



A General Framework to Learn Tertiary Structure for Protein Sequence Characterization

Mu Gao* and Jeffrey Skolnick*

Center for the Study of Systems Biology, School of Biological Sciences, Georgia Institute of Technology, Atlanta, GA, United States

OPEN ACCESS

Edited by:

Jianlin Cheng,
University of Missouri, United States

Reviewed by:

Jie Hou,
Saint Louis University, United States
Xin Gao,
King Abdullah University of Science
and Technology, Saudi Arabia

*Correspondence:

Mu Gao
mu.gao@gatech.edu
Jeffrey Skolnick
skolnick@gatech.edu

Specialty section:

This article was submitted to
Protein Bioinformatics a section
of the journal
Frontiers in Bioinformatics

Received: 01 April 2021

Accepted: 06 May 2021

Published: 21 May 2021

Citation:

Gao M and Skolnick J (2021) A General Framework to Learn Tertiary Structure for Protein Sequence Characterization. *Front. Bioinform.* 1:689960. doi: 10.3389/fbinf.2021.689960

During the past five years, deep-learning algorithms have enabled ground-breaking progress towards the prediction of tertiary structure from a protein sequence. Very recently, we developed SADLSA, a new computational algorithm for protein sequence comparison via deep-learning of protein structural alignments. SADLSA shows significant improvement over established sequence alignment methods. In this contribution, we show that SADLSA provides a general machine-learning framework for structurally characterizing protein sequences. By aligning a protein sequence against itself, SADLSA generates a fold distogram for the input sequence, including challenging cases whose structural folds were not present in the training set. About 70% of the predicted distograms are statistically significant. Although at present the accuracy of the intra-sequence distogram predicted by SADLSA self-alignment is not as good as deep-learning algorithms specifically trained for distogram prediction, it is remarkable that the prediction of single protein structures is encoded by an algorithm that learns ensembles of pairwise structural comparisons, without being explicitly trained to recognize individual structural folds. As such, SADLSA can not only predict protein folds for individual sequences, but also detects subtle, yet significant, structural relationships between multiple protein sequences using the same deep-learning neural network. The former reduces to a special case in this general framework for protein sequence annotation.

Keywords: protein structure prediction, protein folding, sequence alignment, deep-learning, structural alignment

INTRODUCTION

The amino acid sequence of a protein encodes the information for carrying out its function. One essential aspect is the tertiary structure of the protein. Indeed, the prediction of protein tertiary structure from its sequence is a fundamental question in biophysics (Anfinsen and Scheraga, 1975). In order to predict protein structure at high accuracy, one main challenge is to model the long-range, many-body effects that collectively dictate a protein's tertiary structure (Skolnick and Gao, 2021). Over the past several years, exciting breakthroughs have been made to better address these long-range interactions (Skolnick and Gao, 2021). Using a deep residual convolutional neural network, significant success has been demonstrated in predicting contacts between individual residues of a protein sequence (Wang et al., 2017). Such residue-residue contacts yield both the local secondary structure and the global fold, and it is the accurate prediction of their synergy that improves model quality. Subsequently, several groups demonstrated that better residue-residue contact or detailed distance matrix (distogram) predictions led to significantly improved structure predictions,

especially for challenging targets (Gao et al., 2019; Hou et al., 2019; Xu, 2019; Zheng et al., 2019; Senior et al., 2020) In CASP13, a blind biannual protein structure prediction competition, all four top-ranked groups in the most challenging, free-modeling category used residue-residue contacts or distance matrices predicted via deep-learning. Among them, DeepMind's AlphaFold achieved the best performance using a high-quality distogram to derive statistical folding potentials (Abriata et al., 2019). In CASP14, many improved deep-learning approaches using the convolutional residual networks were presented, e.g., (Yang et al., 2020), but AlphaFold2 dominated the competition using a new, end-to-end deep-learning algorithm with an attention mechanism.

A topic closely related to protein structure prediction is protein sequence comparison or alignment (Söding and Rimmert, 2011). In the low pairwise sequence identity regime of less than 30%, two protein sequences may exhibit no apparent sequence similarity yet display significant fold similarity when their structures are revealed and superimposed (Chothia and Lesk, 1986). This observation is attributed to the fact that the structural space of protein folds is very small for both evolutionary (Murzin, 1998) and physical reasons (Skolnick et al., 2014). Traditionally, a variety of sequence alignment approaches have been developed and applied to assist protein structure prediction, e.g., Hidden Markov Model (HMM) (Hildebrand et al., 2009) and "threading" approaches (Jones, 1999; Skolnick et al., 2004; Zhou and Zhou, 2005). These efforts provide the foundation for template-based modeling approaches (Zhang, 2008). Conversely, if the structures encoded in the two sequences are known, their structural alignment generally leads to a more accurate sequence alignment than those from classical sequence alignment approaches. Such an accurate, meaningful alignment is often the key to understanding what a novel protein sequence does, e.g., predicting functional sites (Gao and Skolnick, 2009; Gao and Skolnick, 2013).

Naturally, this leads to a question: can deep-learning be directly applied to generate a protein sequence alignment with an accuracy close to the structural alignment counterpart? If so, this would not only extend the ability to recognize evolutionarily distant sequence relationships but also enable a deeper learning of the folding code. Moreover, it has practical applications for function and structure prediction and possibly evolutionary inference. To answer this question, we recently developed SAdLSA, a sequence alignment algorithm that uses a deep convolutional neural network to learn many thousands of structural alignments (Gao and Skolnick, 2020). The advantage of SAdLSA was demonstrated in benchmark tests against HMM-based HHsearch (Soding, 2005). For challenging cases, SAdLSA is ~150% more accurate at generating pairwise alignments and ~50% more accurate at identifying the proteins with the best alignments in a sequence library. This allowed the program to detect remote relationships that may be useful for genome annotation or functional predictions.

Given the encouraging benchmarking results of SAdLSA, one would like to understand why it performs better than classical sequence comparison approaches. Obviously, the deep-learning

algorithm plays a key role in this improvement, but how does it work? Previously, we have speculated that SAdLSA implicitly learns the protein folding code without offering direct evidence. In this study, we shall further substantiate this claim. We noticed that when the same sequence was input into SAdLSA, SAdLSA aligns the sequence against itself, i.e., self-alignment, and outputs an intra-sequence distogram for the input. This distogram could encode the fold much like a deep-learning algorithm designed to predict the distogram for a single query sequence, e.g., DESTINI (Gao et al., 2019). We shall perform analysis to understand the distogram generated by SAdLSA self-alignment and demonstrate that SAdLSA provides a more general framework to learn protein structures for sequence annotation purposes.

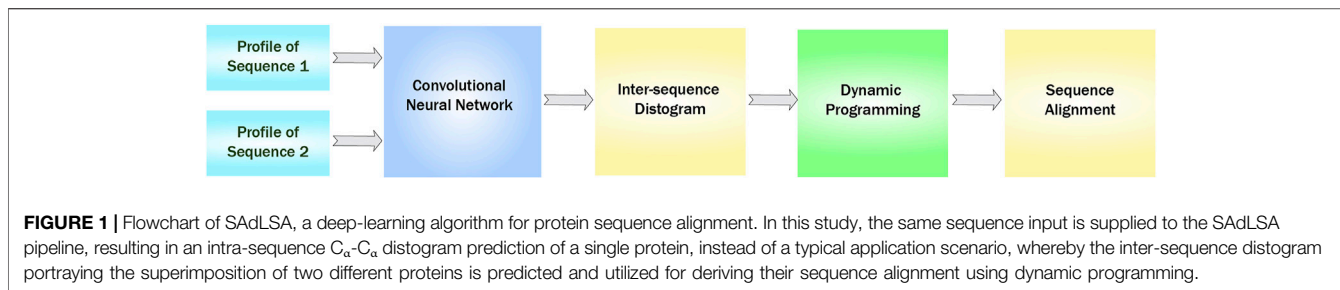
METHODS

For this study, we mainly employ SAdLSA, a deep-learning (DL) based approach for protein sequence alignment. The details of SAdLSA and the benchmark results have been described in detail elsewhere (Gao and Skolnick, 2020). Here, we briefly recapitulate its key features.

SAdLSA

An overview of SAdLSA is presented in **Figure 1**. The inputs to this network are two position-specific sequence profiles, each of dimension $N_k \times 20$, where N_k is the length of the k -th sequence ($k = 1, 2$), and the 20 columns represent the 20 different amino acids at each residue position (hence position-specific). Here, we use the profiles generated from HHblits (Rimmert et al., 2011). In machine-learning language, the sequence profiles are also known as embeddings. The outer product of these two 1D sequence features yields a 2D matrix of features, where at position (i, j) of the matrix the elements are a concatenation of the 20 columns formed from the i -th residue of sequence 1 and the j -th residue of sequence 2. Subsequently, these 2D features are fed into a fully convolutional neural network consisting of up to 34 residual blocks. The main objective of this neural network model is to predict residue-residue distances between the two input sequences that recapitulates their optimal structural alignment, using observed structural alignments as the training ground truth. The training distance labels are created from structural alignments by APoc (Gao and Skolnick, 2013), which takes advantage of a global alignment provided by TM-align (Zhang and Skolnick, 2005).

The neural network is composed of multiple residual blocks, either conventional (Deep residual learning for image recognition, 2016) or dilated (Dilated residual networks, 2017) in slightly different design variants. The residual block design is a key to train a deep neural network model. Within a residual block, each convolutional layer is composed of about 64 filters with a kernel size of 3×3 or 5×5 . After the residual blocks, the last 2D convolutional layer outputs 22 channels, representing 21 distance bins (1–20 at 1 Å each, and >20 Å) and channel 0 which is reserved for ignored pairs (e.g., gap residues missing in a structure, or large distances >30 Å). Finally, a softmax layer calculates the probability scores for each distance bin. Here,



the same network was also applied to the same two input. The mean probability scores of these two runs are the final output score for this DL model. This ensures that the alignment is identical if one swapped the position of two input sequences. For this study, we used the consensus scores from six DL models, including three models with 14 residual blocks and $64 \times 5 \times 5$ kernels in each convolutional layer, and three dilated model with 34 residual blocks (alternating 1, 2, 8, 16, 32, dilate rates) and 50 to $75 \times 3 \times 3$ kernels. The two type of DL models have 2.9 and 2.4 million parameters, respectively.

The outputs from a DL model are the probabilities of distance bins forming an inter-protein residue-residue distance matrix. To build an alignment using dynamic programming (DP), we convert this probability matrix into a mean distance matrix D , whose element

$$d_{ij} = \sum_{k=1}^n p_{ij}^k b_k - c \quad (1)$$

where i, j are target/template sequence positions, p_{ij}^k is the probability for bin k at position (i, j) , b_k are distance labels from the sequence (1, 2, ..., 20, 22). D is subsequently adapted as the scoring matrix to obtain the optimal alignment using a Smith-Waterman-like DP algorithm (Smith and Waterman, 1981). The distance matrix D is also used to calculate an estimated TM-score (Zhang and Skolnick, 2004; Xu and Zhang, 2010) for ranking the significance of an alignment. The constant c is set at one such that a perfect alignment gives an estimated TM-score of one.

Since we study the self-alignment of a given sequence, we feed exactly the same sequence profiles into SAdLSA, and the resulting sequence alignment itself is universally at 100% identity, with a predicted TM-score ~ 1 . We focus on the residue-residue distance matrix D , which is converted from a general inter-sequence scenario into a special intra-sequence scenario, because the two input sequences are the same.

Data Sets

We employed the same test and training sets from the original SAdLSA study (Gao and Skolnick, 2020). Both sets are curated from the SCOP database (Fox et al., 2014). The training set is composed of 79,000 pairs from a SCOP30 set of $\sim 5,000$ domains. These training protein domain sequences share less than 30% identity. The test set is an extrinsic test set of the sequences of 593 randomly selected protein domains from 391 SCOP folds. Homologs of the testing sequences at 30% sequence identity or higher, or with a BLAST E-value < 0.1 , are excluded from the

training set. In this study, we employed SAdLSA models trained on this training set and conducted SAdLSA self-alignments on each of the 593 test sequences.

Analysis

Distogram Assessment

It has been well-established that residue-residue contacts characterize a protein structural fold (Gao et al., 2019; Skolnick and Gao, 2021). The most common definition of a protein contact is based on the distance between C_{α} or C_{β} atoms. That is, a contact between a pair of protein residues i and j is defined if the Euclidean distance between their C_{α} (or C_{β}) atoms is less than a cutoff value, typically at 8 \AA . A popular contact metric adopted by the CASP assessors (Schaarschmidt et al., 2018) is the precision of the top L/k C_{β} - C_{β} contact predictions within short, medium or long ranges, where L is the length of the target and $k = 1, 2$, and 5 , and the sequential distance s_{ij} of residues i and j , $s_{ij} \equiv |i - j|$ defines the range: *short* [$s_{ij} \in (6, 11]$], *medium* [$s_{ij} \in (12, 23]$], and *long* [$s_{ij} \in (24, \infty)$], i.e., nonlocal residue pairs. Since SAdLSA was trained on the distances between C_{α} atoms, we use C_{α} - C_{α} contacts with an 8 \AA cutoff as our definition of protein inter-residue contacts and consider only those belonging in either the medium or long-range regime, i.e., $s_{ij} \in [12, \infty)$. The predictions are ranked by the probability of forming a C_{α} - C_{α} contact. To obtain the probability, one simply sums the probabilities for distance bins from 0 to 8 \AA , since the SAdLSA DL models output a probability matrix D for 21 distance bins. This probability score is then employed for the precision analysis as outlined above. The precision is defined as $TP/(TP + FP)$, where TP is the number of true positives, i.e., native contact observed in an experimental structure, and FP is the number of false positives within the top L/k contact predictions evaluated.

In addition, we introduce the Mean Absolute Error (MAE) of the predicted distogram vs. the ground truth distogram. Specifically, we calculate the MAE using the coordinates of the C_{α} atoms of nonlocal residue pairs, i.e., the sequential distance of residue pairs is no less than six. The overall MAE for a distogram is defined as

$$d_{MAE} = \sum_{i,j} |d_{ij}^{pre} - d_{ij}^{nat}| / M \quad (2)$$

where (i, j) are the indexes of a pair of nonlocal residues separated up to 20 \AA in the native distogram, M the total of such pairs, d_{ij}^{pre} the predicted distance by the SAdLSA self-alignment, and d_{ij}^{nat} the distance observed in the native distogram. Additionally, for each

target, we also obtain the MAE values within each distance bin from 4 to 20 Å. If a target does not have any nonlocal residue pair within a specific 1 Å bin, the MAE calculation is skipped for this bin.

In order to estimate the statistical significance of d_{MAE} , we consider its expected value using a background distribution based on the pairwise residue distances observed in the training distograms. For each distance bin, we count the residue pairs within this bin as observed in the training ground truth distograms, and then obtain the observed frequency f^k for this distance bin by dividing the count against all counts of all bins from 1 to 21 (inter-residue distances between 20 and 30 Å are used for bin 21 as they are what collected for training). Substituting f^k for p_{ij}^k in Eq. 1, we obtain the expected value of d_{ij} , $d_{ij}^{exp} = 15.5$ Å for any (i, j) according to this reference distribution. Note that our training ground truth distograms are of inter-sequence residues, in comparison to the distograms of intra-sequence residues from the self-alignment. But, since the inter-sequence distograms are actually employed for SAdLSA training, it is appropriate to use the distance distribution collected from these distograms as the reference background. This leads to a naïve way to calculate the expect d_{MAE} , $d_{MAE}^{exp} = \sum_{i,j} |d_{ij}^{exp} - d_{ij}^{nat}| / M$, yielding a mean d_{MAE}^{exp} of 3.07 Å and a standard deviation of 0.20 Å by applying the formula to 4,661 structures employed in the training set. Likewise, for comparison the same formula is applied to the test set, including the overall error and errors for individual distance bins. Using the one-tailed test for a normal distribution, one can calculate the p -value of an d_{MAE} value using the background distribution derived from the training set.

t-SNE Analysis

We performed t-distributed Stochastic Neighbor Embedding (t-SNE) (Van der Maaten and Hinton, 2008) to analyze the factors that contribute to the precision of contact predictions by SAdLSA self-alignment. This is a nonlinear dimensionality reduction tool more suitable for this analysis than a classical principal component analysis. Five features were used for this analysis including i) the best training pair for each target measured by their TM-score (Schaarschmidt et al., 2018), ii) the ratio of the total observed native contacts (combining both medium- and long-range contacts, i.e., $s_{ij} \in [12, \infty)$) over the sequence length, whether the best training pair belongs to iii) the same fold or just the iv) superfamily as the target according to the SCOP classification, and v) the sequence diversity of the multiple sequence alignment of the target. The TM-score is a protein length-independent metric ranging from 0 to 1, and a TM-score >0.4 indicates a statistically significant alignment (Zhang and Skolnick, 2005). We use (0,1) to represent the logical variables (e.g., is it a member of the same SCOP fold or not). The sequence diversity is calculated with the Multiple Sequence Alignment (MSA) of the target and defined as $\ln(N_{eff}/\sqrt{L})$ taken from (Gao and Skolnick, 2020), where N_{eff} is the number of effective sequences in the MSA without length normalization. In our t-SNE analysis, we employed the default parameters including the perplexity parameter set at 30.

DESTINI2

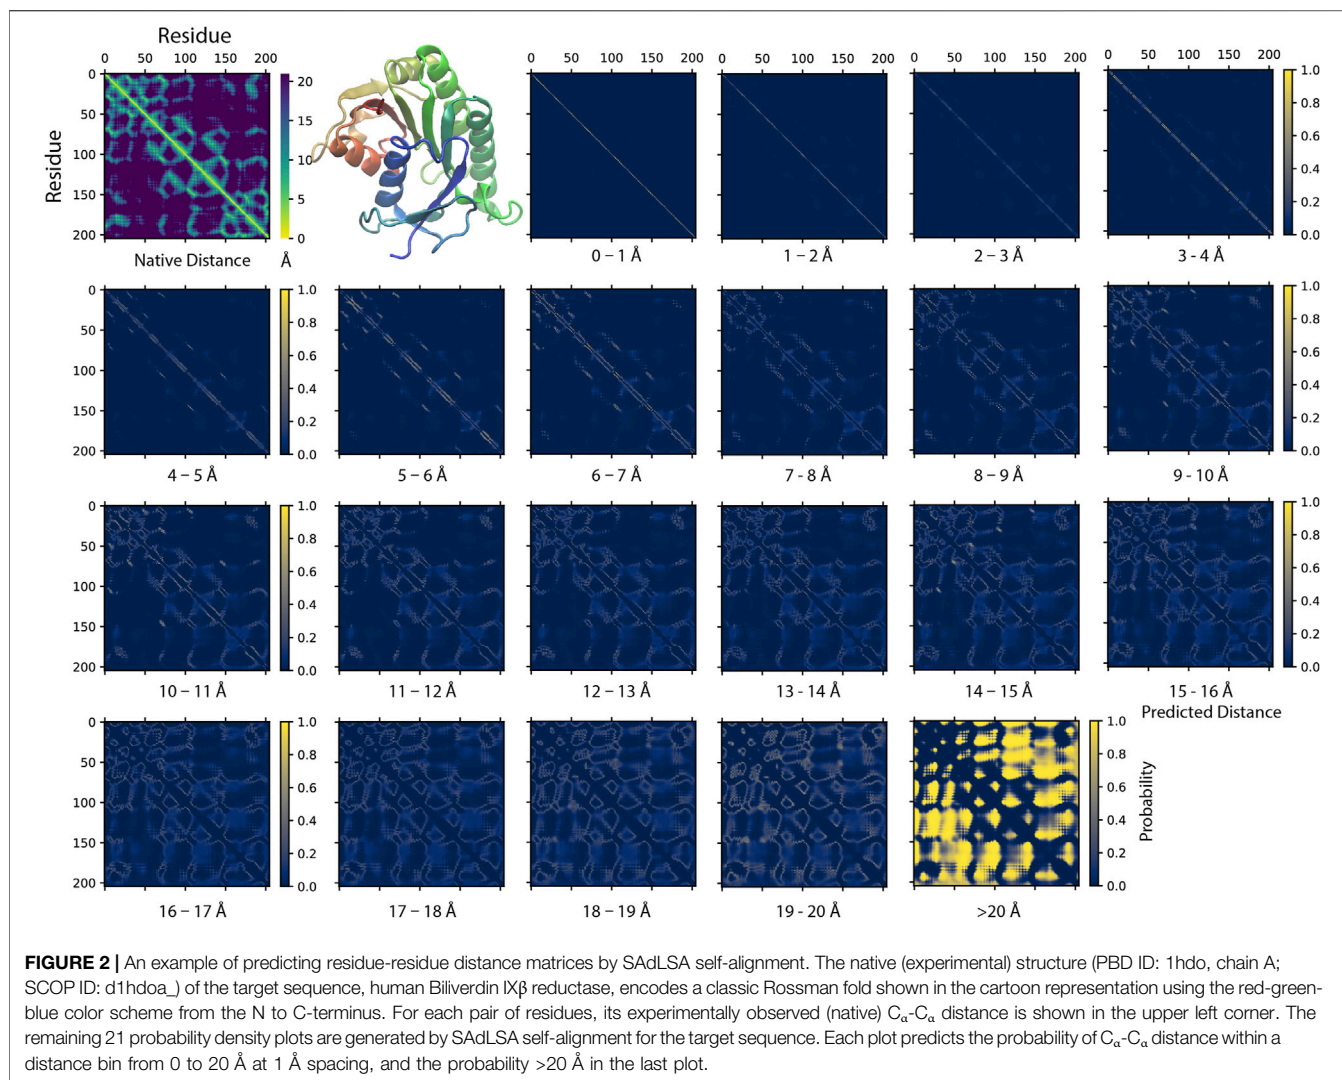
For comparison purposes, we employed DESTINI2 to conduct inter-residue distance prediction on the same test data set, i.e., 593 SCOP domain sequences, used to benchmark SAdLSA. DESTINI2 improves DESTINI by using a deeper, dilated convolutional residual network model. In this study, we used 39 dilated residual blocks similar to the one implemented in SAdLSA. For training, about 10,000 crystal structures with <2.5 Å resolution were taken from a March 2020 release of PISCES (Wang and Dunbrack Jr., 2003), which were curated from the PDB database (Berman et al., 2000). For a fair comparison, we retrained DESTINI2 models by removing the test set entries and their close homologs from the original DESTINI2 training set using the same criteria as that used for the SAdLSA test.

RESULTS

We conducted self-alignments for 593 target sequences with SAdLSA, using deep-learning models trained without obvious homologs to any of the 593 target sequences (see Methods). As one would expect, SAdLSA returns a sequence alignment at 100% identity and a predicted TM-score close to one. This seems trivial. But, if one carefully inspects the residue-residue distogram prediction, it not only contains information giving rise the identical alignment, but also contains inter-residue distance information to structurally characterize the fold encoded by the sequence itself.

SAdLSA Self-Alignment Generates a Fold-Depicting Distogram

Figure 2 illustrates an example of a SAdLSA self-alignment prediction. This 205 AA target sequence encodes a classic Rossmann fold, which is composed of an $\beta\alpha\beta$ alternating secondary structural segments found in many nucleotide-binding proteins (Hanukoglu, 2015). The characteristics of this fold are displayed in the distance plot calculated between the C_α atoms determined in the crystal structure (Figure 2, top left panel). The remaining residue-residue distance plots are generated by the SAdLSA self-alignment of the same sequence. These are from the 21 scoring channels designed to predict the probability of each pair of C_α atoms falling into each distance bin from 0 to beyond 20 Å. The first three channels are straight diagonal lines, which give rise the 100% identity in the sequence alignment and are not the focus of this study. Starting from plots ≥ 4 Å, one recognizes inter-residue distance relationships. First, the immediate neighboring residues are shown between 3–4 Å. Then, the main secondary structure elements including the fold's six α -helices and seven β -strands exhibit their patterns in the 4–5 and 5–6 Å plots. The packing between these secondary structural elements becomes clear in the subsequent three plots from 6–9 Å. The detailed packing orientations among secondary structural elements are further delineated in the remaining distance plots up to 20 Å. Finally, the highlights in the >20 Å plot signal the regions that are distant from each other. The top L medium or long-range



contact predictions for this case have a precision of 86%, which is sufficient to reconstruct a high-resolution structural model whose TM-score >0.7 (Skolnick and Gao, 2021). The mean absolute error d_{MAE} , calculated from the C $_{\alpha}$ -C $_{\alpha}$ distances of all nonlocal residue pairs separated up to 20 Å, is only 1.57 Å from the native distogram.

How accurate is the SADLSA self-alignment for predicting residue-residue distances in general? If one examines the d_{MAE} , the overall number looks good with a mean of 2.43 Å, and 92.7% of targets are below 3 Å (Figure 3 insert). By comparison, if one naively assigns distance distribution according to the observed fractions from the training set (see Methods), one obtains a mean expected error, d_{MAE}^{exp} at 3.08 Å. Overall, 96.6% of targets have a smaller distance prediction error by SADLSA than the value from the naïve reference approach. Moreover, about 408 (69%) targets have a significant d_{MAE} below the p -value cutoff of 0.05. Figure 3 further details the distributions of MAE between the SADLSA predicted and the corresponding native distograms for individual distance bins from 4 to 20 Å. It is clear that residues forming direct contacts within the first five bins are most challenging to

predict and exhibit large variations, with the mean MAE gradually decreasing from 6.6 Å in the 3–4 Å bin to 3.4 Å in the 8–9 Å bin. But these distance predictions by SADLSA are actually highly significant in comparison to the expected values whose MAE errors are up to 8 Å higher on average than SADLSA predictions. On the other hand, the large distance bins from 14 to 18 Å yield relatively small MAE values <2.5 Å, but it is not surprising as the expected MAE is below 2 Å.

From a structural perspective, direct inter-residue contacts are the most important. Moreover, one only needs to predict a fraction of the total number of contacts in order to obtain a correct fold prediction (Gao et al., 2019). We therefore turn our focus to inter-residue contacts. According to the benchmark results on 593 protein sequences, the mean precisions are 40.3, 52.6, and 63.7% for the top L , $L/2$, $L/5$ predictions of medium or long-range inter-residue C $_{\alpha}$ contacts, respectively (Table 1). The detailed distributions of individual predictions are shown in violin plots (Figure 4). For instance, in the middle top $L/2$ plot, about half of the sequences have a precision value >50%, which is sufficient to derive the correct fold for a single-domain

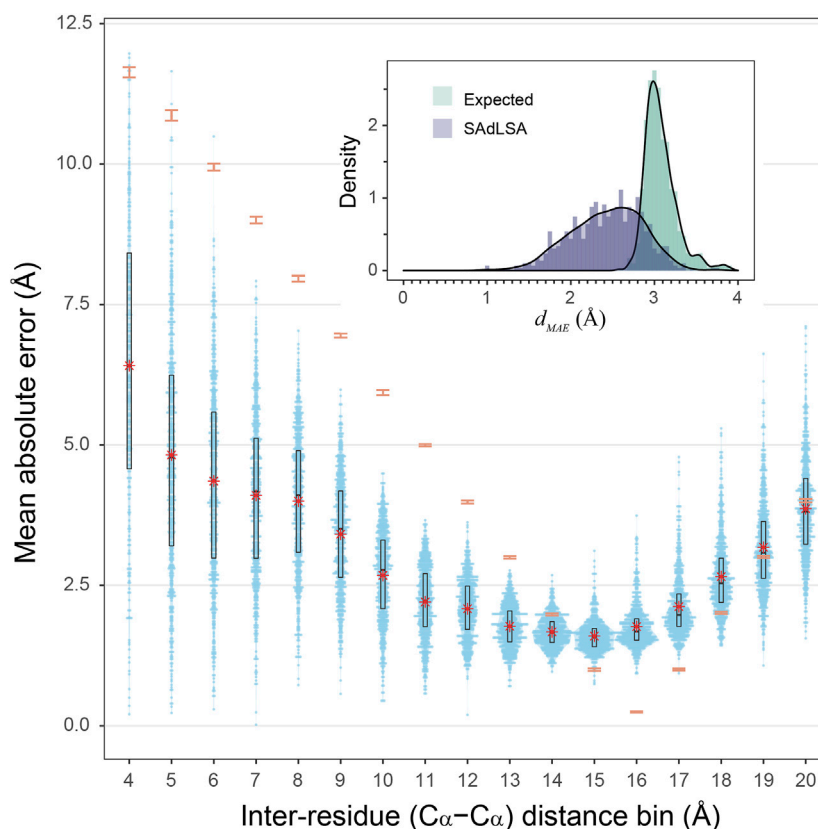


FIGURE 3 | Mean absolute error of predicted distogram vs native distogram for 593 test cases. The main plot displays the distribution of MAE values within each C_{α} - C_{α} distance bin from 4 to 20 Å for nonlocal residue pairs. Each bin has a spacing of 1 Å and each blue point represents a target protein. Violin contours are proportional to the counts of targets at different MAE levels with a bin width of 0.025. Black boxes and bars represent the second and third quartiles (25–75% ranked by MAE values) and the median of the distributions. The red stars represent the mean values. For comparison, the expected MAE distributions are shown in orange error bars, which are centered at the mean and extended to \pm sd. The insert shows the histograms of the d_{MAE} values in two separate assessments, calculated (purple) from the SAdLSA self-alignment and expected (teal) from the background distribution. For each target, its d_{MAE} value is calculated from all nonlocal residue pairs up to 20 Å as observed in the native distogram.

TABLE 1 | Mean precision of Medium/Long-range inter-residue contacts for 593 targets.

Method	Medium/Long-range Contacts*		
	L	L/2	L/5
SAdLSA (C_{α} - C_{α})	0.403	0.526	0.637
DESTINI2 (C_{α} - C_{α})	0.645	0.777	0.858
DESTINI2 (C_{β} - C_{β})	0.678	0.803	0.879

*Medium/Long-range Contacts denote residue-residue contacts whose sequential distance are in either the medium or long-range regime. Contact predictions are converted from the distogram predicted by each method. All deep-learning models were re-trained to exclude the test set and their close homologs from the original training sets (see Methods).

using these predictions (Gao et al., 2019). At L/5, 107 (18%) entries have 100% precision, and the median precision is at 69%. These numbers are not as good as the deep-learning approaches specifically trained to predict residue distograms, e.g., they are about 25% worse than DESTINI2 (Table 1). Given the fact that SAdLSA was not trained to predict residue distances for a single

structure (see more reasoning in Discussions), but rather to recognize the similarity between pairs of structures, these results are encouraging and clearly demonstrate the generality of the SAdLSA framework for learning protein structures.

What Contributes to Successful Distogram Prediction by SAdLSA Self-Alignment?

Next, we seek to understand the factors that affect the accuracy of distogram predictions by SAdLSA self-alignment. Like all machine-learning algorithms, the capability of distogram predictions must come from the training data of SAdLSA. Although SAdLSA was not trained to learn individual protein folds, we hypothesize that the fold-depicting distogram provided in the SAdLSA self-alignment comes from learning the training pairs sharing fold similarity as the target sequence. During SAdLSA training, it mainly focuses on aligned residues across two different sequences, but the network also observes the relative positions among aligned residues. As a result, SAdLSA learns the folding code of individual protein sequences, provided that they exhibit significant fold similarity. The more similar in their




FIGURE 4 | Precision of inter-residue contact predictions via SADLSA self-alignment. The precision of the top L (light green), $L/2$ (green), and $L/5$ (gold) are shown as circles for each target sequence. Violin contours are proportional to the counts of targets at different precision levels with a bin width of 0.01. Black boxes, median bars and means are represented following the same plotting scheme as in **Figure 3**.

structures, ideally the same fold as the training pair, the more likely the distogram pattern is learned for this specific fold.

Presence of the Same Fold as a Target in the Training set Is Important

To explore the validity of this hypothesis, for a target structure and a pair of training structures, we introduce T , defined as the minimal TM-score of the two training structures with respect to the target structure, using the target to normalize the TM-score. For each target, we find the highest T , denoted as T^* , among all pairs in the training set. **Figure 5** shows the correlations between T^* and the precision for the top $L/2$ residue-residue contact predictions. The violin plots demonstrate a clear upward trend in precision from the low T^* regime to high T^* regime. When there is a lack of obvious training structures that are similar to the target, e.g., when $T^* < 0.4$, the median precision is only at 39%. The same metric increases to 52% if $T^* \in [0.5, 0.6)$, and dramatically to 82% if $T^* > 0.7$. The Pearson correlation coefficient between T^* and the precision of all targets is 0.34, which clearly shows the dependency of precision on T^* , despite the indication that other factors are also in play.

For example, **Figure 5** shows that there are still very good distogram predictions in the low T^* regime. How could SADLSA accurately predict a distogram when there is no similar training

structure? There are 11 target sequences with $L/2$ precision $>60\%$ within $T^* < 0.4$. If one examines these structures, some of them are composed of multiple domains or subdomain structures. For example, one target, with SCOP ID d3u7qb, is a 522 AA sequence composed of three Rossmann folds and two helical bundle domains, despite the fact that SCOP defines it as a single domain. Although there is no other structure in the training set that resembles the overall structure of this target, the folds of its individual domains can be learned separately. As shown above, the Rossmann fold is relatively easy to learn (**Figure 2**). As a result, the overall precision prediction is 76% and the d_{MAE} is 2.44 Å (p -value = 8.0×10^{-4}) for this query sequence. More interesting examples are analyzed below.

Evolutionary Relationships Facilitate Fold-learning

What are the other possible contributing factors? In addition to T^* , we further consider four other features including the ratio of total observed native contacts (within either the medium or long-range) over the sequence length, whether the T^* training pair belong to the same fold or superfamily according to the SCOP classification and the sequence diversity of the multiple sequence alignment of the target. The Pearson correlation coefficient between each of these four additional features and the contact precision is 0.55, 0.30, 0.33, and 0.33, respectively. **Figure 6** shows

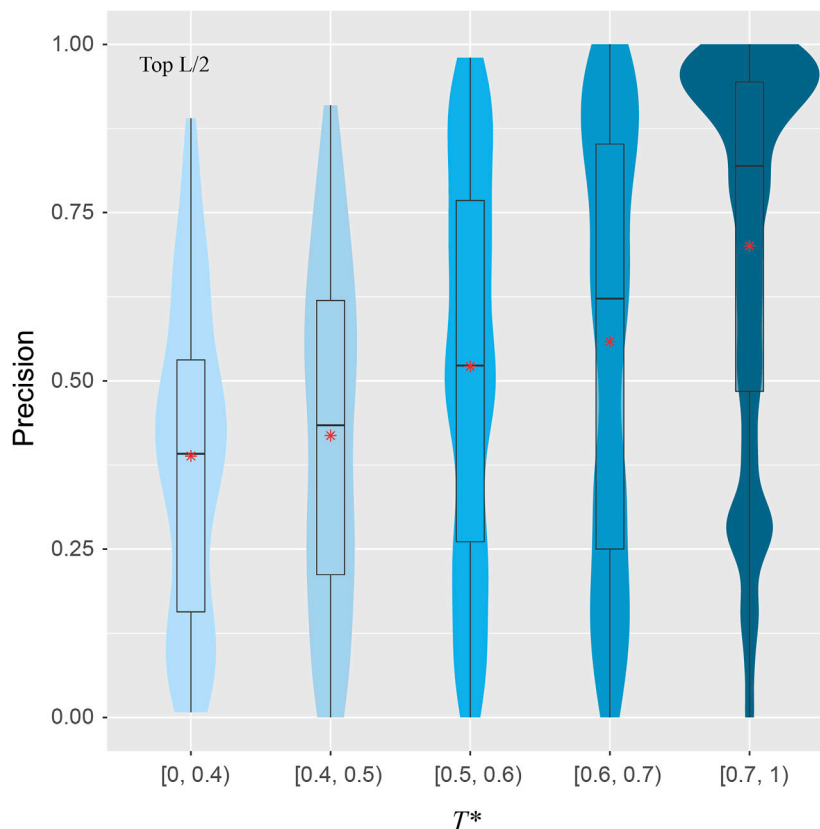


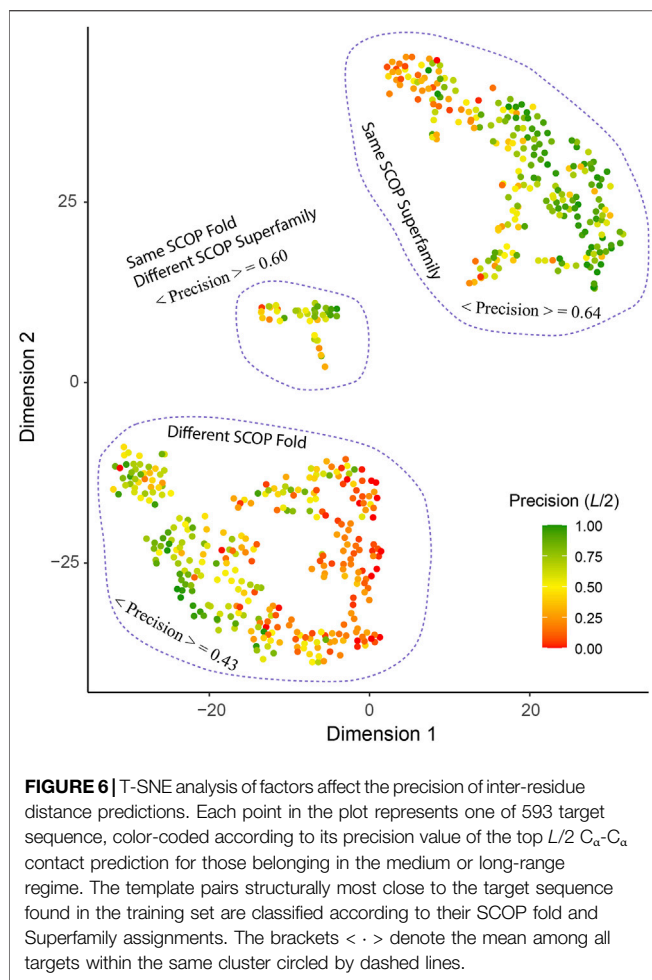
FIGURE 5 | Precision of inter-residue contact prediction vs T^* , the highest TM-score to all structures in the pairs found in the training set. The violin and box plots follow the same scheme as in **Figure 3**.

the results from t-SNE analysis of these five features. Three big clusters emerge and represent targets whose T^* training pair are from the same SCOP superfamily (240 entries), the same SCOP fold but not the same superfamily (40 entries), and different SCOP fold (313 entries). They exhibit different levels of difficulty for contact prediction, at mean precision values of 64%, 60%, and 43%, respectively. This result makes sense as the sequence profiles from the same superfamily are similar and relatively easy to learn for a neural network model, whereas sequence profiles from remote families or those without an apparent evolutionary relationship are much harder to learn. In the same superfamily cluster, there are very few “bad” predictions, e.g., 41 (17%) targets at $L/2$ have a precision <30%. One explanation is the structural variation between the target and the training structures, despite the fact that they are from the same superfamily. On the other hand, even for cases where the evolutionary relationship might not be clear, it is still possible to predict a fold-depicting distogram reasonably well with SAdLSA, as is evident in the clusters where the SCOP superfamilies are different. In particular, we note that SAdLSA performs very well for a target structural fold if there are training structures originated from the same SCOP fold but not necessarily the same SCOP superfamilies, as is evident in the mean precision of 60% among 40 such entries, close to the value of 64% in the cluster representing the same superfamily.

Moreover, the sequence diversity correlates positively with the contact precision at a correlation coefficient of 0.33. There are 53 cases with $T^* < 0.4$, i.e., when there is a lack of similar fold in the training set for these targets. Among them, those with high diversity, e.g., 28 targets with $N_{eff} > 1,000$, have an average $L/2$ contact precision of 47% vs. 30% for 25 cases with smaller sequence diversity. Likewise, for cases when similar folds are present in the training set, e.g., $T^* > 0.6$, but the N_{eff} is low <100, the mean contact precision is at 42%, much lower than 64%, which the mean value for cases at the same $T^* > 0.6$ criterion but with higher sequence diversity.

Delineation of Protein Folds via Deep-learning Across SCOP Folds

Notably, even when the best training structures are from a different SCOP fold, there are still many highly accurate distogram predictions as exhibited in **Figure 6**. Indeed, there are 102 (33%) such cases in the different fold cluster with precision >60% among top $L/2$ predictions. Three representative examples are displayed in **Figure 7**. One main reason for this observation is that they may still have reasonably close or even a highly similar fold present in the training set, despite the different SCOP classifications. For 102 targets, 91/62% of them have $T^* > 0.4/0.5$, respectively. **Figure 7A** illustrates one such case, which is the C-terminal domain from pyruvate kinase



of *Leishmania mexicana* (Morgan et al., 2010). Even though there is no structure belonging to the same SCOP fold as the target, there are 12,239/498/10 training structures with TM-score $>0.4/0.5/0.6$ to the target structure, and the precision for contact prediction is at 98% with d_{MAE} of 1.68 Å (p -value = 2.5×10^{-12}) from the native distogram. A second example is delineated in **Figure 7B** for a domain from *Bacillus subtilis* Q45498 with unknown function. It has a T^* value of 0.47, among 151 training structures in the TM-score regime 0.4 to 0.5 but none higher. SAdLSA self-alignment makes a good prediction at 85% and d_{MAE} of 2.29 Å (p -value = 5.1×10^{-5}). For these examples, it is reasonable to conceive that SAdLSA learns to predict this fold at high precision values through the training on the comparison of these structures.

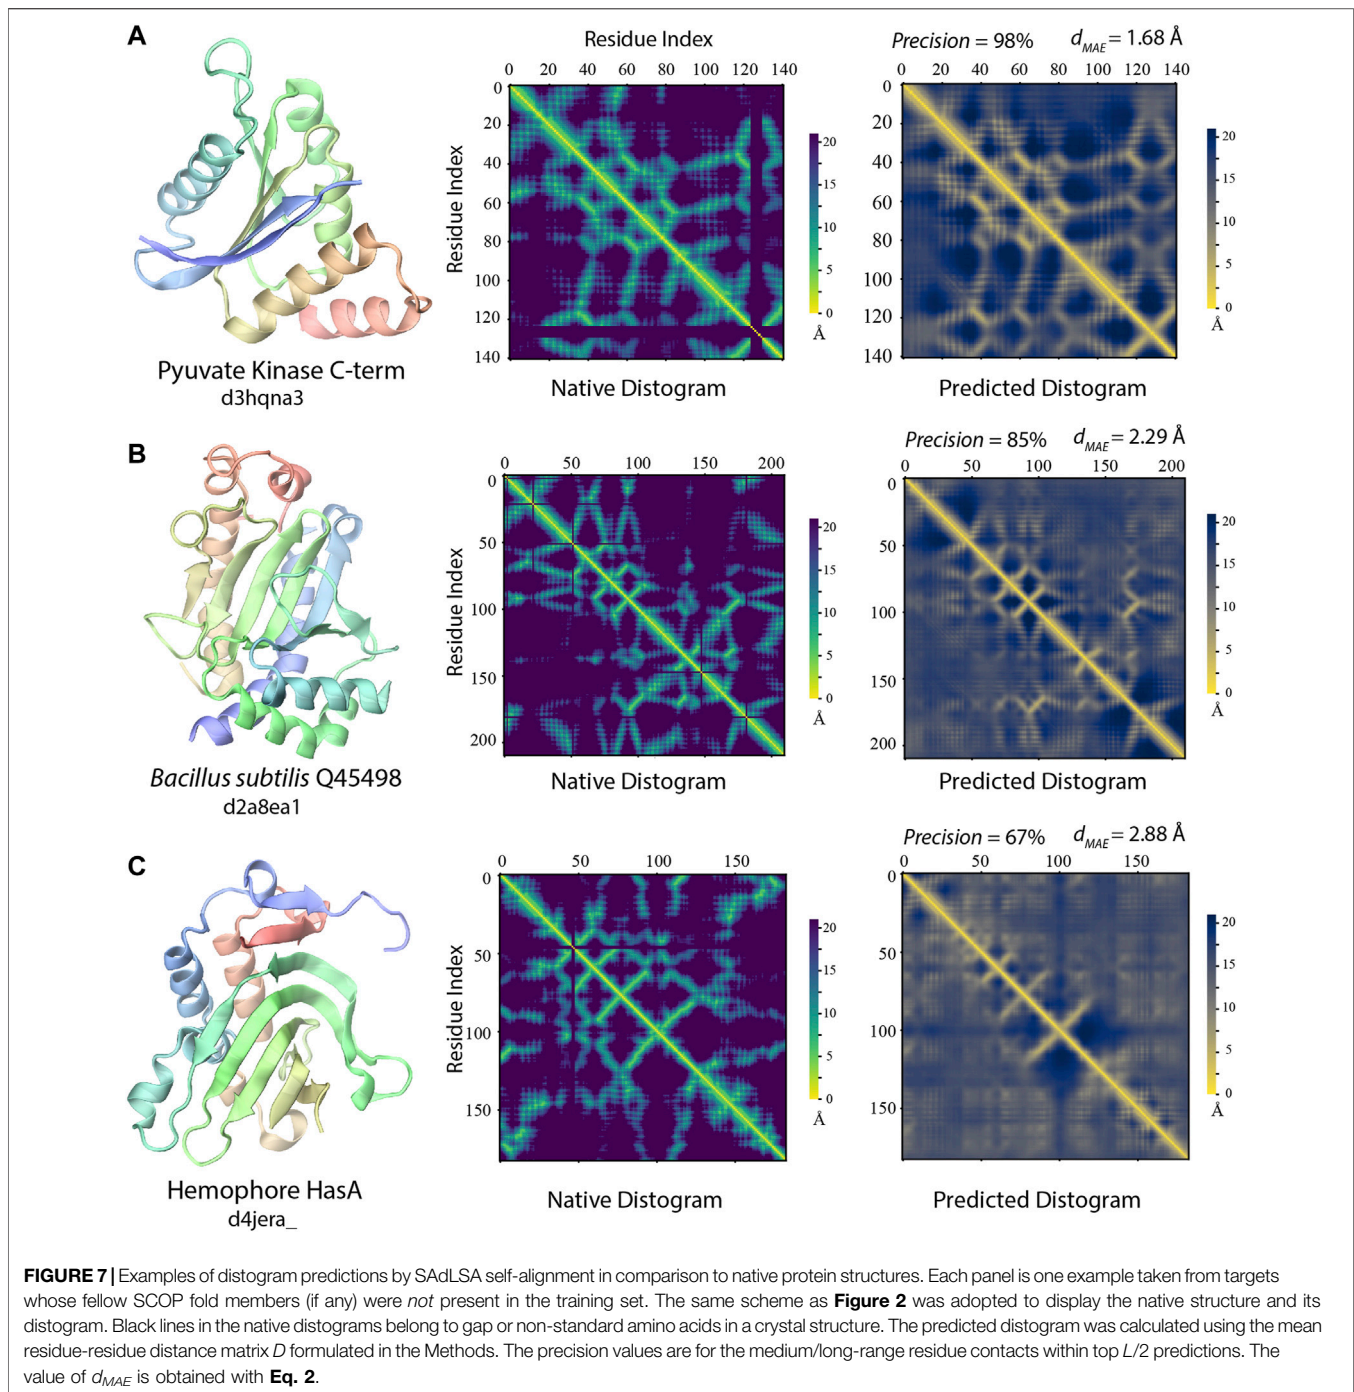
In addition, there are 127 targets that do not share the same SCOP fold with any member in the training set. Yet, for 32/41 cases, SAdLSA can predict residue-residue contacts at $>50/60\%$ precision. One such example is shown in **Figure 7C**. The target sequence is the hemophore HasA from *Yersinia pestis* (Kumar et al., 2013). The distogram prediction via SAdLSA alignment is at 67% with a d_{MAE} of 2.88 Å (p -value = 0.17). While the results are not as good as the above two cases, one can still find correct predictions between the main secondary structure elements.

Here, the long-range interactions by SAdLSA are fuzzy and imprecise. Nevertheless, the result is remarkable giving that no training structure shares a TM-score >0.4 , and its T^* is at 0.38. In these cases, SAdLSA likely learns the packing pattern for the subdomains or fragments of a target sequence from its training structures, which may share relatively low overall structure similarity but high similarity to some individual domains, subdomains or fragments, like the example in 3.2.1, but here it is more general and subtle without a clear definition of the domain.

Lastly, one technical reason why some targets have low precision is due to the definition of this metric, which penalizes the case where very few medium or long-range contacts are present in the observed native structure, e.g., coiled-coil structures. With relatively few or even no true positives in extreme cases, the L/n normalization will bring down the precision value. In fact, the ratio of native contacts over the length of the sequence has the strongest correlation (0.55) with the contact precision among all five features considered. There are 22 cases whose ratio <0.25 between the number of native medium/long-range contacts and the length of protein. Their mean precision at $L/2$ is only 9%. However, 15 of 22 have a significant $d_{MAE} < 2.68$ Å (p -value < 0.05). Moreover, due to the lack of long-range contacts, such structures (e.g., $\#contacts/L < 0.5$) are likely more flexible than compact folds, and therefore, are challenging for structure accurate prediction. In the same superfamily/different fold cluster, there are 12/49 cases forming the reddish subclusters at the edge in **Figure 6**. Of these 62 cases in total, 18 have a $d_{MAE} \geq 2.68$ Å. Thus, the reason for a few of these poor predictions might reflect the intrinsic propensity towards disorder for some proteins.

DISCUSSIONS

What structural information does one wish to obtain from a machine-learning algorithm, given an input protein sequence? Very recently, numerous approaches have employed deep-learning techniques to predict tertiary protein structure, notably through an inter-residue distogram (Skolnick and Gao, 2021). One may argue that a more general machine-learning approach should go beyond the prediction of the tertiary structure for a single sequence to predict structural relationships between multiple protein structures, which may lead to a deeper understanding of their sequence or functional relationships. For this purpose, we introduced SAdLSA, which predicts protein sequence alignments by learning their structural alignment via deep-learning (Gao and Skolnick, 2020). In this contribution via a self-alignment analysis, we extended the previous study and explicitly demonstrate that SAdLSA learns the protein folding code. The key to understanding the deep-learning folding code lies in the analysis of the distogram prediction. Indeed, we obtain distogram predictions at surprisingly high accuracy for many folds, at a mean precision of 52% for the top $L/2$ contact predictions and a mean absolute error of 2.43 Å in inter-residue distogram predictions. In terms of d_{MAE} , about 97% of predicted inter-residue distograms are better



than expected from a background prediction, and 74% are statistically significant. This explains the advantage of SAdLSA over the classic approaches as up to a 100% improvement was observed in benchmark tests (Gao and Skolnick, 2020).

How does SAdLSA obtain its fold-depicting capability? The most important contribution comes from similar fold structures subjected to training. The algorithm was designed to pay attention to the distances between aligned residues. When these two training structures share a similar fold, the

distogram of such fold can be learned, as evident in the correlation between the target and training structures (**Figure 5**). Additionally, if a training sequence is evolutionarily related to a target sequence, even remotely, it facilitates fold learning. More interestingly, SAdLSA can learn from training structures that go beyond the SCOP fold, e.g., cases that share no similar SCOP fold or even no similar overall fold whose TM-score >0.4 in the training set. This seems surprising at first but may be partially explained as follows: First, a target may

have many structural analogs with high structural similarity that escapes human manual classification. Second, the protein's structure may consist of multiple-domains, subdomains, or smaller fragments whose packing pattern may be learned individually and separately from many training structures. As with all machine-learning based algorithms, since the prediction capability of SAdLSA comes from the training set, the key question is how general is the resulting model? The success in this SAdLSA self-alignment benchmark that includes many challenging cases is a good indication of its generality. But ultimately, it needs the validation in real-world, large-scale applications, ideally at the proteome-level.

Despite this success, we note that the fold-prediction ability of SAdLSA self-alignment is not as accurate as deep-learning algorithms specially designed to predict protein structures, e.g., DESTINI2. There are two main reasons for this reduced performance. First, the training of SAdLSA is on the inter-sequence, inter-residue distances between a pair of superimposed structures, rather than the intra-sequence, inter-residue distances observed in a single structure. As such, SAdLSA self-alignment performs well when a pair of structures sharing a highly similar fold are present in the training set. But this requirement is not always true, especially as our training set is derived from ~5,000 SCOP domains, which is half the size of the training set for protein structure prediction, e.g., ~10,000 structures for DESTINI2's training. Only 41% of the testing targets have a close training pair at $T^* > 0.6$. Second, the current version of SAdLSA uses sequence profiles as its only input features, whereas many more features, especially the direct co-evolutionary signals that are essential for the success of DESTINI2 and the like, are not employed for a technical reason. Nevertheless, more recent deep-learning algorithms, such as AlphaFold2, directly learn from the multiple sequence alignment with an "attention" mechanism. Since the same multiple sequence alignment was used to derive the sequence profile used by SAdLSA as in DESTINI2, in principle, one may design a next generation, SAdLSA-like algorithm with a similar attention mechanism. It

will not only predict protein tertiary structure via self-alignments, but also compare structures encoded by two different sequences at high precision. Such work is now underway.

DATA AVAILABILITY STATEMENT

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession numbers can be found below: <https://sites.gatech.edu/cssb/sadlsa/>.

AUTHOR CONTRIBUTIONS

MG and JS contributed to the conception and design of the study. MG performed the research and data analysis and wrote the first draft of the article. All authors contributed to article revision, read, and approved the submitted version.

FUNDING

This work was supported in part by the Division of General Medical Sciences of the National Institute Health (NIH R35GM118039) and the DOE Office of Science, Office of Biological and Environmental Research (DOE DE-SC0021303). The research used resources supported by the Director's Discretion Project at the Oak Ridge Leadership Computing Facility and by the Partnership for an Advanced Computing Environment (PACE) at Georgia Tech.

ACKNOWLEDGMENTS

We thank Bartosz Ilkowski for internal computing support and Jessica Forness for proof-reading the manuscript.

REFERENCES

- Abriata, L. A., Tamò, G. E., and Dal Peraro, M. (2019). A Further Leap of Improvement in Tertiary Structure Prediction in CASP13 Prompts New Routes for Future Assessments. *Proteins* 87 (12), 1100–1112. doi:10.1002/prot.25787
- Anfinsen, C. B., and Scheraga, H. A. (1975). Experimental and Theoretical Aspects of Protein Folding. *Adv Protein Chem*, 29, 205–300. doi:10.1016/s0065-3233(08)60413-1
- Berman, H. M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T. N., Weissig, H., et al. (2000). The Protein Data Bank. *Nucleic Acids Res.* 28 (1), 235–242. doi:10.1093/nar/28.1.235
- Chothia, C., and Lesk, A. M. (1986). The Relation between the Divergence of Sequence and Structure in Proteins. *EMBO J.* 5 (4), 823–826. doi:10.1002/j.1460-2075.1986.tb04288.x
- Dilated Residual Networks (2017). F. Yu, V. Koltun, and T. Funkhouser in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Fox, N. K., Brenner, S. E., and Chandonia, J.-M. (2014). SCOPe: Structural Classification of Proteins-Extended, Integrating SCOP and ASTRAL Data and Classification of New Structures. *Nucl. Acids Res.* 42 (D1), D304–D309. doi:10.1093/nar/gkt1240
- Gao, M., and Skolnick, J. (2020). A Novel Sequence Alignment Algorithm Based on Deep Learning of the Protein Folding Code. *Bioinformatics* 37, 490–496. doi:10.1093/bioinformatics/btaa810
- Gao, M., and Skolnick, J. (2009). A Threading-Based Method for the Prediction of DNA-Binding Proteins with Application to the Human Genome. *Plos Comput. Biol.* 5 (11), e1000567. doi:10.1371/journal.pcbi.1000567
- Gao, M., and Skolnick, J. (2013). APoc: Large-Scale Identification of Similar Protein Pockets. *Bioinformatics* 29 (5), 597–604. doi:10.1093/bioinformatics/btt024
- Gao, M., Zhou, H., and Skolnick, J. (2019). DESTINI: A Deep-Learning Approach to Contact-Driven Protein Structure Prediction. *Sci. Rep.* 9 (1), 3514. doi:10.1038/s41598-019-40314-1
- Hanukoglu, I. (2015). Proteopedia: Rossmann Fold: A Beta-Alpha-Beta Fold at Dinucleotide Binding Sites. *Biochem. Mol. Biol. Educ.* 43 (3), 206–209. doi:10.1002/bmb.20849

- Hildebrand, A., Remmert, M., Biegert, A., and Söding, J. (2009). Fast and Accurate Automatic Structure Prediction with HHpred. *Proteins* 77 (Suppl. 9), 128–132. doi:10.1002/prot.22499
- Hou, J., Wu, T., Cao, R., and Cheng, J. (2019). Protein Tertiary Structure Modeling Driven by Deep Learning and Contact Distance Prediction in CASP13. *Proteins* 87 (12), 1165–1178. doi:10.1002/prot.25697
- Jones, D. T. (1999). GenTHREADER: an Efficient and Reliable Protein Fold Recognition Method for Genomic Sequences. *J. Mol. Biol.* 287, 797–815. doi:10.1006/jmbi.1999.2583
- Deep Residual Learning for Image Recognition (2016). Editors K He, X Zhang, S Ren, and J Sun in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*.
- Kumar, R., Lovell, S., Matsumura, H., Battaile, K. P., Moënne-Loccoz, P., and Rivera, M. (2013). The Hemophore HasA from *Yersinia pestis* (HasAyp) Coordinates Hemin with a Single Residue, Tyr75, and with Minimal Conformational Change. *Biochemistry* 52 (16), 2705–2707. doi:10.1021/bi400280z
- Morgan, H. P., McNaie, I. W., Nowicki, M. W., Hannaert, V., Michels, P. A. M., Fothergill-Gilmore, L. A., et al. (2010). Allosteric Mechanism of Pyruvate Kinase from *Leishmania Mexicana* Uses a Rock and Lock Model. *J. Biol. Chem.* 285 (17), 12892–12898. doi:10.1074/jbc.M109.079905
- Murzin, A. G. (1998). How Far Divergent Evolution Goes in Proteins. *Curr. Opin. Struct. Biol.* 8 (3), 380–387. doi:10.1016/s0959-440x(98)80073-0
- Remmert, M., Biegert, A., Hauser, A., and Söding, J. (2011). HHblits: Lightning-Fast Iterative Protein Sequence Searching by HMM-HMM Alignment. *Nat. Methods* 9 (2), 173–175. doi:10.1038/nmeth.1818
- Schaarschmidt, J., Monastyrskyy, B., Kryshchavych, A., and Bonvin, A. M. J. J. (2018). Assessment of Contact Predictions in CASP12: Co-evolution and Deep Learning Coming of Age. *Proteins* 86 (Suppl. 1), 51–66. doi:10.1002/prot.25407
- Senior, A. W., Evans, R., Jumper, J., Kirkpatrick, J., Sifre, L., Green, T., et al. (2020). Improved Protein Structure Prediction Using Potentials from Deep Learning. *Nature* 577 (7792), 706–710. doi:10.1038/s41586-019-1923-7
- Skolnick, J., and Gao, M. (2021). The Role of Local versus Nonlocal Physicochemical Restraints in Determining Protein Native Structure. *Curr. Opin. Struct. Biol.* 68, 1–8. doi:10.1016/j.sbi.2020.10.008
- Skolnick, J., Gao, M., and Zhou, H. (2014). On the Role of Physics and Evolution in Dictating Protein Structure and Function. *Isr. J. Chem.* 54 (8–9), 1176–1188. doi:10.1002/ijch.201400013
- Skolnick, J., Kihara, D., and Zhang, Y. (2004). Development and Large Scale Benchmark Testing of the PROSPECTOR_3 Threading Algorithm. *Proteins* 56 (3), 502–518. doi:10.1002/prot.20106
- Smith, T. F., and Waterman, M. S. (1981). Identification of Common Molecular Subsequences. *J. Mol. Biol.* 147 (1), 195–197. doi:10.1016/0022-2836(81)90087-5
- Söding, J. (2005). Protein Homology Detection by HMM-HMM Comparison. *Bioinformatics* 21 (7), 951–960. doi:10.1093/bioinformatics/btl125
- Söding, J., and Remmert, M. (2011). Protein Sequence Comparison and Fold Recognition: Progress and Good-Practice Benchmarking. *Curr. Opin. Struct. Biol.* 21 (3), 404–411. doi:10.1016/j.sbi.2011.03.005
- Van der Maaten, L., and Hinton, G. (2008). Visualizing Data Using T-SNE. *J. machine Learn. Res.* 9 (11).
- Wang, G., and Dunbrack, R. L., Jr. (2003). PISCES: a Protein Sequence Culling Server. *Bioinformatics* 19 (12), 1589–1591. doi:10.1093/bioinformatics/btg224
- Wang, S., Sun, S., Li, Z., Zhang, R., and Xu, J. (2017). Accurate De Novo Prediction of Protein Contact Map by Ultra-deep Learning Model. *PLOS Comput. Biol.* 13 (1), e1005324. doi:10.1371/journal.pcbi.1005324
- Xu, J. (2019). Distance-based Protein Folding Powered by Deep Learning. *Proc. Natl. Acad. Sci. USA* 116 (34), 16856–16865. doi:10.1073/pnas.1821309116
- Xu, J., and Zhang, Y. (2010). How Significant Is a Protein Structure Similarity with TM-Score = 0.5? *Bioinformatics* 26 (7), 889–895. doi:10.1093/bioinformatics/btq066
- Yang, J., Anishchenko, I., Park, H., Peng, Z., Ovchinnikov, S., and Baker, D. (2020). Improved Protein Structure Prediction Using Predicted Interresidue Orientations. *Proc. Natl. Acad. Sci. USA* 117 (3), 1496–1503. doi:10.1073/pnas.1914677117
- Zhang, Y. (2008). Progress and Challenges in Protein Structure Prediction. *Curr. Opin. Struct. Biol.* 18 (3), 342–348. doi:10.1016/j.sbi.2008.02.004
- Zhang, Y., and Skolnick, J. (2005). TM-align: a Protein Structure Alignment Algorithm Based on the TM-Score. *Nucleic Acids Res.* 33 (7), 2302–2309. doi:10.1093/nar/gki524
- Zhang, Y., and Skolnick, J. (2004). Scoring Function for Automated Assessment of Protein Structure Template Quality. *Proteins* 57 (4), 702–710. doi:10.1002/prot.20264
- Zheng, W., Li, Y., Zhang, C., Pearce, R., Mortuza, S. M., and Zhang, Y. (2019). Deep-learning Contact-map Guided Protein Structure Prediction in CASP13. *Proteins* 87 (12), 1149–1164. doi:10.1002/prot.25792
- Zhou, H., and Zhou, Y. (2005). SPARKS 2 and SP3 Servers in CASP6. *Proteins* 61 (Suppl. 7), 152–156. doi:10.1002/prot.20732

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2021 Gao and Skolnick. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.