Improved protein structure refinement guided by deep learning based accuracy estimation

Naozumi Hiranuma[1,2,a], Hahnbeom Park[1,a], Ivan Anishchanka[1], Minkyung Baek[1], and David Baker[1,3,*].

1 - Department of Biochemistry and Institute for Protein Design, University of Washington, WA, USA

2 - Paul G. Allen School of Computer Science & Engineering, University of Washington, WA, USA

3 - Howard Hughes Medical Institute, University of Washington, WA, USA

[a] N.H and H.P contributed equally to this work.

Correspondence to: *dabaker@uw.edu

Abstract

We develop a deep learning framework (DeepAccNet) that estimates per-residue accuracy and residue-residue distance signed error in protein models and uses these predictions to guide Rosetta protein structure refinement. The network uses 3D convolutions to evaluate local atomic environments followed by 2D convolutions to provide their global contexts. The network was trained on approximately 1 million alternative local energy minima for 7,510 different proteins exhibiting a wide diversity of errors, and outperforms other methods that similarly predict the accuracy of protein structure models without template or evolutionary information. Overall accuracy predictions for X-ray and cryoEM structures in the PDB correlate with resolution, and the network should be broadly useful for assessing accuracy of both predicted structure models and experimentally determined structures, and identifying specific regions likely to be in error. Guiding protein structure refinement by incorporation of the accuracy predictions at multiple stages in the Rosetta refinement protocol led to improvements in model quality in 63 out of 73 test cases, illustrating how deep learning can improve search for global energy minima.

Significance Statement

We develop a deep learning method to predict the accuracy of protein structure models, and use the method to improve protein structure refinement. Benchmark tests show that both the accuracy prediction method and the protein structure refinement method improve on previously described approaches.

Introduction

Distance prediction through deep learning on amino acid co-evolution data has considerably advanced protein structure prediction (1–3). However, in most cases the predicted structures still deviate considerably from the actual structure (4). The protein structure refinement challenge is to increase the accuracy of such starting models. To date the most successful approaches have been with physically based methods that involve large scale search for low

1

energy structures, for example with Rosetta (5) and/or molecular dynamics (6). This is because any available homology and co-evolutionary information is typically already used in the generation of the starting models.

The major challenge in refinement is sampling; the space of possible structures that must be searched through even in the vicinity of a starting model is extremely large (5, 7). If it were possible to accurately identify what parts of an input protein model were most likely to be in error, and how these regions should be altered, it should be possible to considerably improve the search through structure space and hence the overall refinement process. Many methods for estimation of model accuracy (EMA) have been described, including approaches based on deep-learning such as ProQ3D (based on per-residue Rosetta energy terms and multiple sequence alignments with multi-layer perceptrons (8)), and Ornate (based on 3D voxel atomic representations with 3d convolutional networks (9)). Non-deep-learning methods such as Voro-MQA compare a Voronoi tessellation representation of atomic interactions against pre-collected statistics (10). These methods focus on predicting per-residue accuracy. Few studies have sought to guide refinement using network based accuracy predictions (11); the most successful refinement protocols in the recent blind 13th Critical Assessment of Structure Prediction (CASP13) test either utilized very simple ensemble-based error estimations (5) or none at all (12). This is likely because of the low specificity of most current accuracy prediction methods, which predict which residues are likely to be inaccurately modeled, but not how they should be moved, and hence are less useful for guiding search.

Results

We set out to develop a deep-learning based framework (DeepAccNet) that estimates the signed error in every residue-residue distance along with the local residue contact error, and to use this estimation to guide Rosetta based protein structure refinement. Our approach is schematically outlined in Figure 1.

Development of improved model accuracy predictor

We first sought to develop model accuracy predictors that provide both global and local information for guiding structure refinement. We developed network architectures that make the following three types of predictions given a protein structure model: local measures of structure accuracy measured by per residue C_β local distance difference test (l-DDT) scores (13), a native C_β contact map thresholded at 15 Å (referred to as *mask*), and per residue-pair distributions of signed C_β-C_β distance error against corresponding native structures (referred to as *estograms*; histogram of errors); C_α is taken for GLY. Rather than predicting single error values for each pair of positions, we instead predict histograms of errors (analogous to the distance histograms employed in the structure prediction networks of (1–3)), which provide more detailed information about the distributions of possible structures and better represent the uncertainties inherent to error prediction. Networks were trained on alternative structures ("decoys") with model quality ranging from 50% to 90% in GDT-TS (global distance

test - tertiary structure) (14) generated by homology modeling (15), trRosetta (1), and native structure perturbation (see Methods). ~150 decoy structures were generated for each of 7,510 X-ray crystal structures with resolution better than 2.5 Å lacking extensive crystal contacts and having sequence identity less than 40% to any of 73 refinement benchmark set proteins (see below). Of the approximately one million decoys, those for 297 and 292 of the 7,510 proteins were held out for validation and testing, respectively. More details of the training/test set and decoy structure generation can be found in Methods.

The predictions are based on 1D, 2D, and 3D features that reflect accuracy at different levels. Defects in high resolution atomic packing are captured by 3D convolution operations performed on 3D atomic grids around each residue defined in a rotationally invariant local frame, similar to the Ornate method (9). 2D features are defined for all residue pairs, and they include Rosetta inter-residue interaction terms, which further report on the details of the interatomic interactions, while residue-residue distance and angular orientation features provide lower resolution structural information. At the 1D per residue level, the features are the amino acid sequence, backbone torsion angles, and the Rosetta intra-residue energy terms (see Methods for details).

We implemented a deep neural network, DeepAccNet, that incorporates these 1D, 2D, and 3D features (Figure 1A). The networks first perform a series of 3D convolution operations on local atomic grids in coordinate frames centered on each residue. These convolutions generate features describing the local 3D environments of each of the N residues in the protein. These, together with additional residue level 1D input features (local torsional angles and individual residue energies), are combined with the 2D residue-residue input features (distance and inter-residue energies) by tiling (so that associated with each pair of residues there are both the input 2D features for that pair and the 1D features for both individual residues), and the resulting combined 2D feature description is input to a series of 2D convolutional layers using the ResNet architecture (16). A notable advantage of our approach of tying together local 3D residue based atomic coordinate frames through a 2D distance map is the ability to incorporate full atomic coordinate information in a rotationally invariant way; in contrast a Cartesian representation of the full atomic coordinates would change upon rotation, substantially complicating network for both training and its use. Details of the network architecture, feature generation, and training processes are found in Methods.

Figure 2 shows examples of the predictions of DeepAccNet on two randomly selected decoy structures for each of three target proteins (3lhnA, 4gmqA, and 3hixA) not included in training. In each case, the network generates different signed residue-residue distance error maps for the two decoys that qualitatively resemble the actual patterns of the structural errors (rows of Figure 2). The network also accurately predicts the variations in per residue model accuracy (l-DDT scores) for the different decoys. The two samples from 4gmqA (second row) are closer to the native structure than for the other two proteins, and the network correctly predicts the location of the smaller distance errors and l-DDT scores closer to 1. Overall,

while the detailed predictions are not pixel-perfect, they provide considerable information on what parts of the structure need to move and in what ways to guide refinement.

We compared the performance of the full-feature network to that of a baseline network trained only on residue-residue distances. The performance of the network with the full set of features is considerably better on average for almost all the test set proteins (Figure S1A; Figure 3); it outperforms the baseline model in predicting estograms for residue pairs across different sequence separations and input distances (Figure S1B), and overall and per-residue accuracy particularly for quite inaccurate models and residues (Figure S1CD). For both networks trained on the full set of features and only distance-based features, l-DDT score prediction does not decline substantially with increasing size (Spearman correlation coefficient, or Spearman-r, of -0.11 with p-value < 0.05), but estogram prediction performance clearly declines for larger proteins (Spearman-r of 0.58 with p-value < 0.00001) (Figure S1E) -- for larger proteins with more interactions over long distances, estimating the direction and magnitude of errors is a much harder task while since l-DDT scores only consider local changes at short distances, they degrade less with increasing size.

In addition to distance map features, DeepAccNet takes as input a) amino acid identities and properties, b) local atomic 3D environments for each residue, c) backbone torsion angles, d) residue-residue orientations, e) Rosetta energy terms, and f) secondary structure information. To investigate the contributions of each of these features to network performance, we combined each with distance maps one at a time during training and evaluated performance through estogram cross-entropy loss and l-DDT score mean squared error on test sets (Figure 3A, Table S1). The largest single contribution was from the 3D convolution-based features followed by Rosetta energy terms and amino acid-related features. The amino acid sequence based features contribute more to performance on positions with low true per-residue l-DDT scores but not on ones with high true per-residue l-DDT scores.

An effective accuracy prediction method should be useful for evaluating and identifying potential errors in experimentally determined structures as well as computational models. We investigated the performance of the network on experimental structures determined by X-ray crystallography, nuclear magnetic resonance spectroscopy (NMR), and electron microscopy (EM) that were not included in the training set (details of the dataset can be found in Methods). The predicted C_β l-DDT values are close to 1.0 for high resolution crystal structures, as expected for nearly error free protein structures, and decreases for lower resolution structures (Figure 3C, left panel). A similar correlation between predicted accuracy and resolution holds for X-ray structures of membrane proteins (Figure3C, middle panel; Spearman-r 0.62 with p-value < 0.0001) and cryoEM structures (Figure 3C, right panel; Spearman-r 0.86 with p-value < 0.0001). A list of X-ray structures with low predicted l-DDT despite their high experimental resolution is provided in Table S5. Many of these are heme-proteins; as the network does not consider bound ligands, the regions surrounding them are detected as atypical for folded proteins, suggesting that the network may also be useful for predicting cofactor binding and other functional sites from apo-structures. NMR

structures have lower predicted accuracies than high resolution crystal structures (Figure 3D, right; Figure S3), which is not surprising given they represent solution averages rather than crystalline states.

We compared DeepAccNet to other accuracy estimators that take in a single model structure as an input and do not rely on multiple sequence alignment information (Figure 3B). As is clear from recent CASP experiments, co-evolution information derived from multiple sequence alignments provides detailed structure information; we did not include this as an input to our network for two reasons: first, all available homology and co-evolutionary information is typically already used in generating the input models for protein structure refinement and second, in applications such as *de novo* protein design model evaluation, no evolutionary multiple sequence alignment information exists. We compared the performance of DeepAccNet on the CASP13 EMA data (76 targets with approximately 150 decoy models each) to three models that similarly estimate error from a single structure model without co-evolution or template information, VoroMQA (10), Ornate (group name 3DCNN) (9), and a method from Lamoureux Lab (17)); the latter two use 3D convolutions similar to those used in our single residue environment feature calculations. We calculated (i) the Spearman-r of predicted and actual global l-DDT scores per target protein and (ii) area under receiver operator characteristic (ROC) curve for predicting mis-modeled residues per sample (C_β l-DDT< 0.6) which assesses global local model accuracy estimation respectively. DeepAccNet outperformed the other methods (Figure 3B right) according to both metrics. While this improved performance is very encouraging, it must be noted that our predictions are made after rather than before CASP13 data release so the comparison is not entirely fair: future blind accuracy prediction experiments will be necessary to compare methods on an absolutely even footing. As a step in this direction, we tested performance on structures released from the PDB after our network architecture was finalized that were collected in the CAMEO (Continuous Automated Model EvaluatiOn) (18) experiment between 2/22/2020 to 5/16/2020. We observed consistent improvements in both global (entire model) and local (per residue) accuracy prediction compared to other methods that do not use evolutionary information, -- namely, VoroMQA (10), QMean3 (19), and Equant 2 (20) (Figure 3B left; Comparison to methods that also use evolutionary information are provided in Figure S2). We could not compare signed residue-pair distance error prediction because this is not predicted by the other methods.

Guiding search in protein structure refinement using the accuracy predictor

We next experimented with incorporation of the network accuracy predictions into the Rosetta refinement protocol (5, 21), which was one of the top methods tested in CASP13 (22). Rosetta high resolution refinement starts with a single model, and in a first diversification stage explores the energy landscape around it using a set of sampling operators, and then in a subsequent iterative intensification stage hones in on the lowest energy regions of the space. Search is controlled by an evolutionary algorithm which

maintains a diverse but low energy pool through many iterations/generations. With improvements in the Rosetta energy function in the last several years (23, 24), the bottleneck to improving refinement has largely become sampling close to the correct structure. The original protocol utilized model consensus-based accuracy estimations (i.e. regional accuracy estimated as inverse of fluctuation within an ensemble of structures sampled around the input model) to keep search focused in the relevant region of the space -- these have the obvious downside of limiting exploration in regions which need to change substantially from the input model but are located in deep false local energy minima.

To guide search, predicted estograms and l-DDT scores were incorporated into the Rosetta refinement protocol at three levels (details in Methods). First, the per-residue l-DDT predictions were used to decide which regions to intensively sample or to recombine with other models. Second, the estograms were converted to residue-residue interaction potentials with weight for each pair defined by a function of its estogram prediction confidence, and these potentials were added to the Rosetta energy function as restraints to guide sampling. Third, global l-DDT prediction was used as the objective function during the selection stages of the evolutionary algorithm and to control the model diversity in the pool during iteration.

To benchmark the accuracy prediction guided refinement protocol, 73 protein refinement targets were collected from previous studies (5, 21). The starting structures were generally the best models available from automated structure prediction methods. A separate 7 targets from Park et al (5, 21) were used to tune the restraint parameters and were excluded from the benchmarking in this study.

We found that network-based accuracy prediction consistently improves refinement across the benchmark examples. In Figure 4, refinement guided by the accuracy predictions is compared to our previous protocol in which simpler non-deep learning accuracy estimation was used. Refinement of many proteins in the benchmark set was previously quite challenging due to their size (21); however, with the new protocol, consistent improvements are observed over the starting models regardless of protein size (Figure 4A, the l-DDT improve by 10% on average) and over the models produced with our previous unguided search (Figure 4B; the l-DDT improve by 4% on average). The number of targets with l-DDT improvements of greater than 10% increases from 29% to 48% using DeepAccNet to guide refinement. These improvements are quite notable given how challenging the protein structure refinement problem is (a comparison to other best predictors on the latest CASP targets is shown in Figure S7); for reference best improvements between successive biannual CASP challenges are typically < 2% (22). Tracing back through the refinement trajectory reveals that the progress in both predicted and actual model quality occurs gradually through the stages and are quite well correlated to each other (Figure 4C). Predictions of more detailed per residue model quality also agree well with their actual values (Figure 4E). Improvements were found regardless of the experimental techniques used for structure determination, although the contribution from the network was more pronounced for

those proteins determined by X-ray crystallography (5.0%, compared to 2.4% for those determined by NMR).

The model accuracy improvements occur across a broad range of protein sizes, starting model qualities, and types of errors. Refinement improved models across various secondary structures to similar extents, and corrected secondary structures originally modeled incorrectly, increasing model secondary structure accuracy by almost 10% based on an 8-state definition (25) (Figure 4E). As shown in Figure 4F, improvements involve identification and modifications of erroneous regions when the overall structure is correct (TR776) as well as overall concerted movements when the core part of the model is somewhat inaccurate (5m1mA). The accuracy prediction network promotes this overall improvement in two ways: first, it provides a more accurate estimation of unreliable distance pairs and regions at every iteration of refinement for every model on which sampling can be focused, and second, it provides a means to effectively constrain the search space in the already accurately modeled regions through residue-residue pair restraints -- this is particularly important for refinement of large proteins. The network enables the refinement protocol to adjust how widely to search on a case by case basis; this is an advantage over most previous refinement approaches where search has generally been either too conservative or too aggressive (26).

DeepAccNet is available at https://github.com/hiranumn/DeepAccNet.

Discussion

Representations of the input data are critical for the success of deep learning approaches. In the case of proteins, the most complete description is the full Cartesian coordinates of all of the atoms, but these are transformed by rotation and hence not optimal for predicting rotationally invariant quantities such as error metrics. Hence most previous machine learning based accuracy prediction methods have not used the full atomic coordinates (8, 10, 27). The previous described Ornate method does use atomic coordinates to predict accuracy, and solves the rotation dependence by setting up local reference frames for each residue. As in the Ornate method, DeepAccNet carries out 3D convolutions over atomic coordinates in residue centered frames, but we go beyond Ornate in then integrating together this detailed residue information along with additional individual residue and residue-residue level geometric and energetic information by 2D convolutions over the full N x N residue-residue distance map.

Evaluation of performance on CASP and CAMEO datasets shows that, despite not using homology or multiple sequence alignment derived co-evolution information, our DeepAccNet makes state-of-the-art accuracy predictions, and it is the first to our knowledge to predict signed distance errors for protein structure refinement. Model quality estimations on X-ray crystal structures correlate with resolution, and the network should also be useful in identifying errors in experimentally determined structures (Figure 3C). DeepAccNet performs

well on both cryoEM and membrane protein structures, and it could be particularly useful for low resolution structure determination and modeling of currently unsolved membrane proteins (Figure 3C). We also anticipate that the network will be useful in evaluating protein design models.

Guiding search using the network predictions improved Rosetta protein structure refinement over a wide range of protein sizes and starting model qualities (Figure 4). However, there is still considerable room for improvement in the combined method. To more effectively use the information in the accuracy predictions it will be useful to explore sampling strategies which can better utilize the network predictions and more frequent communication between Rosetta modeling and the accuracy prediction network -- the network is fast enough to evaluate the accuracy of many models more frequently. It is clear that there is also considerably more to explore in using deep learning to guide refinement. For example, selection of which of the current sampling operators to use in a given situation, and the development of new sampling operators using generative models such as sampling missing regions by inpainting. More generally, reinforcement learning approaches should help identify more sophisticated iterative search strategies.

Methods

Data preparation.
Training and test sets for protein model structures (often called decoys) are generated to most resemble starting models of real-case refinement problems. We reasoned that a relevant decoy structure should meet the following conditions: i) has template(s) not too far or close in sequence space; ii) does not have strong contacts to other protein chains, iii) should contain minimal fluctuating (i.e. missing density) regions. To this end, we picked a set of crystal structures from the PISCES server (deposited by Nov 26 2018) containing 20,399 PDB entries with maximum sequence redundancy of 40% and minimum resolution of 2.5 Å. We further trimmed the list to 8,718 chains by limiting their size to 50-300 residues and requiring that proteins are either monomeric or have minimal interaction with other chains (weaker than 1 kcal/mol per residue in Rosetta energy). HHsearch (28) was used to search for templates; 50 templates with highest HHsearch probability, sequence identity of at most 40% and sequence coverage of at least 50% are selected for model generation.

Decoy structures are generated using three methods: comparative modeling, native structure perturbation, and deep learning guided folding. Comparative modeling and native structure perturbation are done using RosettaCM (15). For comparative modeling of each protein chain we repeated RosettaCM 500 times in total, every time randomly selecting a single template from the list. In order to increase the coverage of decoy structures at mid-to-high accuracy regime for targets lacking templates with GDT-TS > 50, 500 models are further generated providing a single template and 40% trimmed native structure as templates. Sampled decoy set for a protein chain is included in training/test data only if the total number of decoys at medium accuracy (GDT-TS to native ranging from 50 to 90) is larger than 50. Maximum 15

lowest scoring decoys at each GDT-TS bin (ranging from 50 to 90 with bin size 10) are collected, then the rest with lowest energy values are filled so as to make the set contain approximately 90 decoys. Native structures are perturbed to generate high-accuracy decoys. 30 models were generated by RosettaCM either by i) combining a partial model of native structure with high-accuracy templates (GDT-TS > 90) or ii) inserting fragments at random positions of the native structure. Deep learning guided folding is done using trRosetta (1). For each protein, 5 subsampled multiple sequence alignments (MSAs) are generated with various depths (i.e. number of sequences in MSA) ranging from 1 to maximum available. The standard trRosetta modeling is run 45 times for each of the subsampled MSAs. The final decoy set collected, consisting of about 150 structures (90 from comparative modeling, 30 from native perturbation, and 30 from deep learning guided folding) per each of 7,510 protein chains (6,921, 297, 292 for training, validation and test datasets), are thoroughly relaxed by Rosetta dual-relax (29) prior to the usage. The distribution of the starting l-DDT values of the test proteins are shown in Figure S6.

Model architectures and input features.
In our framework, convolution operations are performed in several dimensions, and different classes of features come in at different entry points of the network (Figure 1). Here, we briefly describe the network architecture as well as classes of features. More detailed descriptions about the features and model parameters are listed in Table S2 and S3.

The first set of input features to the network are voxelized Cartesian coordinates of atoms per residue, generated in a manner similar to Ornate (9). Voxelization is performed individually for every residue in the corresponding local coordinate frame defined by backbone N,Cα,C atoms. Such representation is translationally and rotationally invariant because projections onto local frames are independent of the global position of the protein structure in 3D space. The second set of inputs are per residue 1D features (e.g., amino acid sequence and properties, backbone angles, Rosetta intra-residue energy terms, and secondary structures) and per residue pair 2D features (e.g. residue-residue distances and orientations, Rosetta inter-residue energy terms).

In the first part of the neural network, the voxelized atomic coordinates go through a series of 3D convolution layers whose parameters are shared across residues. The resulting output tensor is flattened so that it becomes a 1D vector per residue, which is concatenated to other 1D features. The second part of the network matches the dimensionality of the features and performs a series of 2D convolution operations. Let us now denote that there are n residues, f_1 1D features, and f_2 2D features. Then, the input matrix of the 1D features M_1 has the shape of n by f_1, and the input matrix of the 2D features M_2 has the shape of n by n by f_2. We tile M_1 in the first and second axis of M_2, concatenating them to produce a feature matrix of size n by n by $2f_1+f_2$. The third axis of the resulting matrix represents vectors of size $2f_1+f_2$, which contain the 2d features and 1D features of i-th and j-th residues. This data representation allows us to convolve over both backbone chain and pairwise interactions.

The concatenated feature matrix goes through a residual network with 20 residual blocks, with cycling dilation rates of 1, 2, 4, and 8 (see Tables S2). Then, the network branches off to two arms of 4 residual blocks. These arms separately predict distributions of C_β distance errors for all pairs of residues (referred to as estograms) and whether a particular residue pair is within 15 Å in a corresponding native structure (referred to as masks). Estograms are defined over categorical distributions with 15 binned distance ranges; the boundary of bins are at -20.0Å, -15.0Å, -10.0Å, -4.0Å, -2.0Å, -1.0Å, -0.5Å, 0.5Å, 1.0Å, 2.0Å, 4.0Å, 10.0Å, 15.0Å, 20.0Å.

In standard calculation of a C_β l-DDT score of i-th residue of a model structure, all pairs of C_β atoms that include the *i th* residue and are less than 15Å in a reference structure are examined. 0.5Å, 1.0Å, 2.0Å, and 4.0Å cutoffs are used to determine the fractions of preserved C_β distances across the set of pairs. The final C_β l-DDT score is calculated by computing the arithmetic mean of all fractional values (13).

In our setup, we obviously do not have access to reference native structures. Instead, a C_β l-DDT score of i-th residue is predicted by combining the probabilistic predictions of estograms and masks as follows:

$$per\_residue\_LDDT = 0.25 * (\bar{p}_0 + \bar{p}_1 + \bar{p}_2 + \bar{p}_3)/\bar{p}_4$$

$\bar{p}_0$ is the mean of probability that the magnitude of C_β distance errors are less than 0.5Å, across all residue pairs that have i-th residue involved and predicted to be less than 15Å in its corresponding native structure. The former C_β distance errors are obtained from estogram predictions and the latter native distance information are directly obtained from mask predictions. $\bar{p}_1...\bar{p}_3$ are similar quantities with different cutoffs for errors; 1.0Å, 2.0Å, and 4.0Å, respectively. $\bar{p}_4$ is the mean probability that native distance is within 15Å and it is again directly obtained from mask predictions.

The network was trained to minimize categorical cross-entropy between true and predicted estograms and masks. Additionally, as noted, we calculated C_β l-DDT scores based on estograms and masks, and we used a small amount of mean squared loss between predicted and true scores as an auxiliary loss. The following weights on the three loss terms are used.

$$global\_loss = estogram\_loss + 10.0 * LDDT\_loss + 0.25 * mask\_loss$$

The weights are tuned so that the highest loss generally comes from $estogram\_loss$ since estograms are the richest source of information for the downstream refinement tasks. At each step of training, we selected a single decoy from decoy sets of a randomly chosen training protein without replacement. The decoy sets include native structures, in which case the target estograms ask networks to not modify any distance pairs. An epoch consists of a full

10

cycle through training proteins, and the training processes usually converge after 100 epochs (Figure S5). We used an ADAM optimizer with a learning rate of 0.0005 and decay rate of 0.98 per epoch. Training and evaluation of the networks was performed on RTX2080 gpus.

Analyzing the importance of features.
Feature importance analysis was conducted to understand and quantify the contributions from different classes of features to accurately predicting accuracy of model structures. To do this, we combined each feature class with a distance map one at a time during training (or removed them in one particular case) and analyzed loss of predictions on a held-out test protein set. In addition to the network with full features and its variant only with distance map features, we trained 7 types of networks: i) with amino acid identities and properties, ii) with local atomic environments scanned with 3D convolution, iii) with backbone angles, iv) with residue-residue orientations, v) with Rosetta energy terms, vi) with secondary structure information, and vii) without 3D convolution from the full feature network. We trained four replicas of each type of network (including the network with full features and only with distance map features) in order to reduce noise introduced by training stochasticity. We used 3,500 out of the full 7,510 proteins to speed up the training process. Models from the best performing epochs on validation sets are selected and their predictions are averaged to produce ensembled predictions.

We are aware that more sophisticated feature attribution methods for deep networks exist (30); however, these methods attribute importance scores to features per output per sample. Since we have approximately quarter million outputs and near million inputs with a typical 150 residue protein, these methods were not computationally feasible and tractable to analyze.

Comparing with other model accuracy estimation methods.
For the CASP13 datasets, we downloaded submissions of QA139_2 (ProQ3D), QA360_2 (ProQ3D-lDDT (8)), QA187_2 (ProQ3 (27)), QA067_2 (LamoureuxLab (17)), QA030_2 (VoroMQA-B (10)), QA359_2 (Ornate, group name 3DCNN (9)) for the accuracy estimation category. The former four methods submitted their predictions for 76 common targets, whereas the last method, Ornate, only submitted for 55 targets. Thus, we decided to do analysis on the 76 common targets for all methods except for Ornate, which was only evaluated on 55 targets. Evaluation was performed in two metrics; i) Spearman-r of predicted quality scores across decoys of each target, and ii) area under ROC curve for predicting mis-modeled residues of each sample (C_β l-DDT< 0.6). Samples whose residues are all below or above 0.6 C_β l-DDT are omitted. For assessing the performance of methods other than ours, their submitted estimations of global quality score was evaluated against the true full-atom global l-DDT scores. To assess the performance of our model, we first dual-relaxed (29) all models and ran DeepAccNet to estimate the global model qualities. We then evaluated the predicted scores against the same true scores.

For the CAMEO datasets, we downloaded the QA datasets registered between 2/22/2020 to 5/16/2020. This corresponds to 206 targets with approximately 10 modeled structures on average. We downloaded submissions of "Baseline potential", EQuant 2, ModFOLD4, ModFOLD7_LDDT, ProQ2, ProQ3, ProQ3D, ProQ3D_LDDT, QMEAN3, QMEANDisco3, VoroMQA_sw5, and VoroMQA_v2. Some methods did not submit their predictions for all samples, and those missing predictions are omitted from the analysis. Our network was evaluated on all samples using the same two metrics.

Visualizing predictions.

Figure 2 visualizes true and predicted estograms per pair of residues. The images are generated by calculating the expected values of estograms by taking weighted sums of central error values from all bins. For the two bins that encode for errors larger than 20.0 Å and smaller than -20.0 Å, we define the central distance at their boundaries of 20.0 Å and -20.0 Å.

Native structure dataset

Native structures that were not used for model training and validation, monomeric, larger than 40 residues, and smaller than 300 residues for the X-ray and NMR structures, and smaller than 600 residues for EM structures were downloaded from the PDB. For Figure 3C, samples with resolution larger than 4Å and 5Å are ignored for the X-ray and EM structures, respectively. The histograms in Figure 3D are using all samples. In total, 23,672 X-ray structures, 88 EM structures, and 2,154 NMR structures are in the histograms. For NMR structures, regions highly varying across the models were trimmed. Structures were discarded if the number of remaining residues after trimming was less than 40 residues or half of the original chain length.

Dataset for refinement runs.

We took 73 proteins and their starting models from our previous work (5) with a few modifications as described below. Of entire 84 targets used in our previous work, 7 small-sized targets (4zv5A, 5azxA, 5ghaE, 5i2qA, 5xgaA, TR569, T0743) are excluded from the benchmark set and were used for restraint parameter search. 8 additional targets (2n12A, 4idiA, 4z3uA, 5aozA, 5fidA, T0540, TR696, TR857) are excluded after more careful visual inspections as those had potential issues in their native structures (e.g. having contacts with ligands or other chains in crystal structures). 4 new targets were added from previous CASP refinement categories that were not included in the original set (TR747, TR750, TR776, TR884). Model accuracy is evaluated on a subset of ordered residues by trimming less confident residues according to the CASP standard evaluation criteria (22).

Refinement protocol.

Refinement protocol tested in this work inherits the framework from previous study (5). The overall architecture consists of two stages (Figure 1B): first initial model diversification stage, followed by iterative model intensification stages where a pool of structures is maintained during optimization by an evolutionary algorithm. At the diversification stage, following

accuracy estimation of the single starting model, two thousands of independent Rosetta modeling are attempted using RosettaCM (15). In the iterative annealing stage, series of accuracy estimation, new structure generation, and pool selection steps are repeated iteratively. At each iteration, 10 model structures are selected from the current pool, then individual accuracy predictions are made for each of 10 structures in order to guide the generation of 12 new model structures starting from each (total 120). New pool with size of 50 is selected among 50 previous pool members plus 120 newly generated ones with criteria of i) the highest global l-DDT estimated and ii) model diversity within the pool. This process is repeated for 50 iterations. At every 5-th iterations, a *recombination iteration* is called instead of a regular iteration where model structures are recombined with another member in the pool according to the residue l-DDT values predicted by the network (see below).

For modeling of a single structure at both diversification and intensification stages, first *unreliable regions* in the structure are estimated from accuracy prediction (see below). Structural information is removed in those regions and fully reconstructed from scratch. Fragment insertions are carried out in a coarse-grained broken-chain representation of the structure (15) focusing more on unreliable regions (5 times more frequently with respect to the rest part), followed by repeated side-chain rebuilding and minimization (29) in all-atom representation. Both coarse-grained and all-atom stage modeling are guided by *distance restraints* derived from accuracy predictions in addition to Rosetta energy. Details of unreliable region predictions, recombination iteration, and restraints are reported in the following sections.

Unreliable region prediction
Accuracy values predicted from the network are used to identify unreliable regions. We noticed that the l-DDT metric has a preference for helical regions (as local contacts are almost always correct). To fix this systematic bias, we exclude short sequence separation contacts in the contact mask that are within sequence separation of 11 to get corrected residue l-DDT values. Then these values are smoothed through a 9-residue-window uniform weight kernel. The residues at lowest accuracy are determined as unreliable regions. Two definitions of regions are made: in *static* definition, the accuracy threshold is varied until the fraction of unreliable regions lies between 10 to 20% of the entire structure. In *dynamic* definition, this range is defined as a function of predicted global accuracy (i.e. average residue-wise corrected accuracy): from f_{dyn} to f_{dyn} +10% with $f_{dyn} = 20 + 20*(0.55 - Q)/30$, where Q refers to predicted global accuracy. f_{dyn} is capped between 20 to 40%. In the diversification stage, one thousand models were generated for each definition of unreliable regions. Static definition is applied throughout the iterative stage.

Restraints
We classified residue pairs in three confidence levels: *high confidence*, *moderate confidence*, and *non-preserving*. Highly or moderately confident residue pairs stand for those whose distance should be fixed from the reference structure (i.e. starting structure) at different strengths; non-preserving pairs refer to the rest which can freely deviate.

Confident pairs are collected if C_β-C_β distance are not greater than 20Å and whose "probability with absolute estimated error ≤ 1Å", shortly P_{cen}, is above a certain threshold (e.g. 0.7). For those pairs, bounded functions are applied at coarse-grained modeling stage, and sum of sigmoid functions at all-atom modeling stage, minima centering at the original distance d_0 for both cases:

Bounded function:

$$f(d) = \frac{(d-(d_0+tol+s))}{s} + 1 \quad \text{for } d > d_0 + tol + s$$

$$f(d) = \left(\frac{d-(d_0+tol)}{s}\right)^2 \quad \text{for } d_0 + tol \le d \le d_0 + tol + s$$

$$f(d) = 0 \quad \text{for } |d - d_0| < tol$$

$$f(d) = \left(\frac{d-(d_0-tol)}{s}\right)^2 \quad \text{for } d_0 - tol - s \le d \le d_0 - tol$$

$$f(d) = \frac{(d-(d_0-tol-s))}{s} + 1 \quad \text{for } d < d_0 - tol - s$$

Sum of sigmoid function:

$$f(d) = w_{fa} * \left[\frac{-1}{1+exp(-5.0*(d-d_0+tol)/s)} + \frac{1}{(1+exp(-5.0*(d-d_0-tol)/s)} + 1 \right]$$

where s and *tol* stand for width and tolerance of the functions. Thresholds in P_{cen} values for highly confident pairs, P_{high}, and moderately confident pairs, $P_{moderate}$, are set at 0.8 and 0.7, with (s,tol) = (1.0,1.0) and (2.0,2.0), respectively, by analyzing the network test results shown in Figure S4. Restraint weight at all-atom stage modeling, w_{fa}, is set as 1.0. We noticed iterative refinement with these empirically determined parameters (w_{fa}, {P_{high}, $P_{moderate}$}) brought too conservative changes. We therefore ran another iterative refinement with a more aggressive parameter set (0.2,{0.8,0.9}) and chose the trajectory from whichever sampled a higher predicted global l-DDT.

For the rest non-preserving C_β-C_β pairs whose input distances are shorter than 40Å, error probability profiles (estograms) are converted into distance potentials by subtracting error bins from the original distances d_0 and taking log odds to convert probability into energy units. Instead of applying raw probabilities from the network, corrections are made against background probability collected from the statistics of the network's predictions over 20,000 decoy structures in the training set conditioning on sequence separation, original distances d_0, and predicted global model quality. The potential was applied in full form interpolated by spline function at the initial diversification stage, and was replaced by a simpler functional form in subsequent iterative process for efficiency:

$$f(d) = (d - 9) + 1 \quad \text{for } d > 9\,\text{Å}$$

$$f(d) = (d - 8)^2 \quad \text{for } 8 \le d \le 9\,\text{Å}$$

$$f(d) = 0 \quad \text{for } d < 8\,\text{Å}$$

14

for those pairs predicted from estogram as contacting within 10Å. Contacts are predicted when $P_{contact} > 0.8$, with $P_{contact} = sum(P_i)$ over i whose $d_0 + e_i < 10$Å and P_i stands for probability in estogram at bin *i*.

Recombination Iteration

At the recombination iteration, instead of running RosettaCM as the sampling operator, model structures are directly generated by recombining the coordinates from two models according to the predicted residue l-DDT profiles by the network. For a "seed" member, 4 "partners" are identified among the remaining 49 members in the pool that have the most complementarity to the seed in the predicted residue l-DDT profiles. All the members in the pool are recombined individually with their 4 partners, resulting in a total 200 new structural models. For each seed-partner combination, first, "complementary regions" are identified where the seed is inferior to the partner in terms of predicted l-DDT, then coordinates at the regions are substituted to those from the partner. Multiple discontinuous regions are allowed but the total coverage is restricted to a range between 20 to 50% of total residues. Next, Rosetta FastRelax (29) is run by imposing residue-pair restraints from estograms brought from either the partner or the seed interpolated into pair potentials (see above). Restraints from the partner are taken if any residue in the pair is included in complementary regions, and from the seed for the rest pairs. Recombination iterations are called at every 5 iterations to prevent over-convergence in the pool.

Final model selection

A model with the highest predicted global l-DDT is selected among 50 final pool members. Then a pool of structures similar to this structure (S-score (31) > 0.8) are collected from the entire iterative refinement trajectory, structurally averaged, and regularized in model geometry by running dual-relax (29) with strong backbone coordinate restraints with a harmonic constant of 10 kcal/mol Å^2, which was the identical post-processing procedure in our previous work (5). The final model refers to this structurally averaged and subsequently regularized structure. Structural averaging adds 1% l-DDT gain on average.

Implementation availability.

Code and accompanying scripts for the model accuracy predictors are implemented with Python 3 and Tensorflow 1.14 and made available at https://github.com/hiranumn/DeepAccNet.

Acknowledgement

Author contributions

N.H, H.P, and D.B. designed research; N.H., I.A., M.B. contributed in developing and analyzing the deep learning networks; H.P. contributed to application of the network on the refinement process; N.H., H.P., and D.B. wrote the manuscript.

References

1. J. Yang, *et al.*, Improved protein structure prediction using predicted interresidue orientations. *Proc. Natl. Acad. Sci. U. S. A.* 117, 1496–1503 (2020).

2. A. W. Senior, *et al.*, Improved protein structure prediction using potentials from deep learning. *Nature* 577, 706–710 (2020).

3. J. Xu, Distance-based protein folding powered by deep learning. *Proc. Natl. Acad. Sci. U. S. A.* 116, 16856–16865 (2019).

4. A. Kryshtafovych, T. Schwede, M. Topf, K. Fidelis, J. Moult, Critical assessment of methods of protein structure prediction (CASP)—Round XIII. *Proteins: Structure, Function, and Bioinformatics* 87, 1011–1020 (2019).

5. H. Park, *et al.*, High-accuracy refinement using Rosetta in CASP13. *Proteins: Structure, Function, and Bioinformatics* 87, 1276–1282 (2019).

6. L. Heo, M. Feig, Experimental accuracy in protein structure refinement via molecular dynamics simulations. *Proc. Natl. Acad. Sci. U. S. A.* 115, 13276–13281 (2018).

7. M. Feig, Computational protein structure refinement: almost there, yet still so far to go. *WIREs Comput Mol Sci* 7 (2017).

8. K. Uziela, D. Menéndez Hurtado, N. Shu, B. Wallner, A. Elofsson, ProQ3D: improved model quality assessments using deep learning. *Bioinformatics* 33, 1578–1580 (2017).

9. G. Pagès, B. Charmettant, S. Grudinin, Protein model quality assessment using 3D oriented convolutional neural networks. *Bioinformatics* 35, 3313–3319 (2019).

10. K. Olechnovič, Č. Venclovas, VoroMQA: Assessment of protein structure quality using interatomic contact areas. *Proteins* 85, 1131–1145 (2017).

11. D. Bhattacharya, refineD: improved protein structure refinement using machine learning based restrained relaxation. *Bioinformatics* 35, 3320–3328 (2019).

12. L. Heo, C. F. Arbour, M. Feig, Driven to near-experimental accuracy by refinement via molecular dynamics simulations. *Proteins: Structure, Function, and Bioinformatics* 87, 1263–1275 (2019).

13. V. Mariani, M. Biasini, A. Barbato, T. Schwede, lDDT: a local superposition-free score for comparing protein structures and models using distance difference tests. *Bioinformatics* 29, 2722–2728 (2013).

14. A. Zemla, LGA: A method for finding 3D similarities in protein structures. *Nucleic Acids Res.* 31, 3370–3374 (2003).

15. Y. Song, *et al.*, High-resolution comparative modeling with RosettaCM. *Structure* 21, 1735–1742 (2013).

16. K. He, X. Zhang, S. Ren, J. Sun, Deep Residual Learning for Image Recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016) https:/doi.org/10.1109/cvpr.2016.90.

17. G. Derevyanko, S. Grudinin, Y. Bengio, G. Lamoureux, Deep convolutional networks for quality assessment of protein folds. *Bioinformatics* 34, 4046–4053 (2018).

18. J. Haas, *et al.*, Continuous Automated Model EvaluatiOn (CAMEO) complementing the critical assessment of structure prediction in CASP12. *Proteins* 86 Suppl 1, 387–398 (2018).

19. P. Benkert, S. C. E. Tosatto, D. Schomburg, QMEAN: A comprehensive scoring function for model quality assessment. *Proteins: Structure, Function, and Bioinformatics* 71, 261–277 (2008).

20. S. Bittrich, F. Heinke, D. Labudde, eQuant - A Server for Fast Protein Model Quality Assessment by Integrating High-Dimensional Data and Machine Learning. *Communications in Computer and Information Science*, 419–433 (2016).

21. H. Park, S. Ovchinnikov, D. E. Kim, F. DiMaio, D. Baker, Protein homology model refinement by large-scale energy optimization. *Proc. Natl. Acad. Sci. U. S. A.* 115, 3054–3059 (2018).

22. R. J. Read, M. D. Sammito, A. Kryshtafovych, T. I. Croll, Evaluation of model refinement in CASP13. *Proteins: Structure, Function, and Bioinformatics* 87, 1249–1262 (2019).

23. H. Park, *et al.*, Simultaneous Optimization of Biomolecular Energy Functions on Features from Small Molecules and Macromolecules. *J. Chem. Theory Comput.* 12, 6201–6212 (2016).

24. R. F. Alford, *et al.*, The Rosetta All-Atom Energy Function for Macromolecular Modeling and Design. *J. Chem. Theory Comput.* 13, 3031–3048 (2017).

25. W. Kabsch, C. Sander, Dictionary of protein secondary structure: pattern recognition of hydrogen-bonded and geometrical features. *Biopolymers* 22, 2577–2637 (1983).

26. V. Modi, R. L. Dunbrack, Assessment of refinement of template-based models in CASP11. *Proteins: Structure, Function, and Bioinformatics* 84, 260–281 (2016).

27. K. Uziela, N. Shu, B. Wallner, A. Elofsson, ProQ3: Improved model quality assessments using Rosetta energy terms. *Scientific Reports* 6 (2016).

28. V. Mariani, F. Kiefer, T. Schmidt, J. Haas, T. Schwede, Assessment of template based protein structure predictions in CASP9. *Proteins* 79 Suppl 10, 37–58 (2011).

29. P. Conway, M. D. Tyka, F. DiMaio, D. E. Konerding, D. Baker, Relaxation of backbone bond geometry improves protein energy landscape modeling. *Protein Sci.* 23, 47–55 (2014).

30. Y. Sun, M. Sundararajan, Axiomatic attribution for multilinear functions. *Proceedings of the 12th ACM conference on Electronic commerce - EC '11* (2011) https:/doi.org/10.1145/1993574.1993601.

31. A. Ray, E. Lindahl, B. Wallner, Improved model quality assessment using ProQ2. *BMC Bioinformatics* 13, 224 (2012).

32. S. Henikoff, J. G. Henikoff, Amino acid substitution matrices from protein blocks. *Proc. Natl. Acad. Sci. U. S. A.* 89, 10915–10919 (1992).

33. J. Meiler, A. Zeidler, F. Schmaschke, M. Muller, Generation and evaluation of dimension-reduced amino acid parameter representations by artificial neural networks. *Journal of Molecular Modeling* 7, 360–369 (2001).
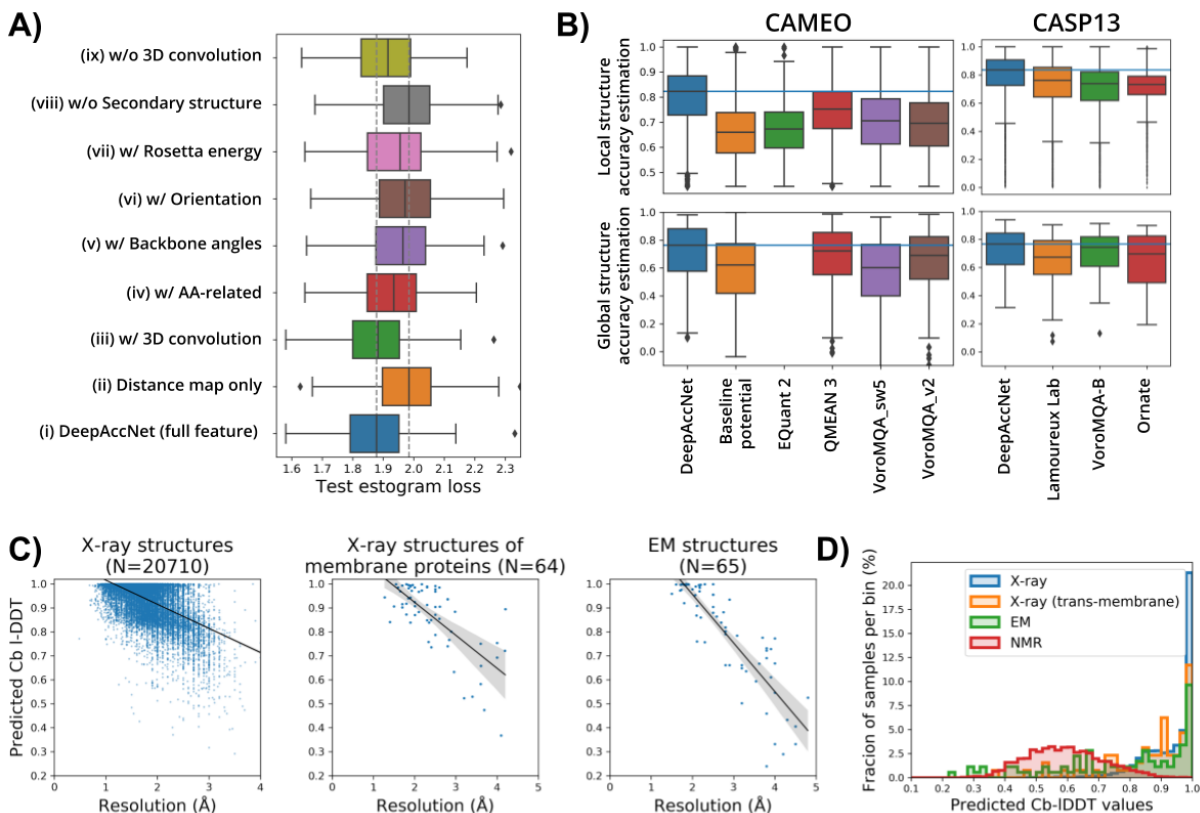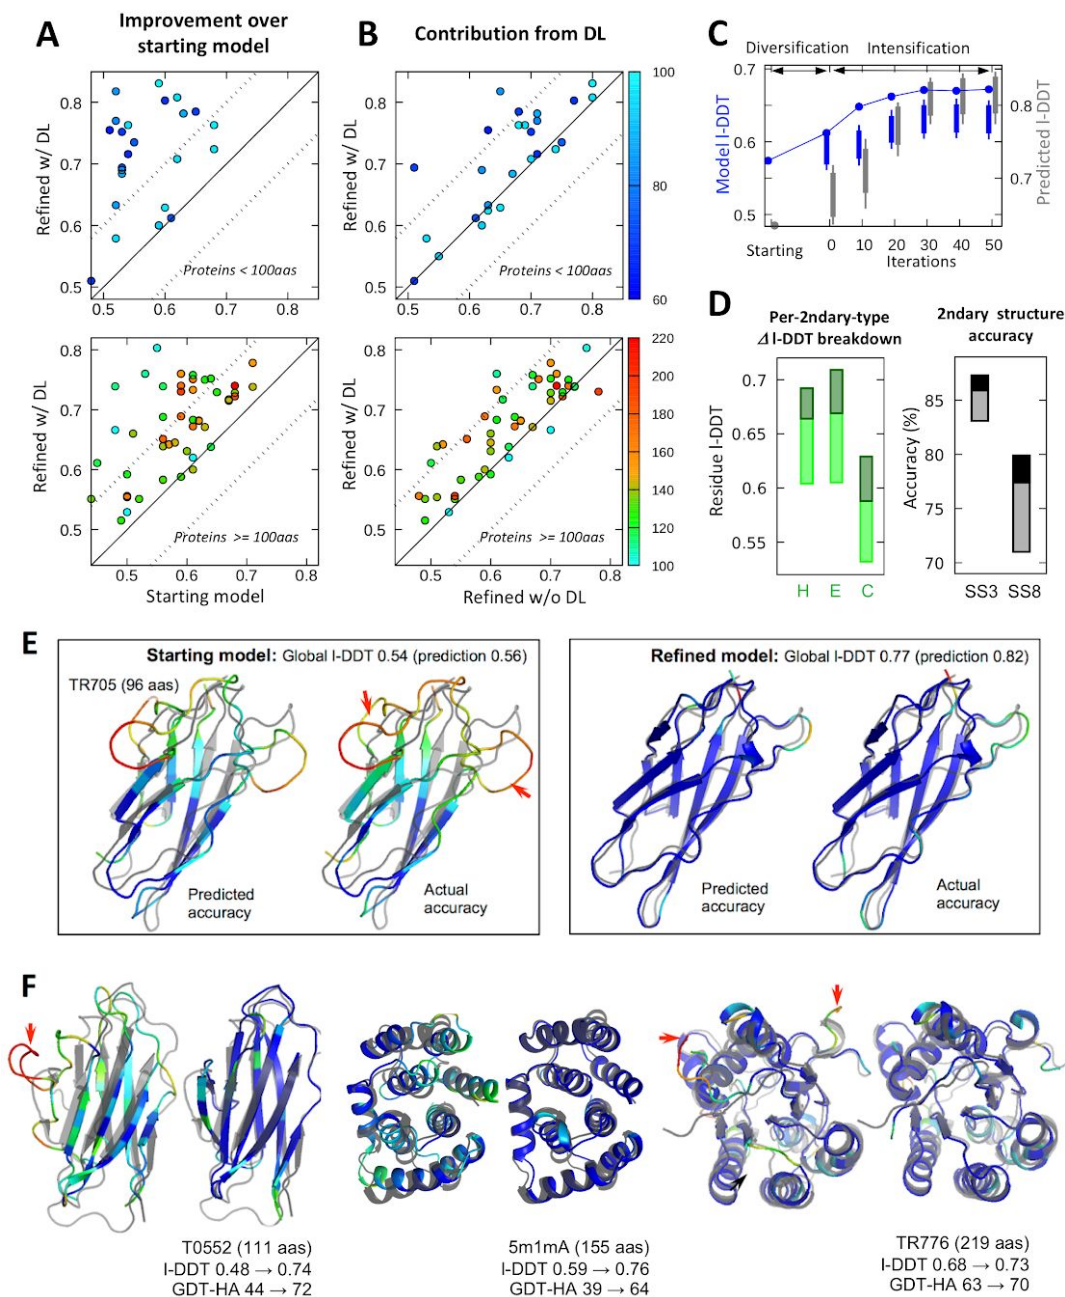
Figures
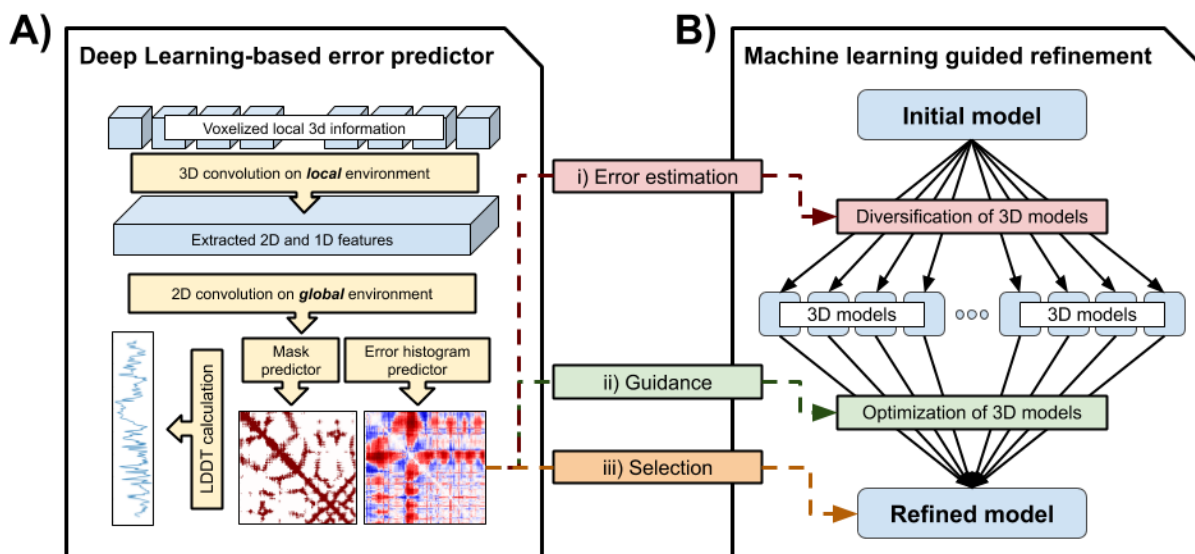
Figure 1: Approach overview. A) The deep learning network (DeepAccNet) consists of a series of 3D and 2D convolution operations. The networks are trained to predict i) the signed C_β-C_β distance error distribution for each residue pair (error histogram or *estogram* in short), ii) the native C_β contact map with a threshold of 15Å (referred to as mask), iii) the C_β l-DDT score per residue; C_α is taken for GLY. Input features to the network include: a) distance maps, b) amino acid identities and properties, c) local atomic environments scanned with 3D convolutions, d) backbone angles, e) residue angular orientations, f) Rosetta energy terms, and g) secondary structure information. Details of the network architecture and features are provided in Methods. B) The machine learning guided refinement protocol uses the trained neural networks in three ways; the estimated l-DDT scores are used to identify regions for more intensive sampling and model recombination, the estimated pairwise error distributions are used to guide diversification and optimization of structure(s), and finally the estimated global l-DDT score, which is mean of per-residue values, to select models during and at the end of the iterative refinement process.

18
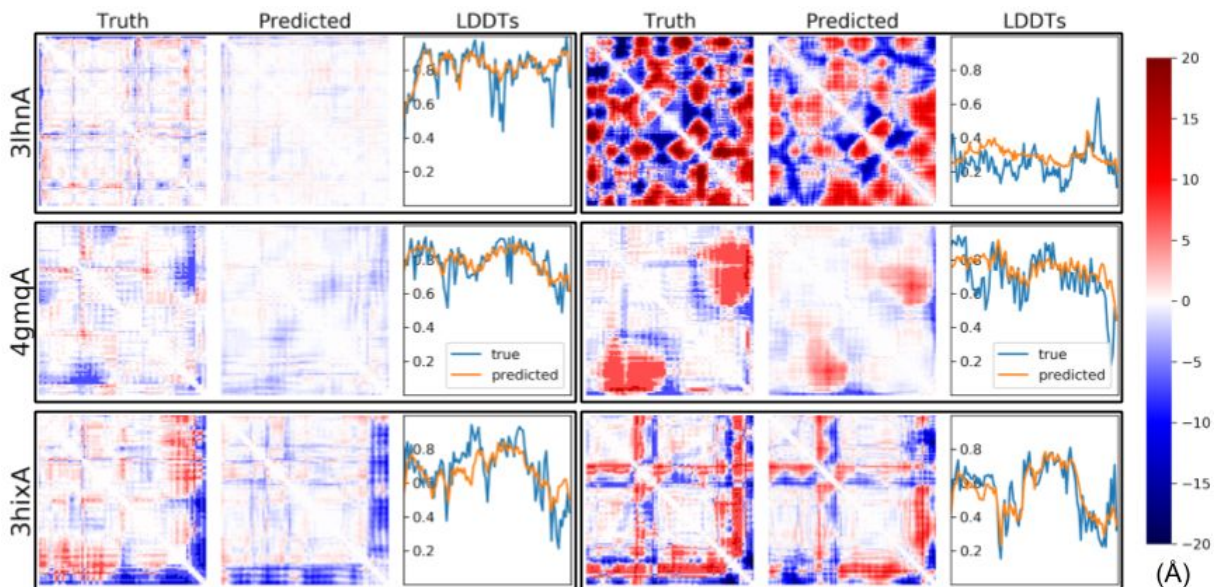
Figure 2: Example estograms and l-DDT score prediction. Model predictions for two randomly selected decoys for three test proteins were randomly selected (3lhnA, 4gmqA, 3hixA; size 108, 92 and 94 respectively; black rectangular boxes delineate results for single decoy). The first and fourth columns show true maps of errors, the second and fifth columns show predicted maps of errors, and the third and sixth columns show predicted and true l-DDT scores. The i,j element of the error map is the expectation of actual or predicted estograms between residues i and j in the model and native structure. Red and blue indicate that the pair of residues are too far apart and too close, respectively. The color density shows the magnitude of expected errors.

Figure 3: DeepAccNet performance. A) Contribution of individual features to network performance; all models include the distance matrix features. Overall, the largest contribution is from the features generated by 3D convolutions on local environments. Estogram (A, cross-entropy) loss values averaged over all decoys for each test protein are shown as one datapoint. The grey dotted line shows the median values from predictors (i) and (ii). B) Comparison of the performance of single model accuracy estimation (EMA) methods on CASP13 and CAMEO data. (top) Performance of local accuracy estimation measured by area under receiver operator characteristic (ROC) curve for predicting mis-modeled residues per sample (C_β I-DDT< 0.6) (bottom) Performance of global accuracy estimation measured by the Spearman correlation coefficient (*r*-value) of predicted and actual global l-DDT scores per target protein. The blue horizontal lines show the median value of our predictor. Equant 2 is not evaluated for global accuracy estimation due to its missing submissions. C) Predicted C_β I-DDT correlates with resolution for X-ray structures (left; Spearman-r 0.43 with p-value < 0.0001), X-ray structures of trans-membrane proteins (middle; Spearman-r 0.62 with p-value < 0.0001), and cryoEM structures (right; Spearman-r 0.86 with p-value < 0.0001). D. X-ray structures have higher predicted I-DDT values than NMR structures.

Figure 4. Consistent improvement in model structures from refinement runs guided by network based accuracy predictions. Refinement calculations guided and not guided by network accuracy predictions were carried out on a 73 protein target set (5, 21) (see Methods for details). A) Network guided refinement consistently improves starting model. B) Network guided refinement trajectories produce larger improvements than unguided refinement trajectories. The accuracy of the refined structure (l-DDT; y-axis) is compared with that of the starting structure in A, and with the final refined structure not using DL-based accuracy predictions in B (5). Top and bottom panels show results for proteins less than 100 residues in length and 100 or more residues in length, respectively. Each point represents a protein target with color indicating the protein size (scale shown at the right side of panel B). C) Actual and predicted model accuracy improvements throughout the refinement trajectory. Model quality (actual in blue and predicted in gray), averaged over 73 test cases, is shown through the refinement process. Points and bars show the model1 quality and the quality range of 50 models in the pool, respectively. D)

Distribution of accuracy improvements. In both panels, bottom, light colored boxes represent improvements without DeepAccNet, while upper, darker regions of the boxes represent additional improvements gained with DeepAccNet; these are calculated over the complete benchmark set. (left panel) Similar improvements are observed across secondary structure types (H:helix, E:extended, C:coil). (right panel) Improvements in model secondary structure accuracy are evaluated on 3- or 8-states following DSSP annotations (25); improvements are evident in both 3 state and 8 state local structure prediction. E) Example of predicted versus actual per-residue accuracy prediction. Predicted and actual l-DDT values are shown before (left) and after refinement (right) with a color scheme representing local l-DDT from 0.0 (red) to 0.7 (blue). Native structure is overlaid in gray color. Red arrows in the panels highlight major regions that have been improved. F) Examples of improvements in refined model structures. For each target, starting structures are shown on the left and the refined model on the right. Color scheme is the same as E, showing the actual accuracy.

Supplementary figures
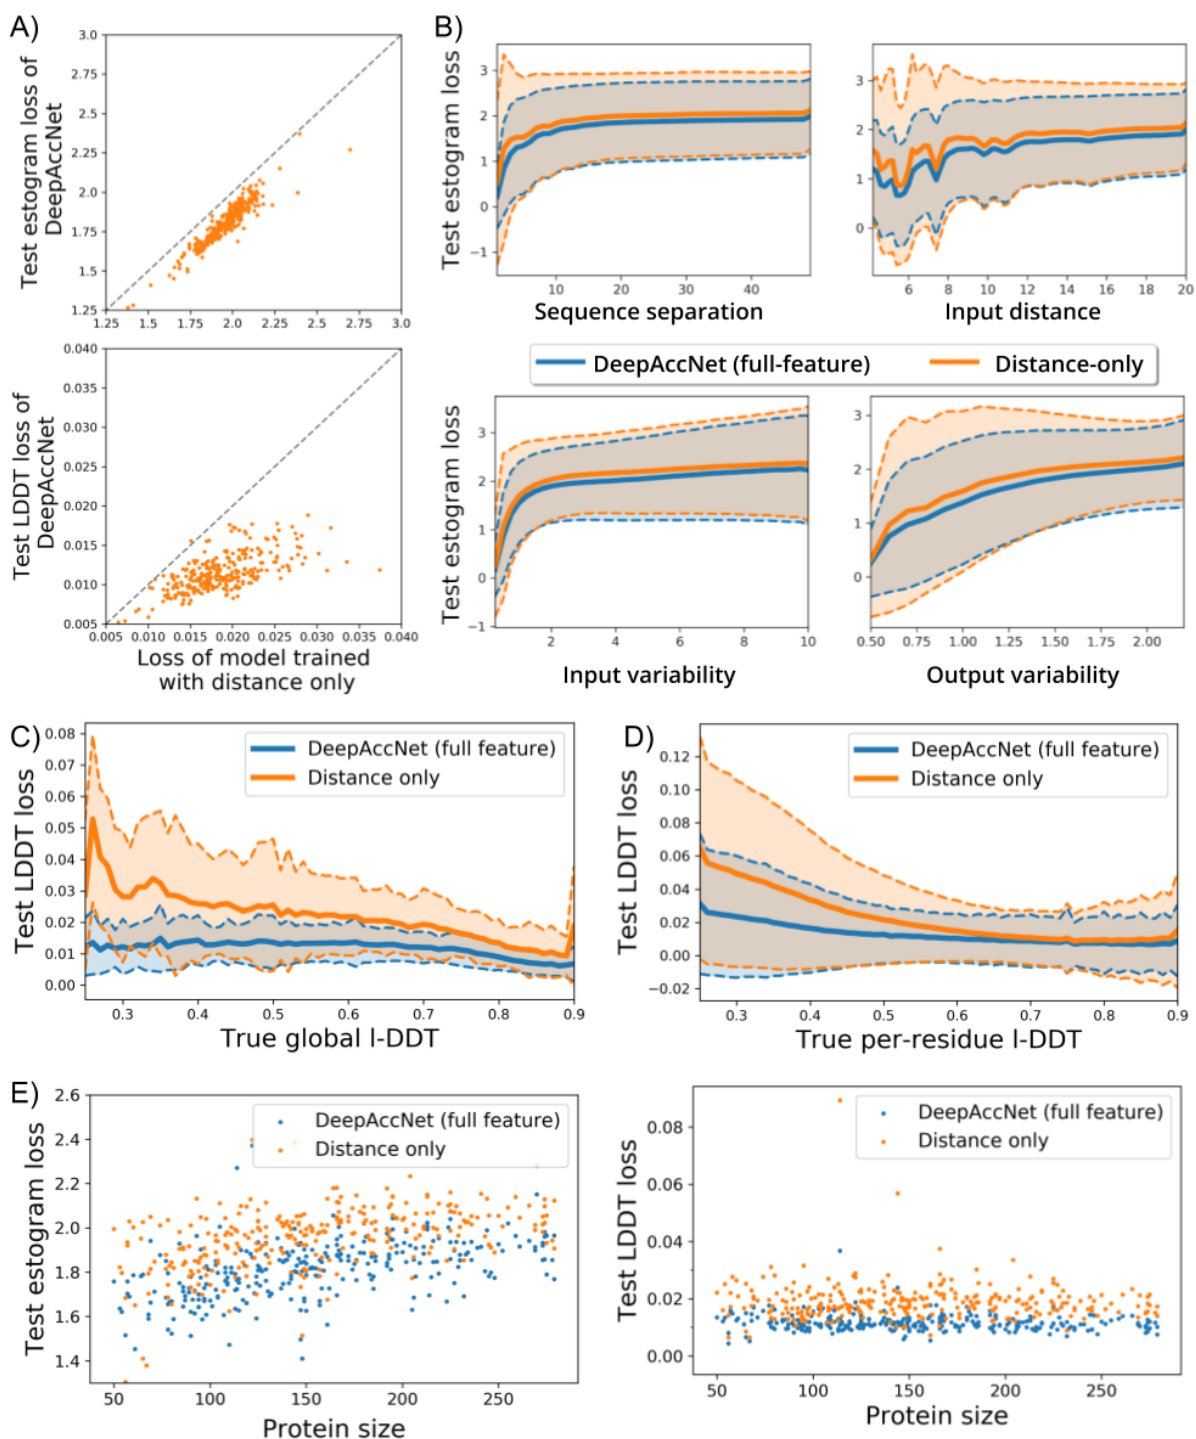
FigureS1. A) Comparison of DeepAccNet to distance-only network on predicted estograms (top) and l-DDT scores (bottom). Each dot represents the loss for a single protein averaged over all decoys. Lower loss values indicate better performance. Estograms are evaluated by cross-entropy loss, and per residue l-DDT scores are evaluated by mean-squared error. B) Test estogram loss plotted against four conditions; sequence separation,

input distance, input variability (standard deviation of input distance across decoys from the same target), and output variability (entropy of true estogram across decoys from the same target). The loss values are binned in terms of x-axis properties. The mean value at each bin is shown on the y-axis, and the range of one z-score is shown with the shaded area. CD) Dependence of l-DDT score loss on true l-DDT per-model (C) and per residue (D). Loss values are binned in terms of the true l-DDT scores. The mean of loss values at each bin is shown on the y-axis as a solid line, and the range of one Z-score is shown with the shaded area. E) Dependence of estogram (top) and l-DDT score per residue (bottom) loss on protein size. Each dot is an average loss value for a single target protein over all decoys.
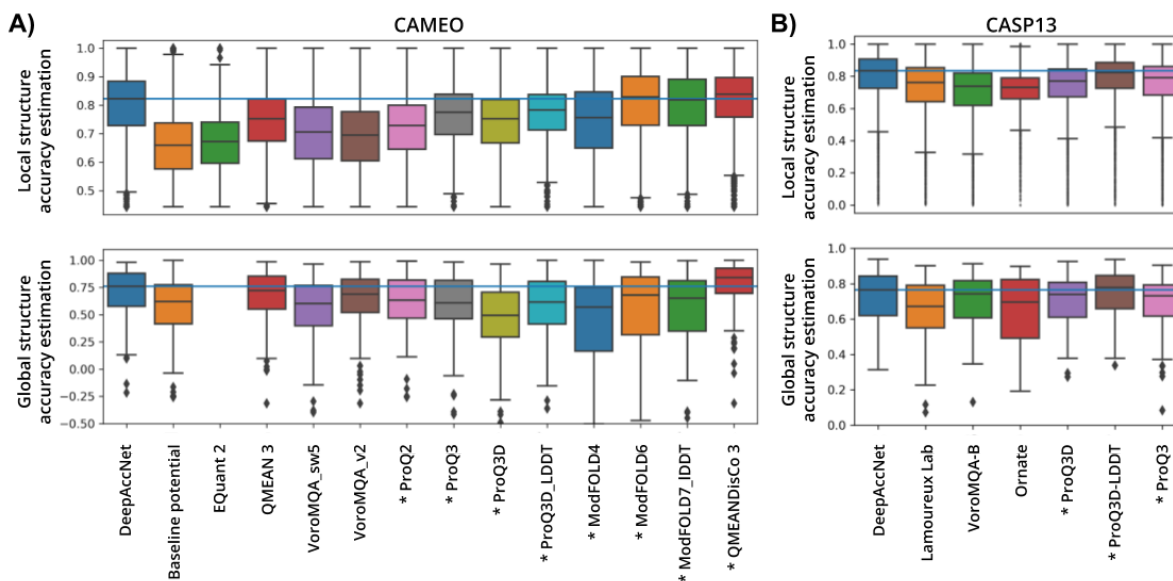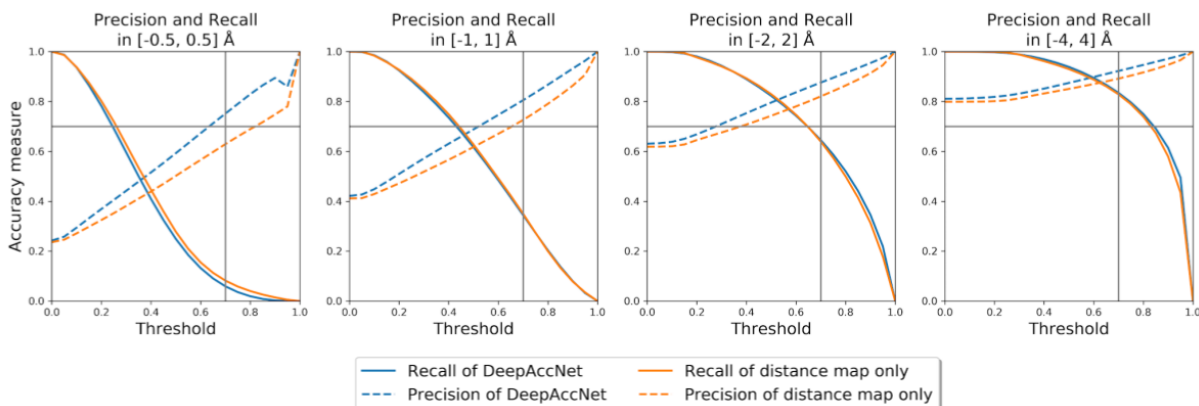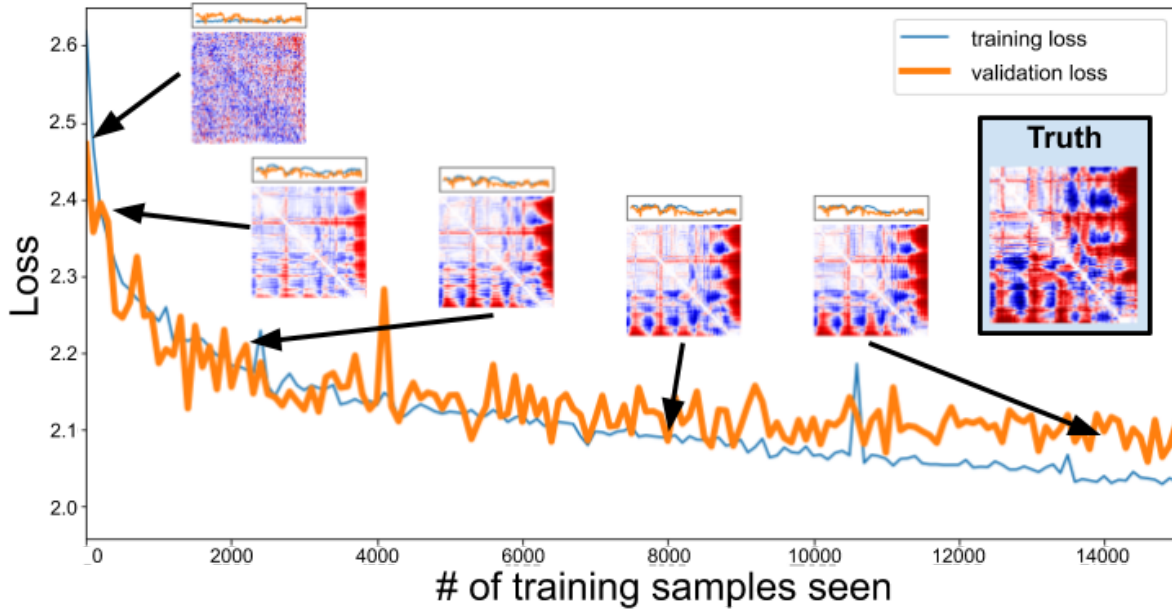

FigureS2. Comparison of the performance of single model accuracy estimation (EMA) methods on CASP13 and CAMEO data. (top) Local accuracy estimation evaluated by area under receiver operator characteristic (ROC) curve for predicting mis-modeled residues per sample (C_β l-DDT< 0.6). (bottom) Global accuracy estimation evaluated by the Spearman correlation coefficient (*r*-value) of predicted and actual global l-DDT scores per target protein. The blue horizontal lines show the median value of our predictor. The methods that predict based on homology or co-evolutionary information are indicated with an asterisk.
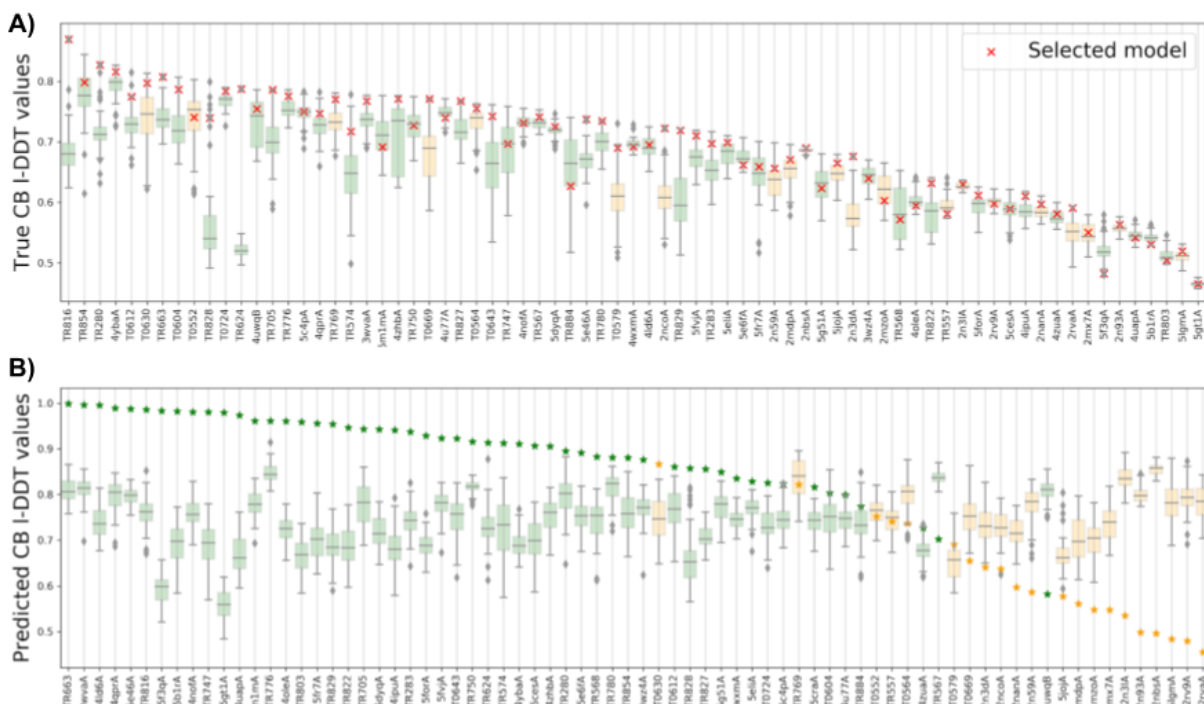
FigureS3. A) Selection of decoy structures at the end of refinement trajectory for each target protein. Red cross indicates samples selected by DeepAccNet (average global I-DDT of 0.685). B) Predicted I-DDT scores of the sampled and native structures. Stars indicate predictions on native target structures. In both A and B, targets for which native structures are determined by X-ray crystallography and NMR are shown in green and yellow, respectively.

FigureS4. Assessment of binary correct/incorrect predictions. Actual error values were grouped into correct and incorrect bins. In each panel, a distance is counted as correct if the actual distance error (from that of the native structure) is within a certain range, while a prediction is counted correct if the sum of probability over the given range in the estogram is above the threshold value (x-axis). Error range definitions are [-0.5, 0.5], [-1, 1] , [-2, 2], and [-4, 4] Å from the left to the right panel. The dotted lines show recall values and solid lines show precision values. The grey lines visualizes the thresholding of 0.7 used in the downstream refinement process.
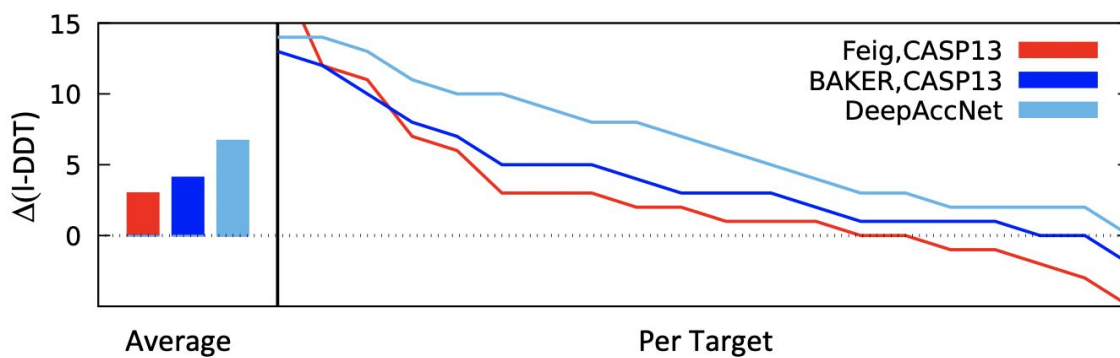
FigureS5. The training and validation loss of DeepAccNet training process. The lines show total loss values which is a weighted sum of estogram, I-DDT, and mask loss values (see Methods). The images show test sample prediction at different epochs.
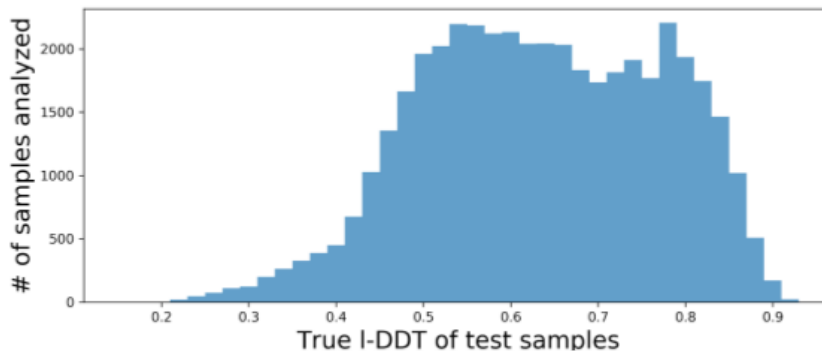
FigureS6. Numbers of samples that participated in loss analysis based on starting I-DDT scores.

26

FigureS7. Performances of the methods on CASP13 refinement category targets. Improvements in I-DDT scores over starting models are shown. Two leading groups in CASP13, Feig and Baker, are brought in for the comparison against refinement with DeepAccNet; Feig group ran long MD simulations, while BAKER group ran the non-DL refinement method presented in the main text with subsequent short MD simulations. Net I-DDT changes for both of these groups range within 3~4%, compared to 7% by DeepAccNet-guided refinement. 9 targets from the CASP13 refinement category are removed from the analysis for which the native structures contain heavy oligomeric contacts or are determined at low resolutions (>3 Å).

Supplementary tables

Models	Held-out proteins (# proteins=194)			True global l-DDT < 0.7 (# samples=16215)			True global l-DDT > 0.7 (# samples=3154)		
	Esto	Mask	I-DDT	Esto	Mask	I-DDT	Esto	Mask	I-DDT
(i) Full features	*1.8691	*0.2434	*0.0124	*1.9100	*0.2620	*0.0122	*1.6412	*0.1394	*0.0134
(ii) Distance only	1.9739	0.2689	0.0169	2.0134	0.2884	0.0158	1.7531	0.1601	0.0231
(iii) w/ 3d convolution	1.8736	0.2448	0.0126	1.9141	0.2633	0.0123	1.6471	0.1414	0.0142
(iv) w/ AA-related	1.9289	0.2565	0.0155	1.9606	0.2740	0.0138	1.7519	0.1593	0.0251
(v) w/ Backbone	1.9559	0.2656	0.0161	1.9966	0.2850	0.0154	1.7282	0.1574	0.0200
(vi) w/ residue Orientation	1.9684	0.2684	0.0168	2.0071	0.2878	0.0156	1.7526	0.1604	0.0233
(vii) w/ Rosetta energy	1.9404	0.2606	0.0146	1.9826	0.2801	0.0141	1.7050	0.1519	0.0172
(viii) w/ Secondary structure	1.9776	0.2701	0.0175	2.0148	0.2891	0.0161	1.7700	0.1642	0.0253
(ix) w/o 3d convolution	1.9092	0.2521	0.0141	1.9488	0.2708	0.0136	1.6885	0.1482	0.0172

Table S1: Performance of the variants of distance-based networks trained with and without a certain class of features. Performance is measured by cross-entropy for estograms and masks and mean squared error for I-DDT scores. We obtained four independent training trajectories for each setting and averaged the prediction from the models with the best validation performance (see Methods). Rows 2 to 4 report the quality of the three predictions averaged over all held-out decoy structures. Rows 5 to 7 report the quality of the predictions on decoys with low true quality (global I-DDT < 0.7). Rows 8 to 10 report the quality of the predictions on decoys with high true quality (global I-DDT > 0.7). The best performing model in each category is indicated with an asterisk. The biggest discrepancy between the model with full features against other models comes from the high-accuracy decoys.

Distance-based	i) C_β to C_β distance map, Cα is taken for GLY, ii) Cα to Tip-atom distance map and its transpose, iii) Tip-atom to Tip-atom distance map, and iv) sequence separation map. The distance maps (i~iv) go through a variance reduction process with *arcsinh(x)*. See Table S4 for the definition of tip atoms.
Amino acid properties	i) One-hot encoded amino acids. ii) Blosum62 scores (32). iii) Per amino-acid feature sets from Meiler et al (33).
Rosetta energy terms	i) Two-body energy terms: fa_atr, fa_rep, fa_sol, lk_ball_wtd, fa_elec, hbond_bb_sc, and hbond_sc. ii) One-body energy terms: p_aa_pp, rama_prepro, omega, fa_dun. iii) Presence of backbone-to-backbone hydrogen bonds.
Backbone angles and lengths	i) Phi, Psi, and Omega angles. ii) Standardized length between backbone atoms.
residue-residue orientations	i) Full 6 degrees of freedom of translation and rotation. ii) cosine and sine of Dihedral and planar angles defined by Yang et al (1).
Secondary structures	1-hot encoded representation of three state secondary structures given by DSSP solver.
Local atomic environments	24 by 24 by 24 voxels of size 0.8Å. In total, it covers an area of size 19.2Å by 19.2Å by 19.2Å. There are 20 channels for 20 atom types defined by Rosetta (See Table S3). The coordinate frame is fixed based on backbone N,Ca,C atoms(9).

Table S2: Generated features for all 7 major feature classes. Some features are scaled and normalized to a reasonable range. Please refer to the code available at github for further details on the normalization scheme.

Layers groups	Descriptions
3D convolution layers	This group has four layers of 3D convolution operations with 20, 20, 30, and 20 filters with sizes of 1, 3, 4, 4, respectively. Elu activation is used. Mean pooling of filter size 4 with stride 4 was performed at the end.
Feature merging	This operation merges flattened 3D conv outputs, 2D, and 1D features (**see Method**). One layer of 2D convolution with 32 filters of size 1 and instance normalization are applied. Elu activation is then used. Finally, the output is upsampled to 256 channels for the following ResNet operations.
Residual blocks 1	Each residual block consists of (i) elu activation, (ii) projection down to 128 channels, (iii) elu activation layer (iv) 3 by 3 convolution, (V) elu activation, (vi) projection up to 256 channels. Instance normalization operations are applied. Residual connection adds inputs to (i) with outputs of (vi). 20 residual blocks are stacked. Dilation is applied to (iv) with a cycling dilation size of 1,2,4,8.
Residual blocks 2 for estograms and masks	Two arms of four residual blocks are applied to predict estograms and masks. The same numbers of channels (256-->128-->256) are used.
l-DDT calculation layers	l-DDT values are calculated within gpu memory based on predicted estograms and masks (**see Methods)**.
Loss	(i) Estograms are evaluated with categorical cross-entropy loss. (ii) Masks are evaluated with binary cross-entropy loss. (iii) l-DDT values are evaluated with mean squared loss. Global loss is defined and shown in Method.

Table S3: Model architectures for the DeepAccNet. Please refer to the code available at github for further details on the implementation.

amino acid	ALA	CYS	ASP	ASN	GLU	GLN	PHE	HIS	ILE	GLY
tip atom	CB	SG	CG	CG	CD	CD	CZ	NE2	CD1	CA
amino acid	LEU	MET	ARG	LYS	PRO	VAL	TYR	TRP	SER	THR
tip atom	CG	SD	CZ	NZ	CG	CB	OH	CH2	OG	OG1

Table S4: Definitions of tip atoms for each residue.

6B17, 3URO, 3TWG, 5DYR, 6HR0, 1P9G, 4G4L, 6EWN, 4HB6, 5JQF, 4U2W, 4HB8, 1MBN, 4HAJ, 1CYC, 1VXB, 3H4N, 2SBT, 1NXB, 4HBF, 1G7V, 2EWI, 1J0O, 2SNS, 4HDL, 3SJ4, 3H34, 4D5M, 1MBS, 1OS6, 2EWU, 1LWK, 1LYZ, 3TRV, 3SJ0, 4Z0W, 1ACX, 1PMK, 3TJW, 1HH5, 1M1R, 6DK5, 2ZVS, 3D6T, 2AOA, 3SEL, 6FM8, 5YP8, 4EFX, 1TGL, 3SJ1, 1TIA, 2EWK, 2XJI, 5HDD, 6CDX, 5VBD, 4HC3, 3NIR, 2YYX, 1HGU

Table S5: List of X-ray native structures with low C_β-lddt despite their high experimental resolution.