3D Reconstruction of Protein Complex Structures Using Synthesized Multi-View AFM Images

Jaydeep Rade¹ Soumik Sarkar¹ Anwesha Sarkar¹ Adarsh Krishnamurthy¹

Iowa State University¹ {jrrade, soumiks, anweshas, adarsh}@iastate.edu

Abstract

Recent developments in deep learning-based methods demonstrated its potential to predict the 3D protein structures using inputs such as protein sequences, Cryo-Electron microscopy (Cryo-EM) images of proteins, etc. However, these methods struggle to predict the protein complexes (PC), structures with more than one protein. In this work, we explore the atomic force microscope (AFM) assisted deep learning-based methods to predict the 3D structure of PCs. The images produced by AFM capture the protein structure in different and random orientations. These multi-view images can help train the neural network to predict the 3D structure of protein complexes. However, obtaining the dataset of actual AFM images is timeconsuming and not a pragmatic task. We propose a virtual AFM imaging pipeline that takes a 'PDB' protein file and generates multi-view 2D virtual AFM images using volume rendering techniques. With this, we created a dataset of around 8K proteins. We train a neural network for 3D reconstruction called Pix2Vox++ using the synthesized multi-view 2D AFM images dataset. We compare the predicted structure obtained using a different number of views and get the intersection over union (IoU) value of 0.92 on the training dataset and 0.52 on the validation dataset. We believe this approach will lead to better prediction of the structure of protein complexes.

1 Introduction

Accurate prediction of a 3D protein structure based on its amino acid sequence has been a challenging research task for more than a five decades [1] and is fundamentally important since the structure defines the functionality of proteins [2]. Researchers have been utilizing and improving experimental methods like x-ray crystallography, nuclear magnetic resonance (NMR) spectroscopy, and cryo-electron microscopy (cryo-EM) to determine the 3D structures of proteins [3–6]. Improvements in experimental methods have increased the availability of structural data for a large number of proteins. In parallel, the advancements in deep learning (DL) have had a significant impact on protein structure prediction, with the development of AlphaFold2 [7] and RoseTTAFold [8]. Combining DL with advancements in cryo-EM imaging of single particle analysis enabled the determination of protein structure and complexes. However, Cryo-EM images can be very noisy, with a signal-to-noise ratio as low as -20dB [9]. In addition, the Cryo-EM experiments are expensive and can cost up to \$1500 per structure [10].

Another way to infer the structure of proteins is to learn the mapping from a protein's sequence of amino acids to its structure. These methods include the use of Position Specific Scoring Matrices (PSSM) [11], which compute the single residue mutation frequency distribution, and Multiple Sequence Alignment (MSA), which determines the co-evolutionary information [12, 13]. AlphaFold2



Figure 1: AlphaFold2 multimer (right) fails to predict the structure of WRC protein complex (left).

uses entire MSAs and a template-directed inference process. However, AlphaFold2 is not as effective in predicting the structure of protein complexes where proteins are bonded together. For example AlphaFold2-multimer [14], which is trained specifically for protein complex structure prediction, fails to predict the structure of the WRC protein complex [15–18] as seen in Figure 1. As *In Silico* methods are unable to infer the structure of protein complexes in certain cases, and Cryo-EM methods can be expensive and time-consuming, we propose to train the DL model with images captured using Atomic Force Microscopy (AFM). AFM is a part of scanning probe microscopy, a non-invasive technique that provides the platform for high-resolution imaging of proteins in their physiological environment without prolonged, complicated sample preparation techniques (freezing, drying, dye-tagging) that ultimately damage these soft samples.

With these advantages of AFM, we propose a novel approach to infer the 3D structure of protein complexes using a neural network that is trained using multi-view AFM images. We use Pix2Vox++ [19] neural network architecture for predicting the 3D protein structure. We explain the details of the architecture in Section 2. Even though AFM is not as expensive as Cryo-EM, experimentally acquiring the AFM image dataset for training is not a practical solution. Hence we develop a virtual AFM imaging pipeline to generate multiple views of the protein complexes using a volume rendering technique. We discuss this virtual imaging process in detail in Section 3. We show our preliminary results in Section 5. Finally, we conclude our work with some future direction in Section 6.

2 Method

We implement the Pix2Vox++ neural network architecture, which was proposed for the 3D reconstruction of the shape of an object using single or multiple images. The goal is to predict the 3D shape represented in the voxel occupancy grid where 0 value indicates the empty voxel and 1 indicates the occupied voxel. It includes four blocks: encoder, decoder, multi-scale context-aware merger, and a refiner. The overall forward pass is summarized in Figure 2. The encoder takes each input image and transforms it into latent space feature maps. Next, the decoder processes each latent feature map and generates a 3D volume for each input image. These generated 3D volumes are fused into one using a multi-scale context-aware fusion block. It adaptively assigns the weight for each voxel in all 3D volumes to select high-quality reconstruction. In the end, the refiner further improves the



Figure 2: *The architecture of Pix2Vox++ neural network which reconstructs the 3D shape from multi-view 2D images.*



Figure 3: Volume rendering of 3D voxelized protein structure to obtain AFM-like 2D images.

reconstruction by rectifying the inaccurate parts in generated fused 3D shape and outputs the final 3D shape.

The encoder uses 2D convolutional layers to transform input images into latent representation feature maps, which decoders will use to generate 3D volume per image. Using 2D convolutional layer followed by non-linearity ReLU function, batch-normalization, and then max-pooling operation to encode the image into $4 \times 4 \times 256$ feature map. The feature map is then reshaped to $2 \times 2 \times 2 \times 512$ and fed into the 3D transposed convolutional layer, followed by the ReLU function, which upsamples the feature map into a 3D volume for each input image. As the inputs are multi-view images, reconstruction from the visible parts is better than the invisible parts, and this observation can be used to adaptively select the high-quality parts from each generated volume. Multi-scale context-aware fusion block combines these selected high-quality parts into a single 3D volume. This block uses convolutional layers with non-linearity functions to generate the score per voxel per 3D volume, a higher score for better quality reconstruction. Finally, the refiner tries to correct the incorrect reconstruction in the fused volume. The refiner uses a network similar to U-Net [20] to output the final 3D shape.

3 Dataset

We have developed the pipeline to generate the virtual multi-view AFM images from a 'PDB' format protein file. First, we download the 'PDB' files of proteins from the Protein Data Bank. Using PyMol software often used to visualize proteins, we transform the 'PDB' format structure into a 3D mesh 'OBJ' file format. Proteins can be visualized in various representations such as surfaces, spheres, sticks, cartoons, or ribbons. We export the 'OBJ' files for surface representations. We then use GPU-accelerated algorithm [21] to voxelized the 3D mesh. The resolution of the voxelization is set to 256 voxels along the longest dimension, and then zero padding is performed along the remaining dimensions, so the final voxelized structure has the resolution of $256 \times 256 \times 256$. For the current work, we down-sample this high-resolution structure to a $32 \times 32 \times 32$ grid.

The GPU-accelerated volume rendering technique is used to get the 2D projections of the voxelized 3D protein structure as shown in Figure 3 to replicate the AFM imaging task. We rotate the 3D structure randomly in 3D space to obtain the multi-view images. A total of 25 multi-view images are saved using an image format. Using this data pipeline, we construct the dataset of around 8K protein samples containing the pairs of voxelized structures and corresponding 25 multi-view virtual AFM images. Additionally, we performed high-resolution (nanometer scale) imaging of the single molecules of the WRC protein-complex [15–18, 22]. We also generate the virtual AFM images of

the WRC complex. These WRC samples, both actual and synthesized images, will be the testing samples instead of testing the trained network on the actual AFM images.

4 Training Details

We split the dataset into training and validation datasets. The 80% of the data forms the training set, and the remaining 20% forms the validation set. As we have generated 25 random views for each protein structure, we can experiment with a different number of views as input to the network. To increase the number of samples, we repeat the samples four times with different views as inputs but the same 3D structure as output. The input RGB image resolution used is $224 \times 224 \times 3$ and the output resolution is $32 \times 32 \times 32$. The total number of parameters in the neural network is around 58.9Mn. A binary cross-entropy (BCE) loss function is used as we are predicting the occupancy 3D voxel grid. We use a batch size of 32 and train for 150 epochs. We use the learning rate of 0.001 with the Adam optimizer [23]. We used NVIDIA A-100 GPUs for running all the experiments.

5 Results and Discussion

We compute the intersection over union (IoU) metric between ground truth and predicted 3D structure; the value is between 0 and 1, while 0 denotes the no overlap and 1 denotes the perfect overlap. We evaluate the IoU value on training, validation data set, and test data with actual AFM images. We also compare the results using a different number of views. We summarize these experiments and corresponding metrics in Table 1. We observe a significant difference between training and validation metrics, which shows that we are overfitting. We tried other techniques to overcome the overfitting issue, such as using L-2 regularization, different learning rates, and a learning rate scheduler, which didn't help to avoid the problem.

#views	Train IoU	Validation IoU	Test IoU
3	0.89	0.50	0.19
5	0.91	0.51	0.24
8	0.92	0.51	0.20
10	0.92	0.52	0.22

Table 1: Comparing the IoU metrics for different number of views.

We tried training the network with the actual images added to the training set to determine if the network could learn from actual AFM images. We visualize the prediction and compare it to the ground truth structure in Figure 4. We also show some anecdotal samples from the training and validation set in Figure 5 and Figure 6, respectively.



Figure 4: Comparing the visualization of ground truth and predicted structure using actual AFM images. We train the network with the actual AFM images added to the training dataset.



Figure 5: Predictions on samples from training dataset.



Figure 6: Predictions on samples from validation dataset.

6 Conclusions

In this work, we propose to use AFM imaging-based deep learning 3D reconstruction approach to predict the 3D structure of protein complexes. We utilized the GPU-accelerated volume rendering technique to replicate the AFM imaging process through which we can generate the virtual AFM-like images from the "PDB" formal file of proteins. We generated 8K protein samples and trained the Pix2Vox++ architecture for 3D reconstruction using multi-view AFM images. Although we are facing the problem of overfitting, we could validate the approach by adding actual AFM images into the training data set and predicting the structure of the WRC protein complex. In future work, we will generate more data samples. We plan to provide additional input to the neural network, such as protein sequence embedding using large language models trained for protein sequences.

References

- Ken A. Dill, S. Banu Ozkan, M. Scott Shell, and Thomas R. Weikl. The protein folding problem. Annual Review of Biophysics, 37(1):289–316, 2008. doi: 10.1146/annurev.biophys.37.092707.153558. URL https://doi.org/10.1146/annurev.biophys.37.092707.153558. PMID: 18573083.
- Ken A. Dill and Justin L. MacCallum. The protein-folding problem, 50 years on. Science, 338(6110): 1042-1046, 2012. doi: 10.1126/science.1219021. URL https://www.science.org/doi/abs/10. 1126/science.1219021.
- [3] Kurt Wüthrich. The way to nmr structures of proteins. *Nature structural biology*, 8:923–5, 12 2001. doi: 10.1038/nsb1101-923.
- [4] Michael C. Thompson, Todd O. Yeates, and José A. Rodríguez. Advances in methods for atomic resolution macromolecular structure determination. *F1000Research*, 9, 2020.
- [5] Xiao chen Bai, Greg McMullan, and Sjors H.W Scheres. How cryo-em is revolutionizing structural biology. *Trends in Biochemical Sciences*, 40(1):49-57, 2015. ISSN 0968-0004. doi: https: //doi.org/10.1016/j.tibs.2014.10.005. URL https://www.sciencedirect.com/science/article/ pii/S096800041400187X.
- [6] Mariusz Jaskólski, Zbigniew Dauter, and Alexander Wlodawer. A brief history of macromolecular crystallography, illustrated by a family tree and its nobel fruits. *The FEBS Journal*, 281, 2014.
- [7] John M. Jumper et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596:583 589, 2021.
- [8] Minkyung Baek et al. Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 373(6557):871–876, 2021. doi: 10.1126/science.abj8754. URL https://www. science.org/doi/abs/10.1126/science.abj8754.
- [9] Harshit Gupta, Thong H. Phan, Jaejun Yoo, and Michael A. Unser. Multi-cryogan: Reconstruction of continuous conformations in cryo-em using generative adversarial networks. In ECCV Workshops, 2020.
- [10] Michael A Cianfrocco and Andres E Leschziner. Low cost, high performance processing of single particle cryo-electron microscopy data in the cloud. *eLife*, 4:e06664, may 2015. ISSN 2050-084X. doi: 10.7554/eLife.06664. URL https://doi.org/10.7554/eLife.06664.
- [11] Bingqian Guo, Christopher O. Audu, Jared C. Cochran, Dale F. Mierke, and Maria Pellegrini. Protein engineering of the n-terminus of nemo: Structure stabilization and rescue of ikkβ binding. *Biochemistry*, 53 (43):6776–6785, 2014. doi: 10.1021/bi500861x. URL https://doi.org/10.1021/bi500861x. PMID: 25286246.
- [12] Andrew W. Senior et al. Improved protein structure prediction using potentials from deep learning. *Nature*, 577:706–710, 2020.
- [13] Jianyi Yang, Ivan Anishchenko, Hahnbeom Park, Zhenling Peng, Sergey Ovchinnikov, and David Baker. Improved protein structure prediction using predicted interresidue orientations. *Proceedings of the National Academy of Sciences*, 117(3):1496–1503, 2020. doi: 10.1073/pnas.1914677117. URL https: //www.pnas.org/doi/abs/10.1073/pnas.1914677117.
- [14] Richard Evans et al. Protein complex prediction with alphafold-multimer. bioRxiv, 2022. doi: 10.1101/ 2021.10.04.463034. URL https://www.biorxiv.org/content/early/2022/03/10/2021.10.04. 463034.
- [15] Vassilis Koronakis, Peter J Hume, Daniel Humphreys, Tao Liu, Ole Hørning, Ole N Jensen, and Emma J McGhie. Wave regulatory complex activation by cooperating gtpases arf and rac1. *Proceedings of the National Academy of Sciences*, 108(35):14449–14454, 2011.
- [16] Klemens Rottner, Theresia EB Stradal, and Baoyu Chen. Wave regulatory complex. *Current Biology*, 31 (10):R512–R517, 2021.
- [17] Ayman M Ismail, Shae B Padrick, Baoyu Chen, Junko Umetani, and Michael K Rosen. The wave regulatory complex is inhibited. *Nature structural & molecular biology*, 16(5):561–563, 2009.
- [18] Baoyu Chen, Klaus Brinkmann, Zhucheng Chen, Chi W Pak, Yuxing Liao, Shuoyong Shi, Lisa Henry, Nick V Grishin, Sven Bogdan, and Michael K Rosen. The wave regulatory complex links diverse receptors to the actin cytoskeleton. *Cell*, 156(1-2):195–207, 2014.

- [19] Haozhe Xie, Hongxun Yao, Shengping Zhang, Shangchen Zhou, and Wenxiu Sun. Pix2vox++: Multi-scale context-aware 3d object reconstruction from single and multiple images. *CoRR*, abs/2006.12250, 2020. URL https://arxiv.org/abs/2006.12250.
- [20] 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.
- [21] Gavin Young and Adarsh Krishnamurthy. Gpu-accelerated generation and rendering of multi-level voxel representations of solid models. *Computers and Graphics*, 75:11–24, 2018. ISSN 0097-8493. doi: https://doi.org/10.1016/j.cag.2018.07.003. URL https://www.sciencedirect.com/science/ article/pii/S009784931830102X.
- [22] Sheng Yang, Yijun Liu, Abbigale Brown, Matthias Schaks, Bojian Ding, Daniel A. Kramer, Li Ding, Olga Alekhina, Daniel D. Billadeau, Saikat Chowdhury, Junmei Wang, Klemens Rottner, and Baoyu Chen. Arf gtpase activates the wave regulatory complex through a novel binding site. *bioRxiv*, 2022. doi: 10.1101/2022.05.13.491857. URL https://www.biorxiv.org/content/early/2022/05/13/2022. 05.13.491857.
- [23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Checklist

- 1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
- 2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [N/A]
 - (b) Did you include complete proofs of all theoretical results? [N/A]
- 3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [No]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
- 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
- 5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Appendix

A Examples of Virtual AFM images

We show some examples of generated virtual multi-view AFM images. We also compare the actual and virtual AFM images of WRC protein complex.





Actual AFM images Virtual AFM images

Figure 7: Comparing the virtual and actual AFM images.



Figure 8: Showing the examples for virtual multi-view AFM images.