

Structural bioinformatics

Atomic protein structure refinement using all-atom graph representations and SE(3)-equivariant graph transformer

Tianqi Wu¹, Zhiye Guo¹, Jianlin Cheng  ^{1,*}

¹Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO 65211, United States

*Corresponding author. Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO 65211, USA.
E-mail: chengji@missouri.edu

Associate Editor: Lenore Cowen

Abstract

Motivation: The state-of-art protein structure prediction methods such as AlphaFold are being widely used to predict structures of uncharacterized proteins in biomedical research. There is a significant need to further improve the quality and nativeness of the predicted structures to enhance their usability. In this work, we develop ATOMRefine, a deep learning-based, end-to-end, all-atom protein structural model refinement method. It uses a SE(3)-equivariant graph transformer network to directly refine protein atomic coordinates in a predicted tertiary structure represented as a molecular graph.

Results: The method is first trained and tested on the structural models in AlphaFoldDB whose experimental structures are known, and then blindly tested on 69 CASP14 regular targets and 7 CASP14 refinement targets. ATOMRefine improves the quality of both backbone atoms and all-atom conformation of the initial structural models generated by AlphaFold. It also performs better than two state-of-the-art refinement methods in multiple evaluation metrics including an all-atom model quality score—the MolProbity score based on the analysis of all-atom contacts, bond length, atom clashes, torsion angles, and side-chain rotamers. As ATOMRefine can refine a protein structure quickly, it provides a viable, fast solution for improving protein geometry and fixing structural errors of predicted structures through direct coordinate refinement.

Availability and implementation: The source code of ATOMRefine is available in the GitHub repository (<https://github.com/BioinfoMachineLearning/ATOMRefine>). All the required data for training and testing are available at <https://doi.org/10.5281/zenodo.6944368>.

1 Introduction

Every cell in the human body contains proteins. Protein participates in most cellular processes, ranging from DNA replications to immune responses. Protein functions are intimately connected with their 3D shapes. Therefore, predicting the protein structure from sequence has been a long-standing grand challenge in computational biology. Recently, AlphaFold (Senior et al. 2020; Jumper et al. 2021) is shown to predict highly accurate tertiary structures for most proteins, which is considered a big advance in the field. However, there are still some limitations in the AlphaFold predicted structures. The recent application of AlphaFold2 (Tunyasuvunakool et al. 2021) to predicting the structures in the human proteome showed that the conformation of 58% of the total residues was of high accuracy with the predicted confident score pLDDT (1) >70, leaving the rest 42% of the total residues with the confidence score pLDDT 70. Besides, a strong correlation between the AlphaFold model quality and the availability of homologous templates in the Protein Data Bank (PDB) has been observed in a few benchmarking studies (Cretin et al. 2021; Pearce and Zhang 2021; Jones and Thornton 2022), suggesting that there is still a room to improve the quality of AlphaFold models, particularly for proteins without homologous templates in the PDB. Moreover, current protein structure prediction methods including AlphaFold have been focused on predicting the backbone structure of proteins correctly without emphasizing

improving the nativeness and all-atom geometry of predicted structures, leaving significant room to improve the all-atom quality of predicted structures (Bhattacharya and Cheng 2013). Therefore, there is a significant need to further refine the protein structures predicted by state-of-the-art methods such as AlphaFold to improve their usability in biomedical research.

Currently, typical model refinement methods apply molecular dynamics (MD) simulation, energy minimization, or fragment assembly to refine input protein structures. Successful MD-based methods (Heo et al. 2013; Mirjalili et al. 2014; Heo and Feig 2018; Lee et al. 2019; Heo et al. 2021) are physics-based approaches to sampling multiple MD trajectories following the physical principles regarding atomic interactions, which are computation-intensive and time-consuming. Energy minimization-based method (Xu and Zhang 2011; Bhattacharya and Cheng 2013) focus on repacking the backbone and side-chain atoms with composite physics and knowledge-based force fields. Fragment assembly-based methods are like knowledge-based methods, taking advantage of template fragment information in the PDB as well as statistical potentials. A notable method is Rosetta (Hirunuma et al. 2021), which uses predicted estimated local structural errors to inform the fragment assembly, followed by side-chain rebuilding and energy minimization in all-atom representation. Though those methods prove to be effective in the refinement of some protein structures, they

require extensive conformation sampling and a lot of computing resources.

Deep learning has recently been applied to improve the geometric property of the protein 3D structure (Senior et al. 2020; Baek et al. 2021; Hiranuma et al. 2021). Graph neural networks were used by GNNRefine (Jing and Xu 2021) to refine the backbone atoms of protein structure. It largely relies on a Rosetta protocol for the full-atom model reconstruction. In the refinement module of RoseTTAFold (Baek et al. 2021), a SE(3)-equivariant graph transformer (Fuchs et al. 2020) is used to refine backbone atoms without directly using machine learning to leverage and improve side-chain atoms in a protein structure. However, it produces a refined model with only backbone atoms and cannot be used as a standalone tool to refine a third-party model.

Inspired by the application of geometric deep learning to molecular structure prediction that can avoid the expensive and extensive conformation sampling, here we present ATOMRefine, a new SE(3)-equivariant transformer network based on a novel all-atom representation of atom types, amino acid types, atom–atom distances, and covalent bonds for refining protein structures in the full-atom scale, which is different from RoseTTAFold. Its graph representation of all the atoms of a protein structure enables the network to leverage sequence-based and spatial information from the entire protein structures to update node and edge features and catch the global and local structural variation from the initial model to the native structure iteratively, which is different from RoseTTAFold’s refinement module that refines backbone atoms only. The 3D-equivariance makes it possible for ATOMRefine to learn essential structural properties regardless of the rotation and translation of the input structure. The network outputs the refined coordinates of all the atoms directly, without using any external protein full-atom reconstruction protocol. To the best of our knowledge, ATOMRefine is the first end-to-end all-atom 3D-equivariant transformer network approach to refine the protein model prediction on the full atom scale.

Evaluated on both AlphaFold and the 14th Critical Assessment of Techniques for Protein Structure Prediction (CASP14) datasets, ATOMRefine improves the quality of both backbone and all atoms of the initial structural model in terms of GDT-TS score, GDT-HA score, RMSD, IDDT, and Molprobity. Noticeably, ATOMRefine can maintain or improve the model quality over the initial models generated by AlphaFold and other structure prediction methods, and generate far fewer model degradation cases than the existing refinement methods.

2 Materials and methods

ATOMRefine is an end-to-end protein refinement method based on a SE(3)-equivariant graph transformer neural network. It directly predicts refined atomic coordinates of all the atoms as output from the initial coordinates of all the atoms in an input structure. To avoid the bond geometry violation, a final relaxation step by Amber (Salomon-Ferrer et al. 2013) is added in the ATOMRefine pipeline. Its simplified version based on the same deep learning architecture is also used to refine the coordinates of backbone atoms only. The overall framework of ATOMRefine is illustrated in Fig. 1. Details of the graph representation, network architecture, training and test data, and evaluation metrics are described as follows.

2.1 All-atom graph-based representation of protein structure and SE(3) graph transformer architecture

In this work, a protein structure is considered as a set of nodes each of which represents an atom in the protein. Each atom i has a 3D coordinate (x_i, y_i, z_i) that can be used to calculate the pairwise spatial relations between atoms. A protein structure is represented as a graph of the nodes in which the edges describe the relationships between the nodes (i.e. atoms).

Each node has atom features including one-hot encoding of atom types (a binary vector indicating 37 atom types) and the types of amino acids that the atom belongs to. Each node also has x, y, z coordinates as variable features that will be updated.

As illustrated in Fig. 1a, each node is connected to the k ($k \leq 128$) nearest neighboring nodes selected by the Euclidean distance between atom 3D coordinates through edges. ATOMRefine employs a neighborhood aggregation approach to enhance predictive accuracy by capturing the local environment of each atom (node) within a neighborhood and refining local conformation during training. By filtering out less relevant information and focusing on important local features, the model can identify and correct local errors. Due to the GPU resource limitation, we use this K-Nearest Neighbors (KNN) graph representation for the atom-level protein structure refinement. Increasing the value of k allows the model to capture more structural information for each node but requires more GPU memory. Therefore, we choose a large k equal to 128 within the limits of the GPU (NVIDIA V100 32 GB memory) available to us. Because each residue can only interact with a maximum of 6–8 other residues (generally less than 100 heavy atoms) due to steric restrictions, the KNN graph with a sufficiently large k (e.g. 128) can capture sufficient local information at all-atom level to predict conformations of atoms and residues under the restriction of the available hardware. Moreover, because the local information is passed from node to node in the graph through the attention mechanism, ATOMRefine can learn global structural features to make prediction.

Besides node features, six edge features are generated, including one distance-based edge feature, one covalent bond edge feature, one relative position edge feature, and three relative orientation edge features. For the distance-based edge feature, we use the radial basis function to convert the distance (d) between two nodes as features:

$$f_{d0} = \exp \left(-\frac{d^2}{r_d^2} \right), \text{ where } d \text{ is the Euclidean dis-}$$

tance between two nodes of an edge, d and r_d are hyperparameters. We set $r_d \approx 0.57$, $d^0 \approx [0, r_d, 2r_d, \dots, 35r_d] / 2 \text{ \AA}$; 20 \AA , following the work of RoseTTAFold. So, for each edge, there are 36 distance-based edges.

We also use a binary covalent bond edge feature to represent the local covalent bond connectivity between atoms. An adjacent bond matrix (M) is calculated from the atom–atom distance matrix (D) to detect if there is a covalent bond between two atoms according to the work of Graphein (Jamasb et al. 2020). We parse the atomic Euclidean distance matrix D into the binary covalent bond adjacent matrix M as shown in Equation (1),

$$M_{ij} = \begin{cases} 1; & 0.4 < D_{ij} < r_0; \\ 0; & \text{otherwise} \end{cases}; \quad (1)$$

where i, j are the atom positions (indices), and the thresholding parameter r is a set of covalent radii based on different

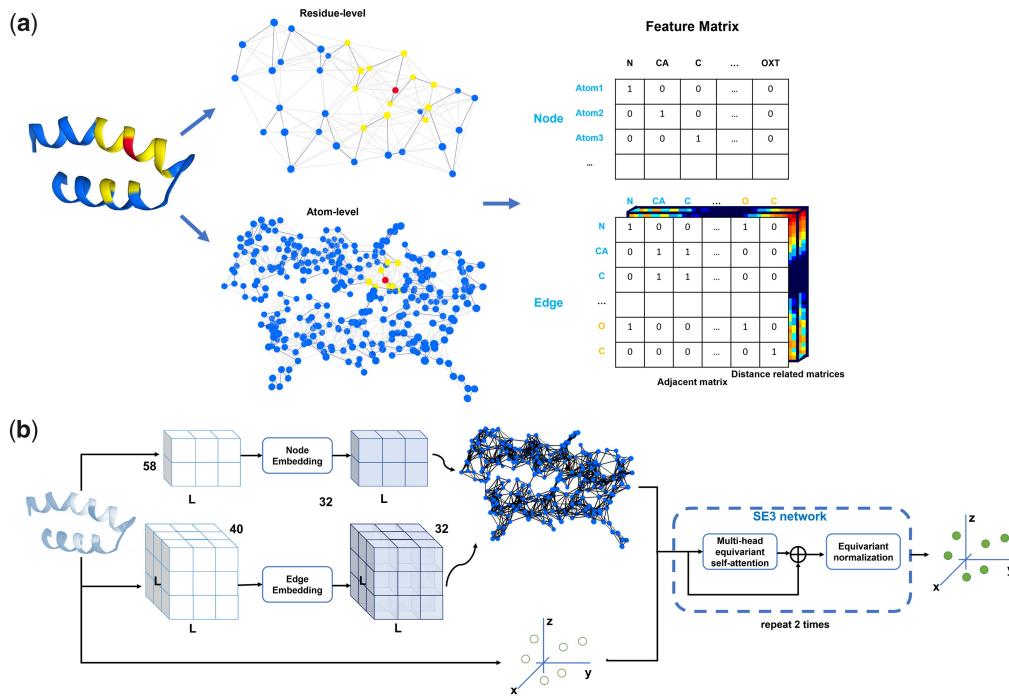


Figure 1. The ATOMRefine framework. (a) ATOMRefine graph representation of a protein structure at all-atom level or backbone atom-level. A node is used to represent an atom. The 3D protein structure is encoded as the atomic features (node features) and inter-atom features (edge features including adjacent bond matrix and distance-related matrices). Each node (e.g. a node in red) is connected to the k nearest neighboring nodes (nodes in yellow) selected by the Euclidean distance calculated from atom 3D coordinates. The covalent bond edge between atoms shown as the solid blue line in the graph is also an edge feature stored in the adjacent bond matrix. (b) The deep learning architecture of ATOMRefine. Each block in the $SE(3)$ transformer network consists of one equivariant graph neural network attention block and one $SE(3)$ -equivariant normalization layer.

atom types. 1 indicates there is a bond between two atoms (see [Supplementary Table S1](#) for details).

Similar to the work of Octavian-Eugen Ganea ([Ganea et al. 2021](#)) and trRosetta ([Yang et al. 2020](#)), we also use the relative position and relative orientation features for edges based on the local coordinate system. We construct the local coordinate system based on each amino acid residue position (index) in a protein model (for atoms of the same amino acid, they share the same local coordinate basis). As shown in [Fig. 2](#), for each residue i , we define the Ca coordinate as the origin, the unit vector pointing from Ca atom to C atom as u_i , and the unit vector pointing from Ca atom to N atom as y_i (on the y axis). The normal of the plane C–Ca–N is defined as z_i (on the z -axis), where $z_i \perp \frac{u_i y_i}{\|u_i y_i\|}$: Naturally, we define $x_i \perp y_i z_i$ (on the x -axis). In total, x_i , y_i and z_i consist of the basis of residue i 's local coordinate system. As shown in [Equation \(2\)](#), the relative position edge feature $p_{im;jn}$ denotes the relative position of atom n in the residue position j to atom m in the residue i . $atom_{jn}$ denotes the coordinate of atom n in the residue position j .

$$p_{im;jn} \triangleq x_i^T y_j^T z_i^T \triangleq atom_{jn} - atom_{im} \quad (2)$$

As shown in [Equation \(3\)](#), relative orientation features $q_{im;jn}$, $k_{im;jn}$, $t_{im;jn}$ denote the relative orientation of atom n in residue position j to atom m in the residue position i .

$$\begin{aligned} q_{im;jn} &\triangleq x_i^T y_i^T z_i^T \triangleq x_j^T \triangleq k_{im;jn} \triangleq x_i^T y_i^T z_i^T \triangleq y_j^T \triangleq t_{im;jn} \\ &\triangleq x_i^T y_i^T z_i^T \triangleq z_j^T \end{aligned} \quad (3)$$

With atom and atom–atom relationship features encoded as node features and edge features above, the protein structure

Feature Matrix

Node	N	CA	C	...	OXT
Atom1	1	0	0	...	0
Atom2	0	1	0	...	0
Atom3	0	0	1	...	0
...					

Edge	N	CA	C	...	O	C
N	1	0	0	...	1	0
CA	0	1	1	...	0	0
C	0	1	1	...	0	0
...						
O	1	0	0	...	1	0
C	0	0	0	...	0	1

Adjacent matrix Distance related matrices

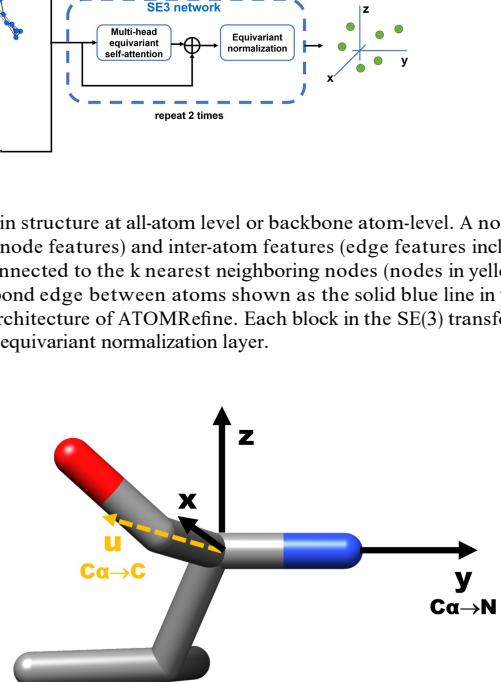


Figure 2. Representation of a single valine residue and its local coordinate system. We define the atom Ca as the origin (at the center). The y -axis points from the atom Ca to the atom N. The x -axis is placed in the plane of C–Ca–N. Following the right-hand coordinate system, the z -axis is the normal of the plane.

can be encoded in the graph. [Figure 1a](#) shows the scheme of the graph representation of a protein model at the atom level. The detailed atomic and residue-based features are presented in [Supplementary Table S1](#).

The general network architecture of ATOMRefine is illustrated in [Fig. 1b](#). We parse the initial protein model as the node and edge features to build a graph representation. The graph is then fed into the $SE(3)$ -transformer to refine the given 3D atom coordinates. All features of each node except for 3D coordinates correspond to $SE(3)$ type 0 node feature, and the 3D coordinates of each node (atom) correspond to $SE(3)$ type 1 node feature. The embedding size of the input node and edge features are set to 32. A $SE(3)$ transformer is used to predict the coordinate shifts between the initial model and the native structure. The target shift is calculated after the initial model and the native structure are superimposed. As the

target shift is invariant to the rotation and translation of either the initial model or the native structure, this helps train a deep learning model invariant/equivariant to the rotation and translation of input and output. It is worth noting that the final refined structure is simply equal to the initial model plus the predicted shift. The SE(3) transformer consists of two SE3 equivariant attention blocks, including one multi-head attention block with 16 channels and 4 attention heads, and one SE3 equivariant normalization layer. For each attention block, queries are the linear projection of the graph node features. Keys and values are from graph edge features computed by TFN (Thomas et al. 2018) layers. ATOMRefine is implemented on top of the deep learning framework PyTorch Lightning (Falcon 2019) and Deep Graph Library (Wang et al. 2019).

2.2 The training dataset and test dataset of ATOMRefine

Training data: We download the predicted protein models from AlphaFoldDB (version 1 released before September 2021). We use MMseqs2 (Steinegger and Söding 2017) to remove the sequence duplication first and then match the remaining protein sequences with the native structures that exist in Protein Data Bank. A structural model is matched with a true structure if the following criteria are met: (i) the model sequence matches with the native sequence; (ii) protein sequence length ≥ 50 . In total, 13 121 are selected as the initial models and their true structures are used as labels. ATOMRefine is trained and validated on the training data via 10-fold cross-validation. During training, protein structures with >1500 residues are cropped to fit the GPU memory. We set Adam as the optimizer with parameters: $b1 \approx 0.9$, $b2 \approx 0.999$, and weight decay ≈ 0.001 . We set the batch size as 1 and use the mean squared error between predicted coordinate shifts and true coordinate shifts of atoms as the loss function. We set the number of training epochs to 50 with the early stopping when there are no improvements in the validation loss for five consecutive epochs. Ten ATOMRefine deep learning models have been trained in this way. We choose five trained deep learning models that have the lowest validation loss as the final deep learning models to make inference. The loss is the mean square error between predicted and true coordinates of atoms. The five deep learning models are used by ATOMRefine to generate five refined protein structural models for an input structural model.

Test data: We use three test datasets to evaluate the methods: an AlphaFoldDB test set containing 193 protein targets retrieved from the AlphaFoldDB, the CASP14 dataset containing 69 regular targets, and the CASP14 refinement dataset (7 protein targets). The CASP14 refinement targets are a subset of the CASP14 regular targets, selected by CASP organizers for challenging predictors to make a structural refinement. Any sequences in the training data that has 30% identity with any sequences in the three test datasets have been removed in the training data preparation so that there is no overlap between the training data and the three all three test sets (e.g. sequence identity $< 30\%$). For each target in the CASP14 dataset, AlphaFold2 is run to generate start models. For the CASP14 refinement dataset, we use the initial models provided by CASP14 organizers as the start models, which were generated by traditional protein structure prediction methods other than AlphaFold2. All the true structures for the targets in the test datasets are obtained from the PDB.

2.3 Evaluation metrics

To compare the model quality of initial models and refined models, we follow the same approach described in Jing and Xu (2021) that generate five refined model from five trained deep learning models respectively to compare with the initial models. We use GDT-TS (Zemla 2003), GDT-HA, RMSD of the Ca atoms, IDDT (Mariani et al. 2013) and Molprobity score (Williams et al. 2018) as five main evaluation metrics. GDT-TS is the global distance score. It ranges from 0 to 100% (or simply from 0 to 1), a higher value indicating better model accuracy. GDT-HA is the high-accuracy version of the GDT-TS score with smaller distance cutoffs. RMSD of the Ca atoms measures the root mean square deviation of the Ca atoms in a protein model from its native structure, describing the accuracy of the positions of the Ca atoms. A lower RMSD means better quality. Local Distance Difference Test (IDDT) uses the distance differences of atom pairs to measure the local conformation quality of each residue (higher, better). IDDT scores of all the residues can be averaged to measure the quality of a protein structural model. The MolProbity score assesses the quality of all the atoms of a model including side-chain atoms. It considers atom contacts, atom clashes, bond lengths and angles, and torsion angles. A lower Molprobity score indicates better model quality.

3 Results

3.1 Comparison of ATOMRefine with other refinement methods in terms of backbone quality

Geometric deep learning-based approaches have been applied to protein structure refinement, among which GNNRefine yields some quality improvement from initial models. However, its machine learning component heavily focuses on the backbone atom refinement, and largely relies on the Rosetta refinement protocol for the final full-atom refinement. In contrast, ATOMRefine applies an all-atom SE(3)-equivariant graph transformer to directly refine all the atoms of a protein structure. Directly refining all the atoms has the benefit of generating an all-atom refined structure in an end-to-end fashion, but it requires a much larger molecular graph to represent all the atoms in a protein structure than that representing only backbone atoms (or only Ca atom). To investigate the trade-off of using a full-atom representation, we implement two versions of our method based on the same SE(3)-equivariant graph transformer architecture: (i) ATOMRefine—the all-atom refinement method and (ii) ATOMRefine_backbone—the backbone atom refinement method. Both of them are trained and validated on the same dataset.

We evaluate ATOMRefine, ATOMRefine_backbone, GNNRefine, and a widely-used energy minimization-based method—ModRefiner (Xu and Zhang 2011) on the AlphaFoldDB test set and the structural models of 69 CASP14 targets. For the AlphaFoldDB test set, the structural models from the AlphaFoldDB are used as the initial models. For the CASP14 dataset, AlphaFold2 is used to predict the structures of the CASP14 targets that are used as the initial models. For each initial model, the best of five refined models produced by each method is selected for evaluation against the true experimental structures. The backbone quality of the initial models and the models refined by these methods is reported in Table 1.

Table 1. Performance of ATOMRefine, ATOMRefine_backbone, GNN-Refine, and ModRefiner on AlphaFoldDB test set and CASP14 dataset in comparison with initial models, respectively.^a

Test set	Method	GDT-TS ["]	GDT-HA ["]	RMSD of Ca#	IDDT ["]	IP (%) ["]
AlphaFoldDB test set	Initial model	81.42	69.84	4.09	0.7945	–
	ATOMRefine	81.59	70.04	4.08	0.7978	58.03
	ATOMRefine_backbone	81.69	70.07	3.81	–	61.14
	GNNRefine	79.46	67.16	4.38	0.7791	31.09
	ModRefiner	81.54	69.72	4.07	0.7764	45.08
CASP14 test set	Initial model	77.97	63.91	4.5	0.7826	–
	ATOMRefine	78.14	64.42	4.49	0.7674	75.36
	ATOMRefine_backbone	78.12	64.4	4.43	–	85.51
	GNNRefine	77.54	63.6	4.73	0.773	50.72
	ModRefiner	77.94	63.97	4.50	0.7706	49.28

^a Bold numbers denote the best results. Improvement percentage (IP) denotes the percentage of the models that have been improved by each method in terms of the GDT-HA score. Because all-atom IDDT is used, ATOMRefine_backbone does not have this score.

On average, both ATOMRefine and ATOMRefine_backbone improve the quality of backbone atoms over the initial models in terms of the GDT-TS score, GDT-HA score, and RMSD of the Ca atoms, while the other two methods do not in most situations. On the AlphaFoldDB test set, ATOMRefine performs generally better than GNNRefine and ModRefiner and only ATOMRefine slightly increases IDDT score after refinement. But on the CASP14 test set, all of the refinement tools decrease IDDT score a bit after refinement. Even though the overall improvement in the backbone quality is small, the results are still significant because the recent 14th community-wide Critical Assessment of Techniques of Protein Structure Prediction (CASP14) (Simpkin et al. 2021) showed that few refinement methods can improve the quality of the backbone of initial models on average. Moreover, the t-test shows that the difference between the initial models and ATOMRefine models in terms of the average GDT-HA score is statistically significant (P-value $\leq 2.61\text{E-}10$ on the AlphaFoldDB test set and P-value $\leq 1.90\text{E-}08$ on the CASP14 dataset). ATOMRefine and ATOMRefine_backbone achieve similar performance on the two datasets, indicating that extending the small backbone representation to the all-atom representation for refinement still maintains the effectiveness of refining the backbones of protein structures despite that ATOMRefine needs to accommodate the extra side-chain atom refinement.

Both ATOMRefine and ATOMRefine_backbone perform better than GNNRefine and ModRefiner in terms of most metrics on average. For instance, on the AlphaFoldDB test set, the average GDT-HA score of ATOMRefine_backbone is 0.35 point higher than the following best external method ModRefiner (Table 1). Supplementary Table S2 shows the target-by-target GDT-HA scores of ATOMRefine, ATOMRefine_backbone, GNNRefine, and ModRefiner on the AlphaFoldDB test set. The mean and minimum GDT-HA score of five refined models generated by the methods are also shown in Supplementary Table S2. In terms of the mean and minimum GDT-HA score of the five refined models, ATOMRefine performs better than ATOMRefine_backbone, GNNRefine and ModRefiner.

On the CASP14 test dataset, the average GDT-HA scores of ATOMRefine are 0.45 point higher than the following best external method ModRefiner (Table 1). The detailed GDT-HA scores of ATOMRefine, ATOMRefine_backbone, GNNRefine and ModRefiner on the CASP14 test set are shown in Supplementary Table S3 target by target. The mean and minimum GDT-HA score of five refined models

generated by the methods are also shown Supplementary Table S3. In terms of mean and minimum GDT-HA score of five refined models, ATOMRefine also performs better than ATOMRefine_backbone, GNNRefine and ModRefiner.

The RMSD of ATOMRefine refined models for the AlphaFoldDB test set and CASP14 dataset is 4.08 and 4.49 Angstrom respectively, lower than 4.38 and 4.73 Angstrom of GNNRefine. On average, out of the four methods, only ATOMRefine and ATOMRefine_backbone improve the backbone atom quality of the initial models in terms of most metrics. Figure 3 illustrates the change in the GDT-HA score of the refined model with respect to the initial model of these methods. ATOMRefine and ATOMRefine_backbone improve the quality of the majority of the initial models (58.03–85.51% of the models) to in terms of GDT-HA score. Furthermore, we show the histograms of the difference in GDT-HA scores between the refined models generated from ATOMRefine, ATOMRefine_backbone, GNNRefine, and ModRefiner and the initial models in Supplementary Fig. S1. A positive change in GDT-HA indicates an improvement after the refinement. It clearly shows that our methods are more effective in improving the quality of backbone structures than GNNRefine and ModRefiner.

3.2 Comparison of ATOMRefine with existing methods in terms of all-atom quality

To further investigate the performance of ATOMRefine as an all-atom model refinement method, we compare ATOMRefine and ModRefiner in terms of MolProbity score based on the analysis of all-atom contacts, bond length, atom clashes, torsion angles, and side-chain rotamers. A lower MolProbity score indicates better all-atom quality and higher nativeness of the protein structure. The MolProbity score has been widely used to assess the geometric correctness and nativeness of experimentally determined protein structures before they are deposited into the PDB. The strength parameter for ModRefiner is set at 80, 85, 90, 95, and 100 respectively to generate five refined models. The strength value is in the range [0,100]. A larger value makes the final model closer to the reference model. The model with the lowest MolProbity score is chosen for comparison. We also include GNNRefine in the full-atom level comparison. Though GNNRefine mainly focuses on refining the predicted distances of the backbone atoms, it constructs the final full-atom protein model by using the Rosetta module FastRelax (Chaudhury et al. 2010).

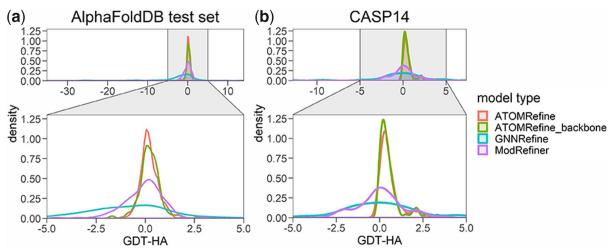


Figure 3. The distribution of quality change (DGDT-HA score) of refined models of ATOMRefine, GNNRefine, ModRefiner, and ATOMRefine_backbone with respect to the initial models. The positive value means the model quality after refinement improves from the starting model and the negative value means the model quality decreases. (a) Results on the AlphaFoldDB test set. (b) Results on the CASP14 dataset.

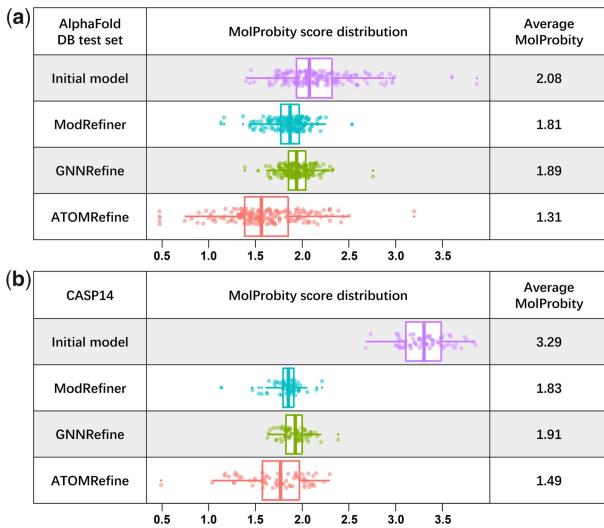


Figure 4. The average all-atom MolProbity score and the MolProbity score distribution of initial models and refined models of ATOMRefine, ModRefiner, and GNNRefine on (a) AlphaFoldDB test set and (b) CASP14 dataset.

The average MolProbity scores of the initial models and the refined models of the three methods on the AlphaFoldDB test set and CASP14 dataset are reported in Fig. 4. The average MolProbity score of the AlphaFoldDB test set and CASP14 dataset is 1.31 and 1.49, much lower than 2.08 and 3.29 of the initial models, indicating a large improvement in the protein geometry and nativeness of the structures predicted by AlphaFold. From the results shown on Fig. 4, ATOMRefine also substantially outperforms GNNRefine and ModRefiner which are also able to improve all-atom quality of AlphaFold2 models in the two datasets to some degree. The larger improvement in all-atom model quality made by ATOMRefine than in backbone model quality is consistent with the previous research (Bhattacharya and Cheng 2013). One possible reason is that ATOMRefine may substantially improve the geometry of the sidechain atoms, such as bond lengths and angles, hydrogen bonding patterns, and the positioning of the atoms.

3.3 Performance of ATOMRefine on different kinds of initial models

The outcome of model refinement is related to the quality of initial models. CASP14 official refinement targets were

carefully selected by CASP organizers to assess the refinement methods considering the quality of initial models and refinement potential. In order to test the room for improvement for different targets, CASP14 selected seven targets each with an initial structure predicted by AF2 (AlphaFold2 group during CASP14 experiment) and a typical structure predicted by one of the other CASP14 groups. Therefore, each target has two different versions (v1/v2: AF2 initial model or other initial model), resulting in 14 models for refinement. In addition, those targets were classified into categories based on their modeling difficulty [FM: free modeling that does not have homologous templates in PDB, hardest targets; FM/TBM: targets in between FM and template-based modeling (TBM), second hardest; and TBM-hard: difficult TBM targets whose homologous templates exist in PDB, but are hard to find, third hardest]. The name, length, classification, and initial model type can be found in Supplementary Table S4. For each target, the GDT-HA scores of the initial models, ATOMRefine, GNNRefine, and ModRefiner are reported in Supplementary Table S4, respectively. Supplementary Table S5 presents the results of Supplementary Table S4 according to the types of the initial models in terms of the GDT-HA score.

Overall, in terms of the average GDT-HA score or the GDT-HA score variation shown in Fig. 5, ATOMRefine outperforms GNNRefine and ModRefiner on most or all targets, respectively. With AF2 models as the initial models shown in Fig. 5a, the average GDT-HA score of ATOMRefine is 70.22, better than the performance of GNNRefine (67.66) and ModRefiner (69.62). ATOMRefine improves the quality of the start AF2 models whose average GDT-HA score is 69.91, but GNNRefine and ModRefiner's GDT-HA score is lower than the GDT-HA score of the start models by 3.22% and 0.41%, respectively. With other CASP14 group models as the initial models shown in Fig. 5b, the average GDT-HA score of ATOMRefine is 40.54, slightly higher than 40.32 of the start models, better than 40.47 of ModRefiner, but slightly lower than GNNRefine 40.75.

In Fig. 5c and d, we also list the per-target GDT-HA score variations by applying the three refinement methods, compared to the initial models starting from either AF2 or other CASP14 groups (the specific variation values for three methods are listed in the figure). For the initial models starting from AF2, ATOMRefine produces much fewer degraded models than the other two methods. Six out of seven ATOMRefine models achieve equal or better model quality, while GNNRefine and ModRefiner show model degradation in most cases. Though GNNRefine performs better than ATOMRefine in terms of the average GDT-HA score on the initial model starting from other non-AlphaFold CASP14 groups, the number of cases achieving equal or better model quality from the two methods are the same. In general, ATOMRefine is able to maintain the model quality or improve the model quality in most cases, regardless of the types of start models.

3.4 Comparison of the speed of ATOMRefine with other methods

In addition to maintaining or improving the model quality, ATOMRefine is also significantly faster than GNNRefine and ModRefiner. We tested the runtime of ATOMRefine, GNNRefine, and ModRefiner on the CASP14 targets with sequence length < 300. Table 2 reports the average runtime for

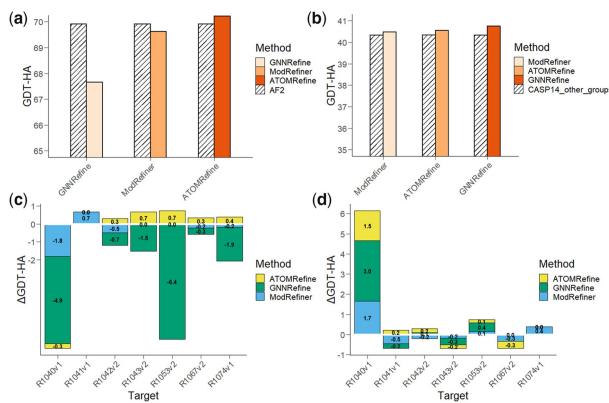


Figure 5. Performance of ATOMRefine, GNNRefine, and ModRefiner on seven CASP14 refinement targets by using two types of initial models as input evaluated by the GDT-HA score. (a) The average performance of the three methods compared with the initial models of AF2. (b) The average performance of the three methods compared with the initial models of other CASP14 groups. (c) GDT-HA score variation of three methods with respect to the initial models of AF2. (d) GDT-HA score variation of three methods with respect to the initial models of other CASP14 groups. (GDT-HA score variation values for each method are listed on (c) and (d); positive values stand for model improvement and negative values stand for model degradation.)

Table 2. Runtime on CASP14 test set.

Method	Average runtime(s)
ATOMRefine	90
GNNRefine	335
ModRefiner	974

each CASP14 target. For a protein with an average length of 156, ATOMRefine typically requires 90 s to complete the entire refinement process on a single Tesla V100 GPU, which is about three times faster than GNNRefine, ten times faster than ModRefiner.

4 Conclusion and future work

In this work, we introduce ATOMRefine, a novel full-atom 3D-equivariant graph transformer method for protein structure refinement. It uses a new full-atom graph to represent atoms, bonds, and coordinates as the node and edge features, which is processed by the equivariant and invariant layers of the SE(3) graph transformer to refine the coordinates of all the atoms. We rigorously evaluate ATOMRefine on three test datasets. Compared to the refinement methods focusing on refining backbone atoms, it has the advantage of directly generating an all-atom refined structure. Moreover, ATOMRefine can improve the quality of both backbones and all atoms including side-chain atoms over the initial input models and outperforms the state-of-the-art deep learning and energy minimization-based methods. The improvement on the backbone of initial models is small but significant, while the improvement on the all-atom conformations is substantial. Finally, once it is trained, ATOMRefine can refine protein structure quickly, making it applicable to proteome-wide protein structure refinement.

We plan to further improve ATOMRefine by training it on a larger dataset consisting of AlphaFold models of more diverse quality, particularly including more low-quality models.

In the current training dataset, 92% of structural models are high-accuracy models, which may limit the amount of improvement that can be made by the deep learning method. Adding more low-quality models into training may make ATOMRefine learn to make larger improvements to the backbone structure on less accurate input.

Supplementary data

Supplementary data are available at [Bioinformatics](https://academic.oup.com/bioinformatics/article/39/5/blad298/7152976) online.

Conflict of interest

None declared.

Funding

This work was supported in part by Department of Energy grants [DE-AR0001213, DE-SC0020400, and DE-SC0021303], NSF grants [DBI2308699, DBI1759934 and IIS1763246], and NIH grants [R01GM093123 and R01GM146340].

References

Baek M, DiMaio F, Anishchenko I et al. Accurate prediction of protein structures and interactions using a three-track neural network. *Science* 2021;373:871–6.

Bhattacharya D, Cheng J. 3Drefine: consistent protein structure refinement by optimizing hydrogen bonding network and atomic-level energy minimization. *Proteins Struct Funct Bioinf* 2013;81:119–31.

Chaudhury S, Lyskov S, Gray JJ. PyRosetta: a script-based interface for implementing molecular modeling algorithms using rosetta. *Bioinformatics* 2010;26:689–91.

Cretin G, Galochkina T, de Brevern AG et al. Pythia: deep learning approach for local protein conformation prediction. *IJMS* 2021;22: 8831.

Falcon W. The PyTorch lightning team. Pytorch Light 2019;3:6.

Fuchs F, Worrall D, Fischer V et al. SE(3)-transformers: 3D rotation-translation equivariant attention networks. *Adv Neural Inf Process Syst* 2020;33:1970–81.

Ganea O-E, Huang X, Bunne C et al. Independent SE(3)-equivariant models for end-to-end rigid protein docking. *arXiv*, arXiv:2111.07786, 2021, preprint: not peer reviewed.

Heo L, Feig M. PREFMD: a web server for protein structure refinement via molecular dynamics simulations. *Bioinformatics* 2018;34: 1063–5.

Heo L, Janson G, Feig M. Physics-based protein structure refinement in the era of artificial intelligence. *Proteins Struct Funct Bioinf* 2021;89: 1870–87.

Heo L, Park H, Seok C. GalaxyRefine: protein structure refinement driven by side-chain repacking. *Nucleic Acids Res* 2013;41: W384–88.

Hirayama N, Park H, Baek M et al. Improved protein structure refinement guided by deep learning based accuracy estimation. *Nat Commun* 2021;12:1340.

Jamasp AR, Lió P, Blundell TL. Graphein—a python library for geometric deep learning and network analysis on protein structures. *bioRxiv*, 2020, preprint: not peer reviewed.

Jing X, Xu J. Fast and effective protein model refinement using deep graph neural networks. *Nat Comput Sci* 2021;1:462–9.

Jones DT, Thornton JM. The impact of AlphaFold2 one year on. *Nat Methods* 2022;19:15–20.

Jumper J, Evans R, Pritzel A et al. Highly accurate protein structure prediction with AlphaFold. *Nature* 2021;596:583–9.

Lee GR, Won J, Heo L et al. GalaxyRefine2: simultaneous refinement of inaccurate local regions and overall protein structure. *Nucleic Acids Res* 2019;47:W451–5.

Mariani V, Biasini M, Barbato A et al. IDDT: a local superposition-free score for comparing protein structures and models using distance difference tests. *Bioinformatics* 2013;29:2722–8.

Mirjalili V, Noyes K, Feig M. Physics-based protein structure refinement through multiple molecular dynamics trajectories and structure averaging. *Proteins Struct Funct Bioinf* 2014;82:196–207.

Pearce R, Zhang Y. Toward the solution of the protein structure prediction problem. *J Biol Chem* 2021;297:100870.

Salomon-Ferrer R, Case DA, Walker RC. An overview of the amber biomolecular simulation package. *WIREs Comput Mol Sci* 2013;3:198–210.

Senior AW, Evans R, Jumper J et al. Improved protein structure prediction using potentials from deep learning. *Nature* 2020;577:706–10.

Simpkin AJ, Sánchez Rodriguez F, Mesdaghi S et al. Evaluation of model refinement in CASP14. *Proteins Struct Funct Bioinf* 2021;89:1852–69.

Steinegger M, Söding J. MMseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nat Biotechnol* 2017;35:1026–8.

Thomas N, Smidt T, Kearnes S et al. Tensor field networks: Rotation- and translation-equivariant neural networks for 3d point clouds. *arXiv*, arXiv:1802.08219, 2018, preprint: not peer reviewed.

Tunyasuvunakool K, Adler J, Wu Z et al. Highly accurate protein structure prediction for the human proteome. *Nature* 2021;596:590–6.

Wang M, Zheng D, Ye Z et al. Deep graph library: A graph-centric, highly-performant package for graph neural networks. *arXiv*, arXiv:1909.01315, 2019, preprint: not peer reviewed.

Williams CJ, Headd JJ, Moriarty NW et al. MolProbity: more and better reference data for improved all-atom structure validation. *Protein Sci* 2018;27:293–315.

Xu D, Zhang Y. Improving the physical realism and structural accuracy of protein models by a two-step atomic-level energy minimization. *Biophys J* 2011;101:2525–34.

Yang J, Anishchenko I, Park H et al. Improved protein structure prediction using predicted interresidue orientations. *Proc Natl Acad Sci USA* 2020;117:1496–503.

Zemla A. LGA: a method for finding 3D similarities in protein structures. *Nucleic Acids Res* 2003;31:3370–4.