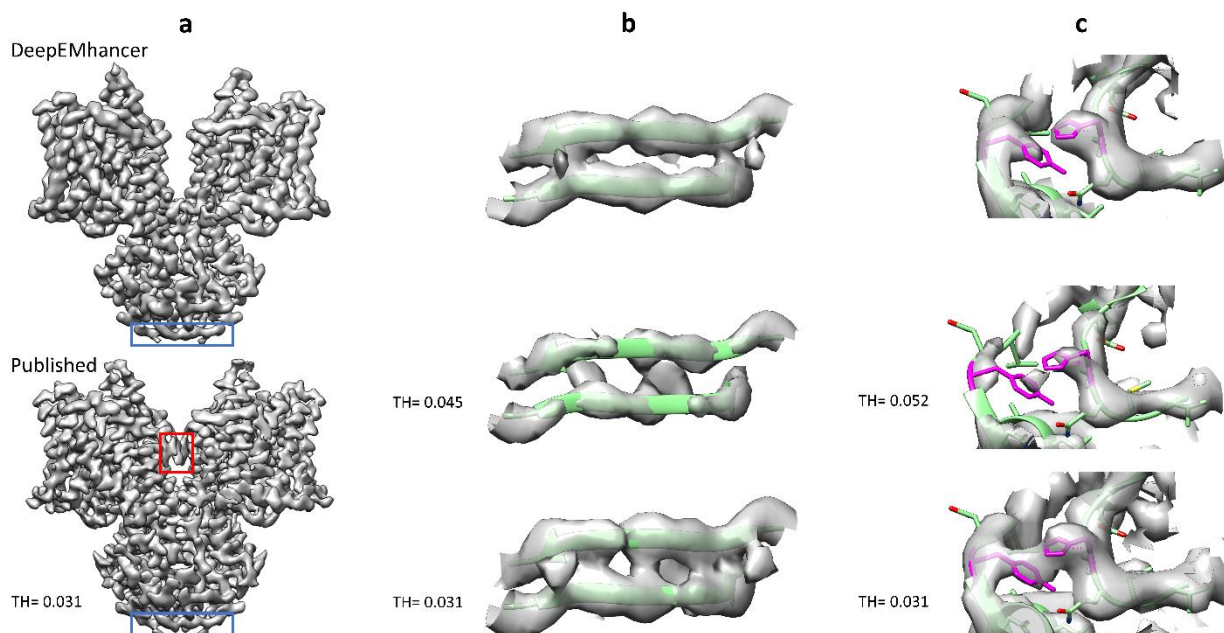
Figure 5. Testing map EMD-4997. **a**, Overview of the published map (B-factor sharpened, shown at the recommended threshold of 0.031), bottom, and the map obtained with DeepEMhacer, top. Red box highlights an artifact that has been automatically removed by DeepEMhacer. Blue box delimits the region showed in b. **b**, Zoom-in of the region marked with a blue box that contains the β-sheet R7-A10, chains A and B. The published volume is shown at the recommended threshold and at the threshold at which the backbone begins to look discontinuous. As it can be appreciated, DeepEMhacer solution resolves better than the published map the two strands of the sheet. **c**, Zoom-in of the region centered at chain B residues H121 and Y361 (colored in magenta). The published volume is shown at the recommended threshold and at the smaller threshold at which the density that connects the two residues disappears. As it can be appreciated, DeepEMhacer post-processed map resolves better than the published map the two residues

automatic masking is a simple process and thus, in order to make the comparison more interesting, we used instead the masks derived from the atomic models. Still, when we evaluated the FSC for the masked regions, only a few maps improved, while many others worsened, leading to a median improvement that was negligible (<0.05 Å) for both FSC and median DeepRes resolution (see Figure 2 and Figure 3).

We acknowledge that the automatic determination of the B-factor can lead to less accurate results than if it were manually selected and it may be the reason behind the poor observed performance. Thus, we have also included in the comparison the post-processed maps deposited in EMDB in which the estimation of B-factor was carried out by the authors. In this case, the improvement in resolution, with median values of ~0.15 Å and ~0.1 Å for DeepRes and FSC respectively, although closer to the values obtained using DeepEMhacer, are still considerably inferior (see Figure 2 and 3). In the light of these results, we

can state that DeepEMhacer maps tend to be much more similar to the atomic models than the ones obtained using B-factor-based methods and thus, more useful for the process of model building.

Visual inspection of testing maps

The purpose of this section is to further explore the results obtained with DeepEMhacer for some of the maps included in the testing set with the aim of illustrating how the improvements in global quality measurements translate to tangible improvements in the quality of the maps.

EMD-7099

The EMD-7099 (Johnson and Chen, 2018) is a high-resolution volume (global resolution 3.1 Å) of a multidrug resistance ATP-driven pump. EMD-7099 presents 17 transmembrane helices and, although the overall quality of the map is excellent, visualizing the transmembrane regions is challenging because of the signal that comes from the lipids. As a result, important parts of the

protein are not traced. Due to the fact that DeepEMhacer was trained to ignore the signal coming from lipidic layers, this example illustrates the unique characteristics of DeepEMhacer when applied to membrane proteins. Thus, as can be observed in Figure 4a-d, DeepEMhacer has been able to suppress the signal coming from the lipid layer in a much more simple and effective way than diminishing the threshold in the raw map or the B-factor-based sharpened maps. The noise suppression effect simplifies the

process of model building, as the researchers do not have to deal with masks or larger thresholds that make the visualization of near to noise level features more difficult. Yet, not only DeepEMhacer produces a noise reduction effect, but also it is able to enhance some parts of the map that under B-factor based sharpening seem noisy and disconnected. Such improvement, although observed in several regions of the map,
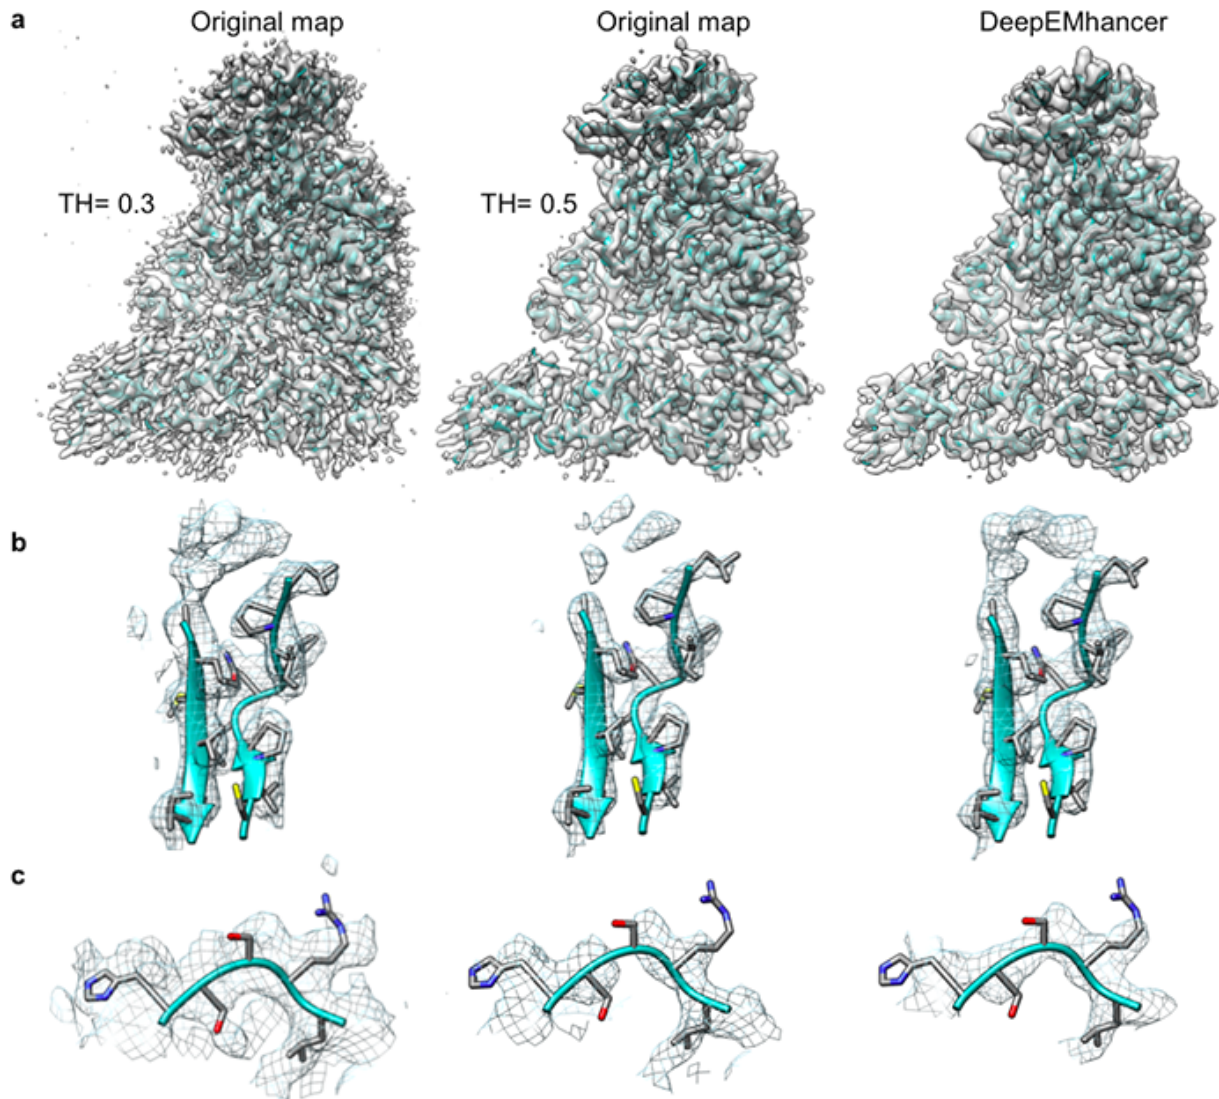
Figure 6. Use case EMD-30178 from SARS-CoV-2 RNA-dependent RNA polymerase. **a**, Overview of the original map displayed with two different thresholds 0.3 (recommended) and 0.5 (left and middle panel, respectively) and processed with DeepEMhacer software (right). 7btf pdb is shown in ribbon, black squares designated the zoomed areas in b panel and blue squares the zoomed areas in c. **b**, Zoom-in and extraction of the density in mesh from the 3D reconstruction of the original map at different thresholds and DeepEMhacer map corresponding to the black squares in a, chain D from residues C114 to I132. **c**, Zoom-in and extraction of the density in mesh from the 3D reconstruction of the original map at different thresholds and DeepEMhacer map corresponding to the blue squares in a, chain A residues H362 to L366. 7btf PDB in b and c is presented in ribbon and the residues in sticks.

is more noticeable at the transmembrane region Thus, the most important enhancement is depicted in Figure 4e-f, in which an important part of the backbone of the protein has been *de novo* traced thanks to DeepEMhacer enhancement, that has restored the densities corresponding to residues A195 to I203 in chain A of PDB 6bhu.

EMD-4997

The EMD-4997 (Walter et al., 2019) is a medium-high resolution volume (4.0 Å) for a murine epithelial anion transporter. As in the previous example, the overall quality of the map is quite good, yet it presents lower quality regions. Figure 5a shows an overview of the published map, displayed at the recommended threshold, and the map obtained with DeepEMhacer. Although it is true that both the original map and the post-processed map look very similar, it is also true that there exist important differences. Firstly, the map processed with DeepEMhacer is cleaner than the original one. Serve as an example the removal of the artifacts that the published map presents near the elbow of the complex (see Figure 5a, red box). More importantly, there can also be found many regions for which the DeepEMhacer post-processed volume resolves better the different residues of the regions. One of such examples can be found near the N-terminal end of the protein complex. Thus, as it is shown in Figure 5b, the densities that correspond to the strands of the beta sheet are better separated than in the published volume. It is important to notice that this better separation is not a consequence of the employed thresholds, as it is proven by the fact that rising the threshold makes the densities corresponding to the backbone discontinuous before the densities for the two strands separate (see Figure 5b). As a result, we can affirm that the quality of this region has been improved by the usage of DeepEMhacer.

Another similar example is displayed in Figure 5c. In this case, two non-contiguous aromatic residues, Y361 and H121, seem connected in the original map. However, when DeepEMhacer is applied, the densities corresponding to the two residues look separated while the backbone remains continuous.

Use case EMD-30178 from SARS-CoV-2 RNA-dependent RNA polymerase

In order to further explore the benefits of the DeepEMhacer algorithm we analyzed more deeply the post-processing of EMD-30178 map from Gao et al. (Gao et al., 2020), corresponding to the SARS-CoV-2 RNA-dependent RNA polymerase. The original map presents detailed structure up to 2.9 Å resolution, however as is often the case in cryo-EM, the resolution of the map is highly heterogeneous. We have chosen this map not only for the importance of this structure in current days but also because of the fact that the heterogeneous quality of the map density presents an ideal case for DeepEMhacer software. As it is shown in Figure 6a, the application of the algorithm reduces the noise and improves the consistency and depiction of the map. To better illustrate these differences, we have chosen two different regions in chains A and D where the differences between the original and the DeepEMHancer map can be appreciated (Figure. 6b-c). While the density in the original map looks noisy or discontinuous depending on the displayed threshold (Figure 6 b-c, left and middle panel), the application of the DeepEMhacer software results in a well-defined continuous density where the side chains are nicely depicted (Figure 6 b-c, right panel). This improvement in the map density allowed us to close the loop between residues C114 to I132 from chain D tracing 4 new residues that were not traced in the original structure (Figure 6b and Figure 7a). The improvement of the density is not only applicable to the edges of the map but it can be also appreciated in its core. Residues H362 to L366 in chain A, traced on the original map where positioned more accurately on the density after map post-processing (Figure 6c and Figure 7b).
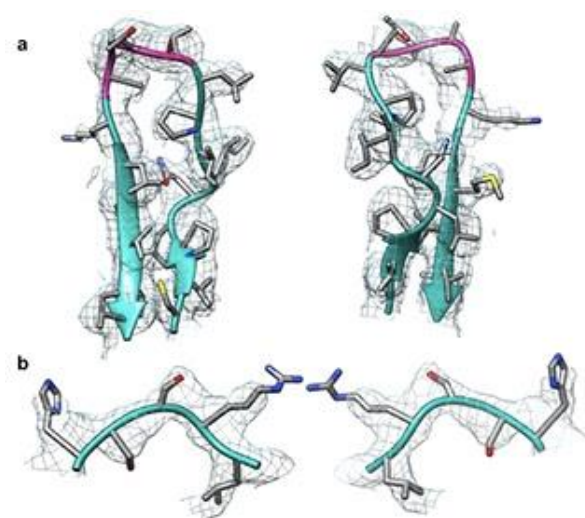
Figure 7. Examples of improvement of the tracing in SARS-CoV-2 RNA-dependent RNA polymerase

DeepEMhacer map. **a**, Zoom-in and extraction of the density in mesh from the 3D reconstruction in two different views of DeepEMhacer map showing 7btf in turquoise and new traced loop region in magenta (chain D residues T123, T124, A125 and A126, in sticks). **b**, Zoom-in and extraction of the density in mesh from the 3D reconstruction of DeepEMhacer map in two different views, 7btf in turquoise shows the repositioning of the chain A residues H362, S363, S364, R365 (in sticks).

Discussion

The number of deposited high resolution cryo-EM maps have soared since the beginning of the 'resolution revolution'. As a result, there is an increasing number of atomic models that are being built using cryo-EM as the primary source of information. However, building atomic models directly from the raw maps is generally not possible. Instead, maps are post-processed in order to enhance the contrast of their high-resolution features.

In this work we have presented Deep cryo-EM Map Enhancer (DeepEMhacer), a new map post-processing method based on deep learning. Trained on pairs of experimental cryo-EM maps and post-processed maps constructed with LocScale using atomic models, DeepEMhacer has learned how to perform a high-quality post-processing operation that reproduces the effects of masking and local sharpening in an automatic fashion.

The performance of our new algorithm has been assessed using a testing set of 20 experimental maps that were not used for training nor during the trial and error process required for its implementation. In all cases, the similarity between the maps obtained from the atomic models and the experimental maps improved after the application of DeepEMhacer. Additionally, we evaluated in detail the performance of DeepEMhacer on two of those maps, showing that, not only DeepEMhacer facilitates the visualization of cryo-EM maps, but also that DeepEMhacer can unveil some details that are not easily recognizable in the raw maps. Finally, we have employed DeepEMhacer on a map of the RNA polymerase of the SARS-CoV 2 virus, improving its quality of the map and the quality of the associated atomic model.

DeepEMhacer is available as an Xmipp plugin for Scipion v3 (currently branch rsg_deepPostProcessing).

Acknowledgments

Author Contributions

Conceptualization, J.V., C.O.S.S., R.S.; Methodology, R.S., J.V. and C.O.S.S.; Software implementation, R.S., J.G.; Evaluation, R.S., J.G., A.C.; Writing, R.S., A.C., J.V., C.O.S.S. and J.M.C.; Supervision, J.V., C.O.S.S.; Funding Acquisition, J.M.C, J.V.

References

Avramov, T., Vyenielo, D., Gomez-Blanco, J., Adinarayanan, S., Vargas, J., Si, D., 2019. Deep Learning for Validating and Estimating Resolution of Cryo-Electron Microscopy Density Maps †. Molecules 24, 1181. https://doi.org/10.3390/molecules24061181

Emsley, P., Cowtan, K., 2004. Coot: Model-building tools for molecular graphics. Acta Crystallogr. Sect. D Biol. Crystallogr. 60, 2126–2132. https://doi.org/10.1107/S0907444904019158

Gao, Y., Yan, L., Huang, Y., Liu, F., Zhao, Y., Cao, L., Wang, T., Sun, Q., Ming, Z., Zhang, L., Ge, J., Zheng, L., Zhang, Y., Wang, H., Zhu, Y., Zhu, C., Hu, T., Hua, T., Zhang, B., Yang, X., Li, J., Yang, H., Liu, Z., Xu, W., Guddat, L.W., Wang, Q., Lou, Z., Rao, Z., 2020. Structure of the RNA-dependent RNA polymerase from COVID-19 virus. Science (80-.). 368, 779–782. https://doi.org/10.1126/science.abb7498

Gupta, H., McCann, M.T., Donati, L., Unser, M., 2020. CryoGAN: A New Reconstruction Paradigm for Single-particle Cryo-EM Via Deep Adversarial Learning. bioRxiv 2020.03.20.001016. https://doi.org/10.1101/2020.03.20.001016

He, K., Zhang, X., Ren, S., Sun, J., 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification, in: Proceedings of the IEEE International Conference on Computer Vision. pp. 1026–1034. https://doi.org/10.1109/ICCV.2015.123

Jakobi, A.J., Wilmanns, M., Sachse, C., 2017. Model-based local density sharpening of cryo-EM maps. Elife 6. https://doi.org/10.7554/eLife.27131

Johnson, Z.L., Chen, J., 2018. ATP Binding Enables Substrate Release from Multidrug Resistance Protein 1. Cell 172, 81-89.e10. https://doi.org/10.1016/j.cell.2017.12.005

Kaur, S., Gomez-Blanco, J., Khalifa, A., Adinarayanan, S., Sanchez-Garcia, R., Wrapp, D., McLellan, J.S., Bui, K.H., Vargas, J., 2020. Local computational methods to improve the interpretability and analysis of cryo-EM maps. bioRxiv 2020.05.11.088013. https://doi.org/10.1101/2020.05.11.088013

Kimanius, D., Forsberg, B.O., Scheres, S.H., Lindahl, E., 2016. Accelerated cryo-EM structure determination with parallelisation using GPUs in RELION-2. Elife 5, e18722. https://doi.org/10.7554/eLife.18722

Lawson, C.L., Patwardhan, A., Baker, M.L., Hryc, C., Garcia, E.S., Hudson, B.P., Lagerstedt, I., Ludtke, S.J., Pintilie, G., Sala, R., Westbrook, J.D., Berman, H.M., Kleywegt, G.J., Chiu, W., 2015. EMDataBank unified data resource for 3DEM. Nucleic Acids Res. 44, D396–D403. https://doi.org/10.1093/nar/gkv1126

Maddhuri Venkata Subramaniya, S.R., Terashi, G., Kihara, D., 2019. Protein secondary structure detection in intermediate-resolution cryo-EM maps using deep learning. Nat. Methods 16, 911–917. https://doi.org/10.1038/s41592-019-0500-1

Pettersen, E.F., Goddard, T.D., Huang, C.C., Couch, G.S., Greenblatt, D.M., Meng, E.C., Ferrin, T.E., 2004. UCSF Chimera - A visualization system for exploratory research and analysis. J. Comput. Chem. 25, 1605–1612. https://doi.org/10.1002/jcc.20084

Ramírez-Aportela, E., Mota, J., Conesa, P., Carazo, J.M., Sorzano, C.O.S., 2019. DeepRes: A new deep-learning- and aspect-based local resolution method for electron-microscopy maps. IUCrJ 6, 1054–1063. https://doi.org/10.1107/S2052252519011692

Ramírez-Aportela, E., Vilas, J.L., Glukhova, A., Melero, R., Conesa, P., Martínez, M., Maluenda, D., Mota, J., Jiménez, A., Vargas, J., Marabini, R., Sexton, P.M., Carazo, J.M., Sorzano, C.O.S., 2020. Automatic local resolution-based sharpening of cryo-EM maps. Bioinformatics 36, 765–772. https://doi.org/10.1093/bioinformatics/btz671

Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation, in: Medical Image Computing and Computer-Assisted Intervention-MICCAI. pp. 234–241. https://doi.org/10.1007/978-3-319-24574-4_28

Rosenthal, P.B., Henderson, R., 2003. Optimal determination of particle orientation, absolute hand, and contrast loss in single-particle electron cryomicroscopy. J. Mol. Biol. 333, 721–745. https://doi.org/10.1016/j.jmb.2003.07.013

Si, D., Moritz, S.A., Pfab, J., Hou, J., Cao, R., Wang, L., Wu, T., Cheng, J., 2020. Deep Learning to Predict Protein Backbone Structure from High-Resolution Cryo-EM Density Maps. Sci. Rep. 10, 1–22. https://doi.org/10.1038/s41598-020-60598-y

Tenthorey, J.L., Haloupek, N., López-Blanco, J.R., Grob, P., Adamson, E., Hartenian, E., Lind, N.A., Bourgeois, N.M., Chacón, P., Nogales, E., Vance, R.E., 2017. The structural basis of flagellin detection by NAIP5: A strategy to limit pathogen immune evasion. Science (80-.). 358, 888–893. https://doi.org/10.1126/science.aao1140

Terwilliger, T.C., Sobolev, O. V., Afonine, P. V., Adams, P.D., 2018. Automated map

sharpening by maximization of detail and connectivity. Acta Crystallogr. Sect. D Struct. Biol. 74, 545–559. https://doi.org/10.1107/S2059798318004655

Vilas, J.L., Vargas, J., Martinez, M., Ramirez-Aportela, E., Melero, R., Jimenez-Moreno, A., Garduño, E., Conesa, P., Marabini, R., Maluenda, D., Carazo, J.M., Sorzano, C.O.S., 2020. Re-examining the spectra of macromolecules. Current practice of spectral quasi B-factor flattening. J. Struct. Biol. 209, 107447. https://doi.org/10.1016/j.jsb.2020.107447

Wagner, T., Merino, F., Stabrin, M., Moriya, T., Antoni, C., Apelbaum, A., Hagel, P., Sitsel, O., Raisch, T., Prumbaum, D., Quentin, D., Roderer, D., Tacke, S., Siebolds, B., Schubert, E., Shaikh, T.R., Lill, P., Gatsogiannis, C., Raunser, S., 2019. SPHIRE-crYOLO is a fast and accurate fully automated particle picker for cryo-EM. Commun. Biol. 2. https://doi.org/10.1038/s42003-019-0437-z

Walter, J.D., Sawicka, M., Dutzler, R., 2019. Cryo-EM structures and functional characterization of murine Slc26a9 reveal mechanism of uncoupled chloride transport. Elife 8. https://doi.org/10.7554/eLife.46986

Wang, F., Gong, H., Liu, G., Li, M., Yan, C., Xia, T., Li, X., Zeng, J., 2016. DeepPicker: A deep learning approach for fully automated particle picking in cryo-EM. J. Struct. Biol. 195, 325–336. https://doi.org/10.1016/j.jsb.2016.07.006

Wu, Y., He, K., 2020. Group Normalization. Int. J. Comput. Vis. 128, 742–755. https://doi.org/10.1007/s11263-019-01198-w

Yang, W., Zhang, X., Tian, Y., Wang, W., Xue, J.H., Liao, Q., 2019. Deep Learning for Single Image Super-Resolution: A Brief Review. IEEE Trans. Multimed. https://doi.org/10.1109/TMM.2019.2919431

Zhong, E.D., Bepler, T., Davis, J.H., Berger, B., 2019. Reconstructing continuous distributions of 3D protein structure from cryo-EM images.

Zhu, Y., Ouyang, Q., Mao, Y., 2017. A deep convolutional neural network approach to single-particle recognition in cryo-electron microscopy. BMC Bioinformatics 18, 348. https://doi.org/10.1186/s12859-017-1757-y

Zivanov, J., Nakane, T., Forsberg, B.O., Kimanius, D., Hagen, W.J.H., Lindahl, E., Scheres, S.H.W., 2018. New tools for automated high-resolution cryo-EM structure determination in RELION-3. Elife 7. https://doi.org/10.7554/eLife.42166