Learning Super-Resolution Electron Density Map of Proteins using 3D U-Net

Baishali Mullick[†], **Yuyang Wang**[‡], **Prakarsh Yadav**[‡], **Amir Barati Farimani**[‡§*]
[†]Department of Electrical and Computer Engineering
[‡]Department of Mechanical Engineering
[§]Machine Learning Department
Carnegie Mellon University, Pittsburgh, PA 15213, USA

Abstract

A well-established protein structure is essential for understanding protein molecular mechanism, phenotypic implication and drug discovery. Recent development of cryo-Electron Microscopy (cryo-EM) offers the advantage of easy sample preparation and not requiring crystallized protein for structural biology. However, the resolution of cryo-EM electron density maps used to determine protein structure, is not at par with X-ray diffraction (XRD) or NMR. In this work, we propose to leverage a deep learning-based model to increase the resolution of low-quality electron density maps. The model is built upon U-Net with 3D convolutional layers, which contains three components: encoder, bottleneck, and decoder. To get paired maps of different resolutions, we collect high-resolution maps from XRD as ground truth labels. While the low-resolution maps are obtained through a noise model which combines dilation operations, Gaussian filters and Gaussian noise. We also introduce data augmentation techniques during model training, like random cropping, rotation, and flipping. Experiments show that when applied to low-resolution electron maps, the U-Net model can improve the resolution in the metric of EMRinger score, which redesigns the map so that it resolves the regions of ambiguity to offer greater certainty in the position of amino acids.

1 Introduction

The elucidation of protein structures is crucial for understanding the biological processes. Proteins carry out numerous biological functions such as catalysis, cell signalling, structural support and immunological response [1, 2]. To better understand these functions it is important to know what properties of the protein enable it to carry out its function. To gain this insight, researchers rely on the structural information of the proteins. The structure of protein is used to explain its mechanism of action, influence of specific mutations on structure of the protein and its associated phenotype [3, 4]. Additionally, the knowledge of the structure of the proteins is seminal for many biotechnology applications like protein engineering, gene therapies and antibody design [5].

The experimental process of determining protein structure is time-consuming and requires isolation of large amounts of soluble protein. The conventionally used methods of protein structure elucidation are X-Ray diffraction (XRD) and Nuclear Magnetic Resonance (NMR) [6]. Both approaches offer good level of structure resolution of small size of proteins, **Fig. 1a**, where the accurate conformation of amino acid side-chains are clearly distinguishable. However, both approaches suffer from their drawbacks. NMR is not able to accurately resolve the structure of proteins larger than 35 kDa [7], whereas XRD requires crystallization of protein molecules. In addition, both the approaches are

*Corresponding author: barati@cmu.edu.

Machine Learning for Structural Biology Workshop, NeurIPS 2020, Vancouver, Canada.

incapable of resolving structures of macro-molecular protein assemblies [8, 9]. Since, majority of the biological function occur due to interaction of multiple proteins, it is important to be able to accurately capture these macro-molecular interactions accurately.

The recent development in Cryo-Electron Microscopy has (Cryo-EM) enabled the researchers to gain insights into the structure of such macro-molecular assemblies [10]. Cryo-EM has greatly expanded the structural biologist's toolkit and its adoption by the scientific community has generated a large dataset of Cryo-EM structures and experimental data. However, Cryo-EM suffers from limitation of the protein structure resolution that it can generate. [11]. This limitation translates to the Cryo-EM protein structures to having very limited information on amino acid side chain orientation, **Fig. 1c**. The orientation of side chains is a crucial information for accurate description of the molecular basis of enzymatic action and protein-protein interactions [12]. Attempts have been made to circumvent this limitation by processing the Cryo-EM micrographs using sharpening algorithms but only an incremental improvement in the resolution has been observed.

The current methods for improvement of protein structure resolution also fail to leverage the plethora of data that is available on online servers, like RCSB-PDB and EMDB (EMDB; http://emdb-empiar.org) [13]. Deep Learning approaches which make use of such data can offer an alternative approach to improve upon the Cryo-EM based structure resolution. By training on thousands of high-resolution protein structures (with resolution of 1.2 Å to 2 Å) it is possible to featurize the protein structure data and by using Convolution Neural Networks (CNNs) it is possible to learn the important features associated with the high-resolution structures [14, 15, 16]. The CNNs can be trained on the feature space of the electron density maps, where they learn the electron density patterns associated with high resolution structures. The advantage of such an approach is that it is possible to train on a large number of high resolution structures from XRD. This allows learning the important features of the high-resolution XRD protein structures and leverage this learned information to improve the resolution of Cryo-EM structures. Since high-resolution XRD protein structures contain information the amino acid side chain orientation, it is possible to leverage this information to resolve the ambiguities in side chain orientation of Cryo-EM structures [17]. In this paper, we introduce 3D U-Net. A deep learning model through which we have aimed to improve the resolution of Cryo-EM structures by training on the electron density maps of protein structures. In addition, we introduce a novel approach for defining the feature space of the training dataset which makes use of both XRD and Cryo-EM experimental data. Through an adequately trained 3D U-Net, we were able to improve the resolution of simulated Cryo EM maps generated from our noise model and we were able to provide evidence that our model was consequently able to improve the quality of CryoEM maps.

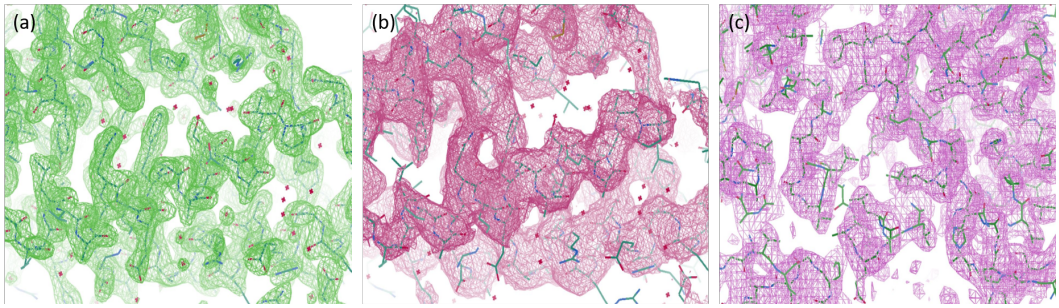

Figure 1: Visualization of randomly selected electron density maps used in this study, by using Coot. The mesh grid displays the electron density surface for the structure and the underlying line diagram represents the protein structure from the PDB file. **(a)** The electron density map of **PDB ID: 6AR7**(Green), with resolution of 2.10 Å. **(b)** The electron density map of **PDB ID: 6AR7** after applying the noise transform to generate model input (Red). **(c)** The electron density map of cryo-EM determined protein **PDB ID: 5ayy**(purple), with resolution of 3.09 Å.

2 Related Works

2.1 Electron Density Map Sharpening and Super-Resolution

Over the years, various techniques have been devised in order to sharpen the quality of Cryo Electron Microscopy (EM) maps. The most extensively used approach is the B-factor correction [18]. In this method, noise-dependent structure factor weight is established along with the temperature factor to prevent noise amplification thereby reducing contrast loss at high resolution. Sharpening methodologies can be classified as local and global [19] . Global sharpening algorithms determine a B factor value over the entire EM map to perform amplitude correction at high resolution. Examples of this technique are RELION post-processing [20] and AutoSharpen in the Phenix Package [21]. RELION [20] has introduced a semi automated particle picking and sorting algorithm which is an upgradation of the original method [22] Contrary to the above global approach, it has been observed that many of the recent EM maps have areas that vary tremendously in their resolution in comparison to the entire map due to sub units in their structures such as in the case of membrane proteins or macromolecules with high flexibility [23]. The LocScale [24] algorithm tackles this issue with its local sharpening approach. LocScale locally scales amplitudes iteratively to match the radial average amplitude of a corresponding atomic model in the Fourier domain. This procedure takes into account the difference in resolution over the entire EM map. The LocalDeblur [19] is another local sharpening algorithm for Cryo EM maps. It uses a gradient descent based iterative approach with Wiener filter. The map transformation takes place with the help of filter banks using bandpass raised cosine filters weighted by the difference of spatial frequency these filter centers and the respective local resolution. Both LocScale and LocalDeblur have proved to be promising techniques towards EM map sharpening and map quality enhancement. However, LocScale, requires a refined atomic model for comparison and local scaling calculation. In this sense this approach is biased and restrictive to the availability of such a model. Furthermore, LocDeblur presently its shortcomings through various theoretical assumptions and limitations. The DeepSharpen [17] provides a data-driven approach to enhance the quality of CryoEM maps. The dataset is obtained from the atomic structure models available in the Protein Data Bank. A pair of high and low resolution density map is generated from the model using a gaussian filtering and same is used to train a reverse model, with 3D convolutional neural networks, to convert the low resolution map back to the higher resolution counterpart. This data-driven approach shows some promising results. However it is heavily dependent on an initial atomic structure model and thus has a bias. Furthermore, same is restrictive to the availability of such a model. The high resolution and low resolution pair of maps are generated with a gaussian blur filter which results in lower quality maps compared to the original EM map and hence the model training is done on lower resolution values. This results in curbing the performance of the overall model in certain scenarios.

2.2 Deep Learning for 3D voxels

To expand deep learning techniques from 2D images to 3D, representations are required to describe the 3D models. One of the widely used represenation is voxel [25], and each voxel represents a value on a regular grid in 3D space. In VoxNet [26], 3D convolutional neural network (3D CNN) is implemented on the voxel data to solve the object recognition. Further, [27] extended 2D CNNs to 3D and systematically analyzed the relationship between 3D CNNs and multi-view CNNs. To decrease the memory and computational cost of 3D CNN, O-CNN [28] makes use of octree representations and perform 3D CNN operations only on the sparse octants. Also, generative adversarial networks (GANs) have been leveraged in synthesize realistic 3D voxel models. In 3D-GAN [29], 3D voxel objects are synthesized from a probabilistic space in combination of 3D volumetric convolution and GAN. Similarly, octree can be utilized to improve the efficiency of voxel model generation [30]. Besides, V-Net [31] and 3D U-Net [32] introduce fully connected volumetric convolutional layers for 3D voxel segmentation.

3 Method

3.1 Data Pre-processing

The DeepSharpen has used a 3D Gaussian filter of varied sigma values to represent the EM map at a lower map quality level [17]. LocDeblur has assumed a Gaussian approximation for the noise

distribution used to define the Wiener Filter. This Wiener filter is then used as the transformation matrix between the observed EM map and the sharpened counterpart of the EM map [19]. Along the same lines, we have introduced a Gaussian filtering and a low-level Gaussian noise for defining the noise distribution in order to generate lower resolution maps from XRay density maps which will mimic the quality of Cryo EM maps, **Fig. 1b**. XRay maps have a higher resolution indicated by sharper electron density mesh clouds compared to Cryo EM maps that have a distended electron density mesh grid indicative of lower confidence in certain regions of the Cryo EM map, **Fig. 1**. In order to make a more accurate representation of the noise model we have used dilation morphological transform with a ball structuring element in order to inflate the electron mesh clouds of the XRay maps such that they resemble the mesh clouds of Cryo EM maps. Our noise model works well to simulate Cryo EM maps and same is corroborated by the EMRinger scores of the corresponding low resolution input maps, **Table 1**.

To train the model based on local resolutions across the map, we have taken random patches of size 128 **x** 128 **x** 128 in order to enable local sharpening of the map. Random flip as well as random rotation of the image patch along all three axis help augment the dataset, resulting in large number of patches from each training data sample which reduces the risk of overfitting.

3.2 3D U-Net

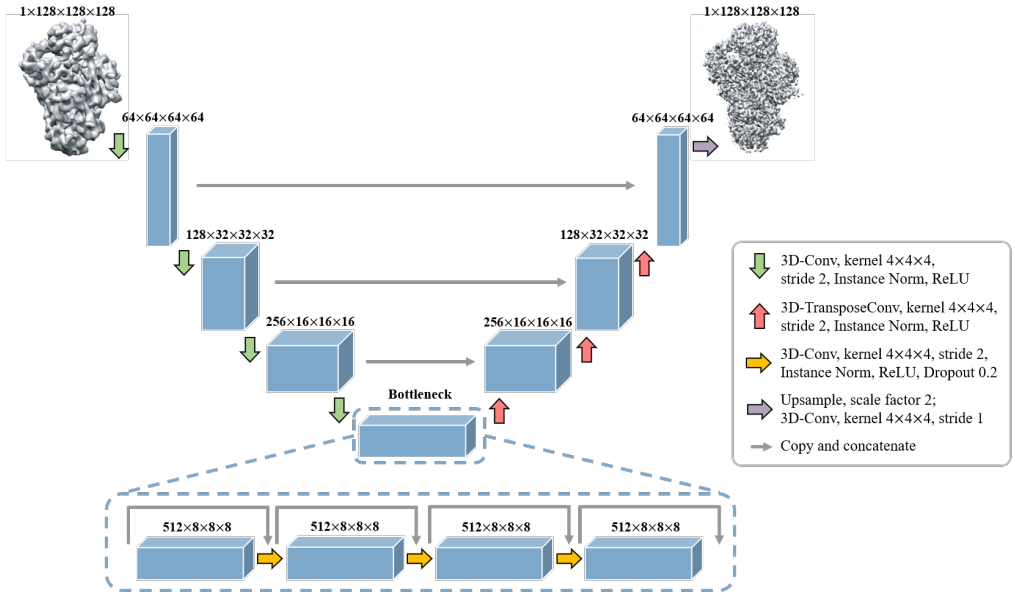

Figure 2: Overview of the deep learning model for electron density map super-resolution built upon 3D U-Net, which consists of the encoder, bottleneck and decoder.

We build our electron density map super-resolution model based upon 3D U-Net [33]. The model consists of three components: downsampling, bottleneck, and upsampling. For each layer in downsampling, 3D volumetric convolution with kernel size $4 \times 4 \times 4$ and stride size 2 is implemented. The bottleneck has the same kernel size for the 3D volumetric convolution, however, the stride size is set as 1. Also, dropout layers with probability 0.2 is utilized after each 3D convolutions in the bottleneck. For upsampling, each layer uses 3D transpose convolution operation with kernel size $4 \times 4 \times 4$ and stride 2. ReLU [34] is used as activation function in all the hidden layers in the downsampling, bottleneck and upsampling. Also, Instance normalization [35] is implemented after each convolution operation as shown in **Fig. 2**. The output layer of the model first upsamples the 3D feature map by scale factor 3 and then uses a 3D volumetric convolution with kernel size $4 \times 4 \times 4$ and stride size 1.

During training, we set the batch size to be 4 and total epochs to be 100. The model is trained using the Adam optimizer [36] with the learning rate 2×10^{-4}, $\beta_1 = 0.5$, and $\beta_2 = 0.999$. The loss

function to minimize during training is built on dice coefficient [31]:

$$D = \frac{2\sum_i p_i g_i}{\sum_i p_i^2 + \sum_i g_i^2},$$

where p_i is the prediction of each voxel value and g_i is the ground truth.

4 Experiments

A challenge with training deep learning models, which sharpen the maps, on Cryo EM electron density maps is that there are no high resolution structures which can be used as ground truth during the training of the model. To circumvent this and simultaneously increase the training dataset size, we applied the afore-mentioned data preprocessing method to more than 1000 XRD electron density maps. This generated the Cryo EM resembling training dataset for the training of 3D U-Net, the original high-resolution, XRD maps were considered to be the ground truth for training of the model. This dataset was sufficiently diverse to enable the model to learn the important features associated with the high-resolution XRD electron density maps. We then applied the fully trained model to improve the resolution of low quality XRD maps obtained from our noise model as well as on Cryo EM maps. We then assessed the model's performance on this test data.

4.1 Evaluation of XRay Diffraction Maps

To train the 3D U-Net to extract and learn the important features associated with XRD maps, we used the training dataset in batches of size 4 and optimized the model using the Adam. The model was fully trained in 200 epochs and had a validation MSE loss of 0.028. After successfully training the model we evaluated its performance on a dataset of about 200 XRD electron density maps. The applicability of deep learning models in structural biology is dependent on their interpretability and the intuitive understand of the model's output. To measure the quality of the sharpened maps, we used the EMRinger score and through the model we were able to improve the resolution of noise simulated XRD maps. The model successfully improved the EMRinger score by 2-fold on average across the testing dataset.

Table 1: EMRinger Score for XRD Maps

PDB-ID	Input	Our Prediction	Ground Truth
2BV8	2.439	6.781	8.057
6AR7	2.439	5.620	7.651
1A5S	2.036	4.479	4.234

The training data consisted of multiple types of proteins ranging from multi-mers, protein-ligand complexes, protein-DNA complexes, etc. Such diversity in training data ensured that the features learnt by the model would not be constrained only to the amino acids in proteins, instead they can be translated to electron density contaning a wide variety of biomolecules. The novel data preprocessing approach we have described also increases the size of the training dataset as we are able to mimic the features of Cryo EM electron density maps in case of XRD maps as well. In **Table 1**, we present EMRinger scores for 3 randomly selected XRD maps from the training dataset. EMRinger score is defined in [37] and it is computed via the Phenix package [38]. "EMRinger score measures the match between an atomic reference model and the density map. It detects and validates based on the presence of rotameric peaks in the side-chains. " [17]. In case of EMRinger a higher numerical value indicates a structure with better atomic resolution. From the table it can be seen that our data-preprocessing method is able to successfully lower the quality of XRD maps. Furthermore, after evaluating the model performance on XRD maps we observed that the model was able to improve the atomic resolution of the electron density maps by more than 2-folds. We investigated the effect of our data-preprocessing on the distribution of electron density in the map and the effect of 3D U-Net on this distribution. To visualize these effects we looked at a randomly selected slice of the **PDB ID: 2BV8** electron density map, **Fig. 3**. We observed that the data-preprocessing was able to successfully introduce a blurring effect in the map by extending the boundary of protein surface (green), **Fig. 3a**. The model was able to successfully sharpen the electron density map of the protein by reducing the

spread of surface and constricting electron density values into regions which overlapped with the ground truth, **Fig. 3b & 3c**. These results evidenced that the 3D U-Net is able to learn the important features associated with high resolution structures and utilizes these features effectively to sharpen the XRD maps. These results motivated the application of fully trained 3D U-Net to testing dataset of Cryo EM maps.

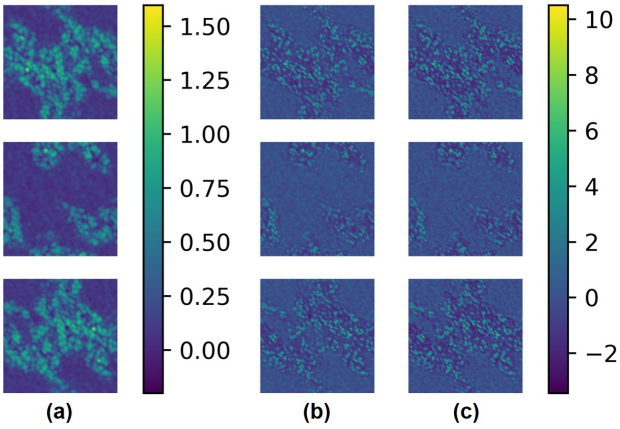

Figure 3: Electron density map slices of 2BV8: **(a)** Low-resolution input to the 3D U-Net model. **(b)** Predicted high-resolution electron density maps. **(c)** Ground truth electron maps from X-ray.

An advantage of 3D U-Net, is that it is trained to improve the electron density distribution of the map within the electron density feature space [17], in their work presented a model which sharpens the Cryo EM maps, however their model suffers from bias, in which it requires the atomic model protein structure. Our model does not require such an input, instead it operates strictly in the electron density map feature space. This can greatly extend the application of 3D U-Net, as it can be applied to elucidate the structure of novel proteins.

4.2 Evaluation of Cryo-EM Maps

We evaluated the performance of our model on several Cryo EM electron density maps. To measure the quality of the sharpened maps, we used the EMRinger score and with the model we were able to improve the resolution on some of these maps, although, we encountered some failed cases for CryoEM maps. To achieve the same we analyzed the output for 5 randomly selected Cryo-EM maps from the test set, **Table 2**. The **EMDB-20358** structure showed striking improvement of resolution as was the case with 20365 and 20359. However, for 20360 the score was nearly the same and in case of 20357 the score reduced. The inability of the model to improve the map in these cases can be attributed to previously unseen features present in these maps. The model would not be able to properly featurize these maps and as a result it was not able to perform very well. It is possible to improve the model's generalizablity by training on a much larger dataset which would enable the model to learn to featurize even more aspects of the electron density maps.

Table 2: EMRinger Score for Cryo-EM Maps

EMDB-ID	Ground Truth	Our Prediction
20357	7.048	5.817
20358	4.378	9.198
20359	6.587	7.966
20360	7.181	7.125
20365	7.362	9.198

Similar to evaluation of XRD maps, we also investigated distribution of the electron density values in the map of **EMDB-20358** and the effect of 3D U-Net on distribution of electron density values, **Fig. 4**. Here we observed that the electron density values were distributed over a larger area (yellow),

Fig. 4a. This dispersed distribution is indicative of a structure with lower resolution and matches the distribution observed in case of noise added XRD map used for training of model, **Fig. 3a**. This visual representation allowed us to ensure that the data preprocessing pipeline we have developed is capable of accurately mimicking the features of Cryo EM maps in XRD electron density distributions. This Cryo EM map was then passed through the 3D U-Net, and the model was able to successfully reduce the dispersion of the electron density values and limit it to specific regions. Here again we observed that the electron density distribution in case of model output, **Fig. 4a**, was closely resembling the distribution observed in case of XRD, **Fig. 3c**. These observations allowed us to conclude that the model is able to take in a Cryo EM electron density map and sharpen the distribution of electron densities by making use of the features it learnt from high resolution XRD maps.

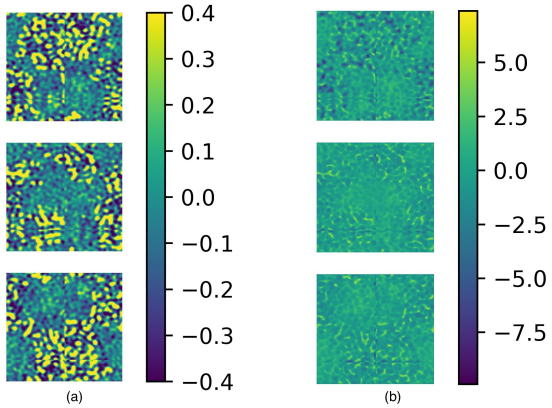

Figure 4: Electron density map slices EMDB-20358: **(a)** Low-resolution input to the 3D U-Net model (**4.378**). **(b)** Predicted high-resolution electron density maps (**9.198**).

5 Conclusion

In this work we present, a 3D U-Net inspired deep learning model, which is capable of sharpening the cryo-EM electron density maps. We also propose a novel data-preprocessing pipeline to augment the XRD electron density maps to mimic the cryo-EM maps, thereby greatly extending size of the training dataset. By training on this extended dataset we were able to obtain a model that was able to successfully sharpen the electron density maps in case of both XRD and cryo-EM maps. Future work to improve this model would be aimed at enhancing the generalizability of the model to a more diverse set of cryo-EM maps. This model is not limited by the requirement of the atomic model of proteins and offers advantage of being applicable to novel proteins, which existing deep learning based approaches fail to do. To evaluate the model performance we have used both quantitative (EMRinger) and qualitative (visualization of map slice) approaches. These approaches have enabled us to present evidence that our model is able to sharpen the cryo-EM maps.

References

[1] A Keith Dunker, Celeste J Brown, J David Lawson, Lilia M Iakoucheva, and Zoran Obradović. Intrinsic disorder and protein function. *Biochemistry*, 41(21):6573–6582, 2002.

[2] Martin Karplus and John Kuriyan. Molecular dynamics and protein function. *Proceedings of the National Academy of Sciences*, 102(19):6679–6685, 2005.

[3] Theresa M Cabrera-Vera, Jurgen Vanhauwe, Tarita O Thomas, Martina Medkova, Anita Preininger, Maria R Mazzoni, and Heidi E Hamm. Insights into g protein structure, function, and regulation. *Endocrine reviews*, 24(6):765–781, 2003.

[4] Peter E Wright and H Jane Dyson. Intrinsically unstructured proteins: re-assessing the protein structure-function paradigm. *Journal of molecular biology*, 293(2):321–331, 1999.

[5] David Baker and Andrej Sali. Protein structure prediction and structural genomics. *Science*, 294(5540):93–96, 2001.

[6] Mark R Wormald, Andrei J Petrescu, Ya-Lan Pao, Ann Glithero, Tim Elliott, and Raymond A Dwek. Conformational studies of oligosaccharides and glycopeptides: complementarity of nmr, x-ray crystallography, and molecular modelling. *Chemical Reviews*, 102(2):371–386, 2002.

[7] Hongtao Yu. Extending the size limit of protein nuclear magnetic resonance. *Proceedings of the National Academy of Sciences*, 96(2):332–334, 1999.

[8] Gregory A Petsko and Dagmar Ringe. Fluctuations in protein structure from x-ray diffraction. *Annual review of biophysics and bioengineering*, 13(1):331–371, 1984.

[9] K Ravi Acharya and Matthew D Lloyd. The advantages and limitations of protein crystal structures. *Trends in pharmacological sciences*, 26(1):10–14, 2005.

[10] Maya Topf, Keren Lasker, Ben Webb, Haim Wolfson, Wah Chiu, and Andrej Sali. Protein structure fitting and refinement guided by cryo-em density. *Structure*, 16(2):295–307, 2008.

[11] Jean-Paul Renaud, Ashwin Chari, Claudio Ciferri, Wen-ti Liu, Herve-William Remigy, Holger Stark, and Christian Wiesmann. Cryo-em in drug discovery: achievements, limitations and prospects. *Nature reviews Drug discovery*, 17(7):471–492, 2018.

[12] Xing Zhang and Z Hong Zhou. Limiting factors in atomic resolution cryo electron microscopy: no simple tricks. *Journal of structural biology*, 175(3):253–263, 2011.

[13] Helen M Berman, John Westbrook, Zukang Feng, Gary Gilliland, Talapady N Bhat, Helge Weissig, Ilya N Shindyalov, and Philip E Bourne. The protein data bank. *Nucleic acids research*, 28(1):235–242, 2000.

[14] Chanqin Quan, Lei Hua, Xiao Sun, and Wenjun Bai. Multichannel convolutional neural network for biological relation extraction. *BioMed research international*, 2016, 2016.

[15] Sarah Webb. Deep learning for biology. *Nature*, 554(7693), 2018.

[16] Christof Angermueller, Tanel Pärnamaa, Leopold Parts, and Oliver Stegle. Deep learning for computational biology. *Molecular systems biology*, 12(7):878, 2016.

[17] Zhizhen Zhao Mona Zehni, Minh N. Do. Deepsharpen: Deep-learning based sharpening of 3d reconstruction map from cryo-electron microscopy. *IEEE*, 2020.

[18] Richard Henderson Peter B. Rosenthal. Optimal determination of particle orientation, absolute hand, and contrast loss in single-particle electron cryo-microscopy. *J. Mol. Biol., 333, 721–745*, 2003.

[19] E. Ramırez-Aportela et al. Automatic local resolution-based sharpening of cryo-em maps. *Bioinformatics, 36(3), 2020, 765–772*, 2019.

[20] Sjors H.W. Scheres. Semi-automated selection of cryo-em particles in relion-1.3. *J. Mol. Biol., 333, 721–745*, 2015.

[21] T.C. et al Terwilliger. Automated map sharpening by maximization of detail and connectivity. *Acta Crystallographica Section D, vol. 74, no. 6, pp. 545–559*, 2018.

[22] (Zhu et al. Automatic particle selection: results of a comparative study. *J. Struct. Biol 145(1-2), 3-14*, 2004.

[23] McMullan G Scheres SH Bai XC, Fernandez IS. Ribosome structures to near-atomic resolution from thirty thousand cryo-em particles. *eLife 2:e00461*, 2013.

[24] Carsten Sachse Arjen J Jakobi, Matthias Wilmanns. Model-based local density sharpening of cryo-em maps. *eLife, vol. 6, pp. e27131*, 2017.

[25] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015.

[26] Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 922–928. IEEE, 2015.

[27] Charles R Qi, Hao Su, Matthias Nießner, Angela Dai, Mengyuan Yan, and Leonidas J Guibas. Volumetric and multi-view cnns for object classification on 3d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5648–5656, 2016.

[28] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-cnn: Octree-based convolutional neural networks for 3d shape analysis. *ACM Transactions on Graphics (TOG)*, 36(4):1–11, 2017.

[29] Jiajun Wu, Chengkai Zhang, Tianfan Xue, Bill Freeman, and Josh Tenenbaum. Learning a probabilistic latent space of object shapes via 3d generative-adversarial modeling. In *Advances in neural information processing systems*, pages 82–90, 2016.

[30] Maxim Tatarchenko, Alexey Dosovitskiy, and Thomas Brox. Octree generating networks: Efficient convolutional architectures for high-resolution 3d outputs. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2088–2096, 2017.

[31] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. IEEE, 2016.

[32] Özgün Çiçek, Ahmed Abdulkadir, Soeren S Lienkamp, Thomas Brox, and Olaf Ronneberger. 3d u-net: learning dense volumetric segmentation from sparse annotation. In *International conference on medical image computing and computer-assisted intervention*, pages 424–432. Springer, 2016.

[33] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.

[34] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*, 2015.

[35] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.

[36] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[37] B. A. Barad et al. Emringer: side chain-directed model and map validation for 3d cryo-electron microscopy. *Nat. Methods, 12, 943–946.*, 2015.

[38] (D. Liebschner et al. Macromolecular structure determination using x-rays, neutrons and electrons: recent developments in phenix. *Acta Crystallographica Section D, vol. 75, no. 10, pp. 861–877*, 2019.