

Identification of residue pairing in interacting β -strands from a predicted residue contact map

Wenzhi Mao^{1,2}, Tong Wang^{1,2}, Wenxuan Zhang^{1,2} and Haipeng Gong^{1,2,*}

¹MOE key laboratory of bioinformatics, School of Life Sciences, Tsinghua University, Beijing, China

²Beijing Innovation Center of Structural Biology, Tsinghua University, Beijing, China

*Corresponding author

E-mail: hgong@tsinghua.edu.cn (HG)

Abstract

Despite the rapid progress of protein residue contact prediction, predicted residue contact maps frequently contain many errors. However, information of residue pairing in β strands could be extracted from a noisy contact map, due to the presence of characteristic contact patterns in β - β interactions. This information may benefit the tertiary structure prediction of mainly β proteins. In this work, we introduce a novel ridge-detection-based β - β contact predictor, RDb₂C, to identify residue pairing in β strands from any predicted residue contact map. The algorithm adopts ridge detection, a well-developed technique in computer image processing, to capture consecutive residue contacts, and then utilizes a novel multi-stage random forest framework to integrate the ridge information and additional features for prediction. Starting from the predicted contact map of CCMpred, RDb₂C remarkably outperforms all state-of-the-art methods on two conventional test sets of β proteins (BetaSheet916 and BetaSheet1452), and achieves F1-scores of ~62% and ~76% at the residue level and strand level, respectively. Taking the prediction of the more advanced RaptorX-Contact as input, RDb₂C achieves impressively higher performance, with F1-scores reaching ~76% and ~86% at the residue level and strand level, respectively. According to our tests on 61 mainly β proteins, improvement in the β - β contact prediction can further ameliorate the structural prediction.

Availability:

All source data and codes are available at <http://166.111.152.91/Downloads.html> or at the GitHub address of <https://github.com/wzmao/RDb2C>.

Author summary

Due to the topological complexity, mainly β proteins are challenging targets in protein structure prediction. Knowledge of the pairing between β strands, especially the residue pairing pattern, can greatly facilitate the tertiary structure prediction of mainly β proteins. In this work, we developed a novel algorithm to identify the residue pairing in β strands from a predicted residue contact map. This method adopts the ridge detection technique to capture the characteristic pattern of β - β interactions from the map and then utilizes a multi-stage random forest framework to predict β - β contacts at the residue level. According to our tests, our method could effectively improve the prediction of β - β contacts even from a highly noisy contact map. Moreover, the refined β - β contact information could effectively improve the structural modeling of mainly β proteins.

Introduction

Since Anfinsen's dogma [1] was firstly introduced, prediction of the tertiary structures of proteins has become the Holy Grail in structural bioinformatics. Although practical tertiary structure prediction generally requires intensive sampling in the conformational space, the computational consumption could be greatly alleviated with the knowledge of residue contacts in the native conformation. Consequently, protein residue contact prediction has attracted more and more attention, particularly with the significant improvement of prediction accuracy in recent years [2, 3]. Theoretically, native residue contacts that are essential for protein structure or function could be inferred from correlated mutations of residue pairs in evolution. With sequence

data accumulated at an unprecedented speed, extraction of such coevolution information from multiple sequence alignment (MSA) has become more and more practicable [4-7].

Many early residue contact prediction methods were derived from statistics and information theory, like OMES [8], MI [9], MIp [10] and SCA [11]. However, these methods ignore the transitive correlation between residues and thus generate many false positive results. The inverse covariance matrix and pseudo-likelihood maximization were introduced subsequently to eliminate transitivity in methods such as DCA [12], PSICOV [13], plmDCA [14], GREMLIN [15], CCMpred [16], FreeContact [17] and PconsC2 [18]. These methods effectively reduce false positive predictions by globally considering all inter-residue correlations. More recently, methods like MetaPSICOV [19], SAE-DNN [20], DeepConPred [21], NeBcon [22] and RaptorX-Contact [23] integrated sophisticated machine-learning techniques to further enhance the prediction accuracy. In the latest CASP12 competition, RaptorX-Contact achieved the best performance in the category of template-free modeling targets.

In spite of the general improvement, none of existing methods can attain a robust and steady prediction among all protein targets, mainly because the reliability of coevolution information is guaranteed only when a sufficiently large number of homologous sequences are present in the MSA. Indeed, many protein families lack enough homologous sequences for reliable inference of residue contacts [21], and the predicted residue contact maps of these targets may be dominated by false positives, which hinders the subsequent protein structure prediction/modeling. However, even in the highly noisy residue contact maps for these small-family protein targets, characteristic patterns of specific structural motifs could be identified, because a collective pattern of multiple residue contacts is less likely to be perturbed by individual prediction errors and therefore could be more reliably identified than a single residue contact. Good exemplar structural motifs include parallel and anti-parallel β strands, where consecutive residue pairs from individual β strands establish repetitive contacts in the diagonal and off-diagonal directions on a residue contact map, respectively. Hence, it is possible to identify the residue pairing in interacting β strands from a predicted residue contact map. Identification of β - β pairing would greatly benefit the structural prediction of mainly β proteins, a group of challenging protein targets with complicated topologies. Arguably, structural models of mainly β proteins are reported to be less accurate than those of mainly α proteins, when constructed from residue contact information with comparable levels of accuracies [24].

A great variety of β - β pairing prediction methods have been developed since 1990s [25], including BetaPro [26], MLN/MLN-2S [27], CMM [28] and BCov [29]. Among these methods, the more recent ones, CMM and BCov, make predictions based on coevolution features extracted from the sequence data. Unfortunately, all these previous methods are constructed with the knowledge of native secondary structures and therefore perform unsatisfyingly when fed with predicted secondary structures, which limits their usefulness in practical protein structure prediction. As the first pure predictor modeled without any native structural information, bbcontacts [30] utilizes hidden Markov models to identify β - β pairing from the residue contact map predicted by CCMpred and exhibits a remarkable improvement in performance over all previous algorithms.

Here, we proposed a new approach to predict β - β pairing using ridge detection, a conception that has been well-developed in image processing to capture the axis of an elongated object. Ridge detection was firstly proposed by Haralick [31] in 1983, and was then applied to medical image

analysis by Pizer and his co-workers [32, 33]. Lindeberg introduced γ -normalized derivatives and scale-space ridges [34] to better depict the detailed feature of a ridge.

Unlike bbcontacts, in this work, we treated the predicted residue contact map as a raw image and employed the ridge detection to characterize the pattern of consecutive residue contacts for interacting β strands. We designed a multi-stage random forest framework to integrate all ridge-related properties and a number of additional features to predict the β - β contacts. Starting from contact maps predicted by CCMpred [16], our algorithm RDb₂C (Ridge-Detection-based β - β Contact predictor) shows significant improvements over bbcontacts at both residue and strand levels. Moreover, when connected with the more advanced residue contact predictor RaptorX-Contact [23], RDb₂C reaches an impressively high level of prediction powers, and the improvement in β - β contact prediction further ameliorates the structure prediction of mainly β proteins.

Results and Discussion

Brief introduction of the model

Theoretically, consecutive residue pairs from interacting β strands should present continuous contact points in the diagonal or off-diagonal directions on a native contact map. Even when disguised by prediction noises, the relative strong signals from these β - β contacts are likely to exhibit continuous elongated distributions on a predicted contact map. Here, we adopted the ridge detection, a computer algorithm to identify elongated objects on a 2D image, to capture the characteristic pattern of β - β interactions from predicted contact maps. The ridge information was extracted using the γ -normalized ridge detection method introduced by Lindeberg [34].

Given the original predicted contact map and extracted ridge information, we then developed a novel multi-stage random forest framework to further refine the prediction of β - β contacts. Fig 1 shows the general architecture of the whole algorithm. RDb₂C starts from a residue contact map predicted based on the amino acid sequence of the target protein, e.g. by CCMpred or by RaptorX-Contact. Besides ridge features, general properties of the input contact map and position of the target residue pair within the map are abstracted as map property features and position features, respectively. The predicted secondary structure probabilities (from DeepCNF [35]) are incorporated as additional features. All features are fed into a 3-stage random forest framework to predict residue pairing in interacting β strands.

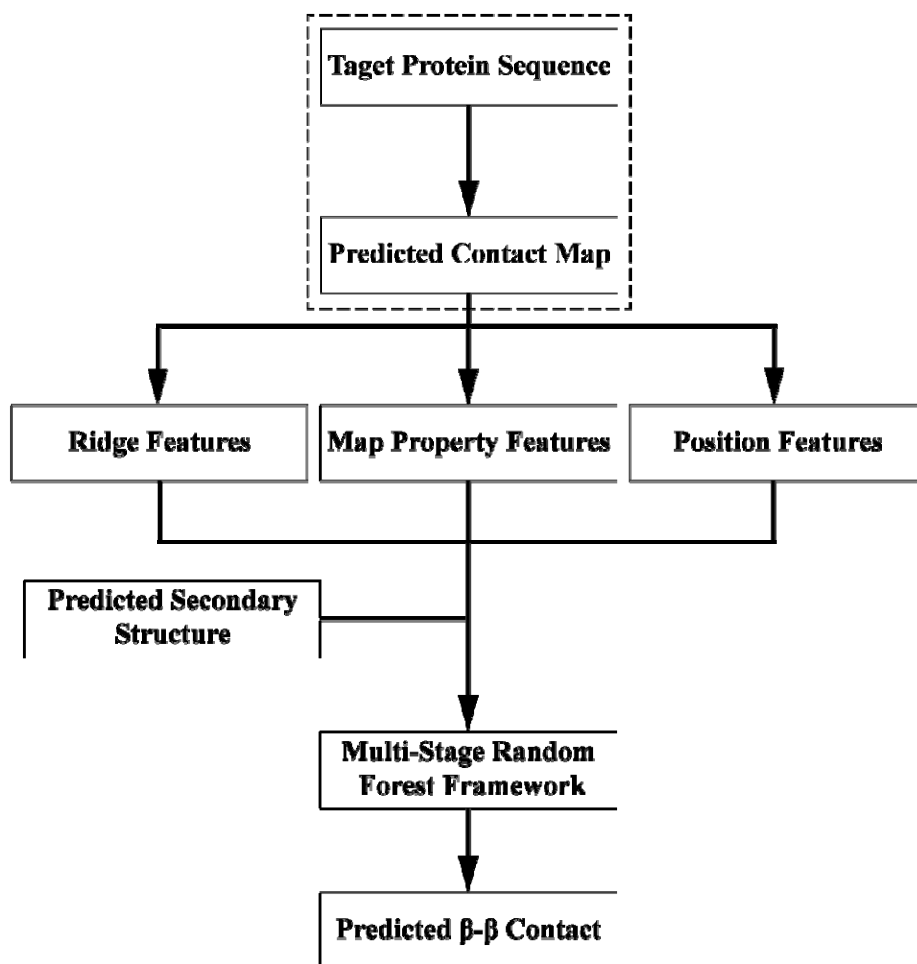


Fig 1. The general flow chart of RDb₂C.

Specifically, we constructed 4 random forest models with different window sizes (3×3, 5×5, 7×7 and 9×9) at the first stage. The prediction results of the first stage models were then combined in the second stage and further optimized in the third stage by taking the preceding-stage results as input features. The model optimization of each stage was performed using 5-fold cross-validation on a training set containing 493 proteins. Further testing and performance evaluation were conducted on two well-established datasets, BetaSheet916 [26] and BetaSheet1452 [29]. Notably, redundancy between the training and test datasets has been carefully removed.

Performance evaluation of the model

The performance of RDb₂C models at all stages was evaluated in the cross-validation as well as the BetaSheet916 and BetaSheet1452 test sets. Table 1 summarizes the residue-level performance. Clearly, all models show robust and balanced performance between the two independent test sets, which indicates appropriate model training. It is noticeable that cross-validation exhibits lower F1-scores than the test sets. This difference may be attributed to the presence of more small-family proteins in the training set than in the test sets (Fig 2): 18.05% of the training set proteins have less than L sequences in the MSA (L is the protein length),

whereas the percentage reduces to only 7.21% and 1.31% in the BetaSheet916 and BetaSheet1452 sets, respectively.

Table 1. Residue-level F1-scores of all models in the 5-fold cross-validation, BetaSheet916 and BetaSheet1452 sets.

Evaluation		1 st stage	2 nd stage	3 rd stage
Cross-validation	3×3	44.40%		
	5×5	45.44%		
	7×7	44.80%	55.08%	55.87%
	9×9	44.30%		
BetaSheet916	3×3	49.41%		
	5×5	50.58%	60.17%	61.19%
	7×7	49.86%		
	9×9	48.80%		
BetaSheet1452	3×3	49.92%		
	5×5	50.97%	61.09%	62.38%
	7×7	50.18%		
	9×9	49.10%		

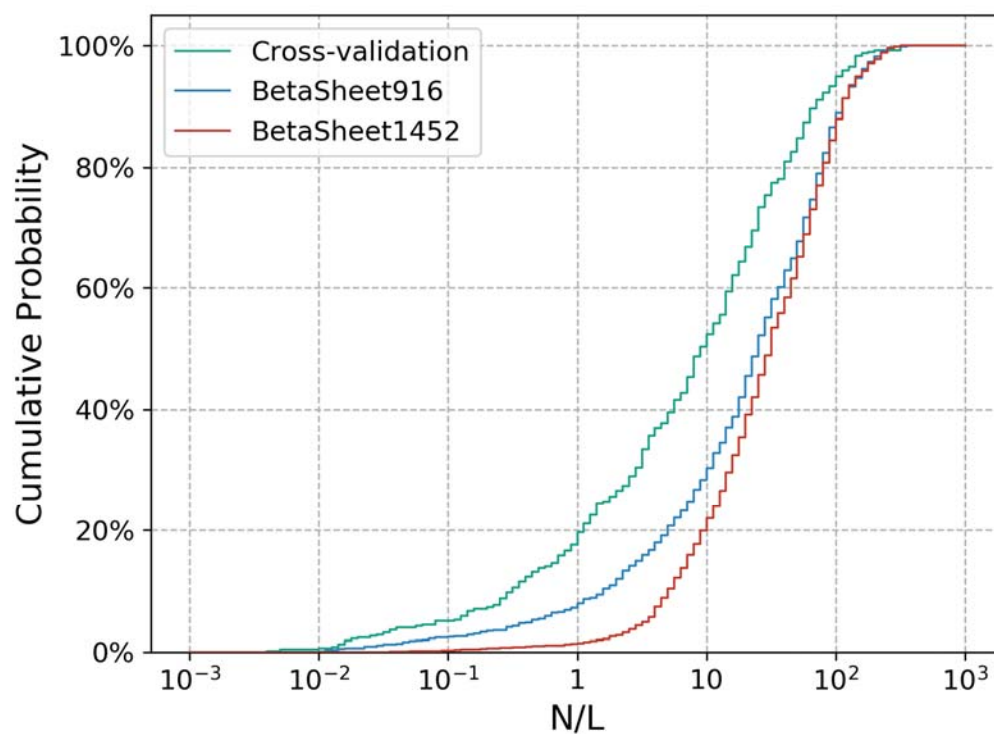


Fig 2. The cumulative distributions for training and test sets with the respect of N/L . N is the number of sequences in the MSA and L is the protein length. There are more proteins in the training set with limited numbers of homologous sequences ($N/L < 1$) than in the BetaSheet916 and BetaSheet1452 sets.

The first-stage models attain the optimal performance at the window size of 5 in both cross-validation and test sets. We suspect that the larger windows include more useful information but also introduce more noises that eventually impairs the model performance, and that balance of useful information and noise may be achieved at the window size of 5. However, models constructed at various window sizes could provide complementary information. Accordingly, the second-stage models that combine information achieved at all window sizes exhibit significant improvement (~10 percentage points) in F1-scores over the first-stage ones. At the third stage, further optimization slightly improves the F1-score to 61.19% and 62.38% in the BetaSheet916 and BetaSheet1452 sets, respectively.

To justify the effectiveness of novel features we proposed in this work, we evaluated the feature importance for all first-stage models. The feature importance was evaluated by re-conducting the model optimization and cross-validation without the corresponding features. As shown in Table 2, all features are essential for the model, since removal of each type weakens the performance. Moreover, all first-stage models exhibit a uniform trend: the ridge features and the original CCMpred map jointly make the major contribution to the prediction power (see the loss of >20 percentage points after removal of both features). Although the ridge features are derived from the CCMpred map, removing ridge features alone significantly deteriorates the F1-score, especially for models of small window sizes, possibly because these features are capable of summarizing the local information and depicting the local shape character of a predicted contact map. Therefore, the ridge features introduced in this work effectively capture the residue contact pattern of β - β interactions. In addition, the secondary structure information predicted by DeepCNF is also constructive to our model, which is reasonable considering that proper assignment of β residues are the prerequisite for the prediction of β - β contacts.

Table 2. The feature importance in the first-stage models.

Window size	1 st stage	-Ridge	-CCMpred	-Ridge -CCMpred	-DeepCNF	-Map Features	-Position Features
3×3	44.40%	36.33%	34.64%	14.33%	37.84%	42.75%	43.92%
5×5	45.44%	39.18%	36.81%	17.30%	38.93%	44.27%	44.66%
7×7	44.80%	40.04%	37.35%	19.42%	37.99%	44.24%	44.54%
9×9	44.30%	40.02%	37.22%	21.01%	37.09%	43.31%	43.60%

The table lists F1-scores of the re-conducted cross-validation without the corresponding features. Winner in each category is highlighted in bold.

As expected, when using the native secondary structures assigned by DSSP [36] instead of the predicted ones as input, the DSSP-based models provide improvement of ~10 percentage points to the residue-level predictions (Table 3). Thus, more accurate secondary structure prediction algorithm could further improve the performance potentially. Table 4 summarizes the strand-level performance in the BetaSheet916 and BetaSheet1452 sets. Notably, the strand-level performance was only evaluated using the DSSP-based framework due to the requirement of exact secondary structure information in the assignment of β strands. Similar to residue-level results (see Table 1), the strand-level models are progressively refined with stages, with the final F1-scores reaching 75.40% and 76.55% in the BetaSheet916 and BetaSheet1452 sets, respectively.

Table 3. Residue-level performance of RDb₂C constructed with DeepCNF-predicted and DSSP-assigned secondary structure information.

Secondary structure	Models	BetaSheet916			BetaSheet1452		
		Precision	Recall	F1-score	Precision	Recall	F1-score
Predicted	1 st stage	63.94%	41.84%	50.58%	57.61%	45.71%	50.97%
	2 nd stage	65.03%	55.99%	60.17%	64.50%	58.02%	61.09%
	3 rd stage	68.00%	55.62%	61.19%	67.91%	57.69%	62.38%
DSSP	1 st stage	69.92%	49.94%	58.26%	62.71%	54.22%	58.16%
	2 nd stage	75.79%	64.00%	69.40%	75.74%	66.07%	70.58%
	3 rd stage	76.28%	65.94%	70.74%	76.56%	67.86%	71.95%

Performances of the models with the window size of 5 are listed here as the representatives of the first-stage models. Winner in each category is highlighted in bold.

Table 4. Strand-level F1-scores of all models in the 5-fold cross-validation, BetaSheet916 and BetaSheet1452 sets.

Evaluation		1 st stage	2 nd stage	3 rd stage
Cross-validation	3×3	67.31%		
	5×5	67.39%		
	7×7	66.33%	77.60%	78.80%
	9×9	65.84%		
BetaSheet916	3×3	65.78%		
	5×5	66.80%		
	7×7	67.51%	74.49%	75.40%
	9×9	67.16%		
BetaSheet1452	3×3	64.50%		
	5×5	65.92%		
	7×7	65.93%	75.62%	76.55%
	9×9	65.72%		

Comparison with bbcontacts

Here, we mainly compared RDb₂C with bbcontacts, the best predictor so far among all previous methods. The performance of RDb₂C and bbcontacts could be fairly compared since both methods take CCMpred contact maps as input. Fig 3 presents the Precision-Recall (PR) curves of RDb₂C and bbcontacts at the residue and strand levels in the BetaSheet916 and BetaSheet1452 sets, respectively. At the residue level, RDb₂C outperforms bbcontacts on the whole range, especially in the region of high-Precision values. Specifically, with the sacrifice of Recall, RDb₂C could approach the Precision level of 90-100%, which means that top-scored predictions of RDb₂C are almost error-less and thus can be directly applied to practical structure prediction. In contrast, bbcontacts can only access the Precision level of 70-80%. As for the strand-level results, despite the crossing of PR curves, RDb₂C outperforms bbcontacts in most ranges, particularly at

the high-Precision region that reflects the quality of top-scored predictions.

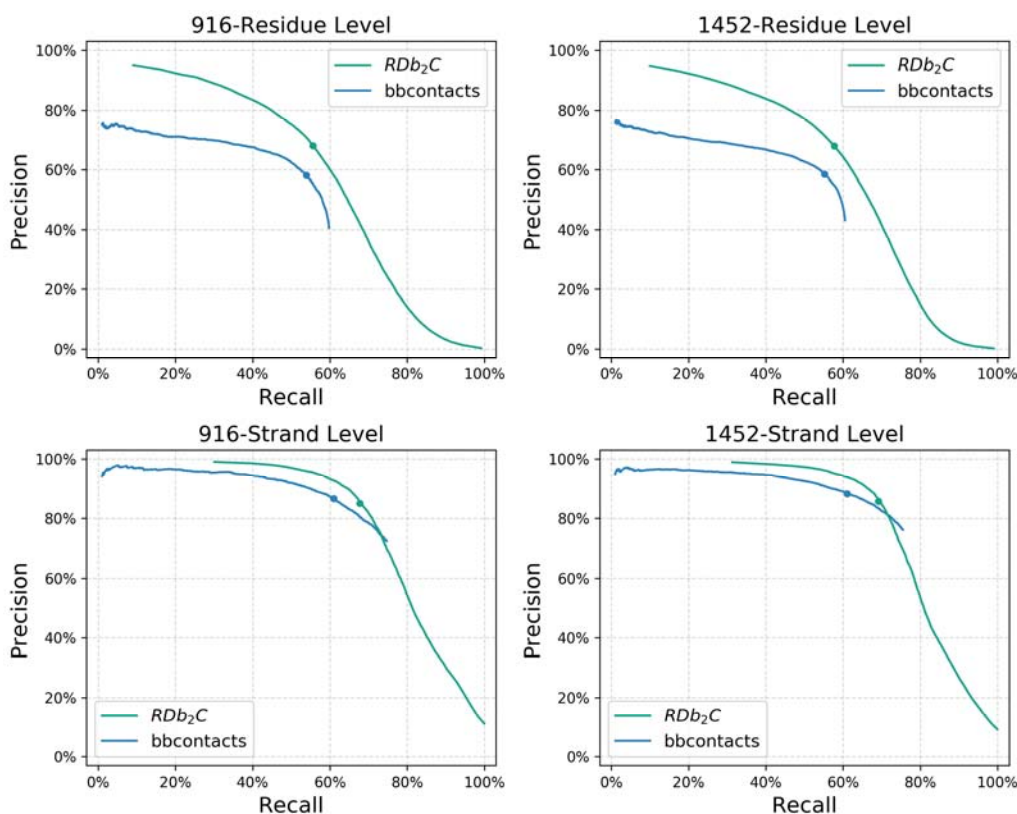


Fig 3. The PR curves in the BetaSheet916 and BetaSheet1452 sets. The comparison is shown for RDb₂C (green) and bbcontacts (blue), at the residue level (top row) and strand level (bottom row) as well as in the BetaSheet916 (left column) and BetaSheet1452 (right column) sets, respectively. Performances at the suggested cutoffs are marked as dots on the PR curves.

Detailed comparison of the two methods at their respective suggested cutoffs is listed in Table 5. Both RDb₂C and bbcontacts are quite robust between the BetaSheet916 and BetaSheet1452 sets. In comparison to the reported numbers in the original paper, performance of bbcontacts increases substantially (residue-level F1-score of ~56% vs. ~50% in the paper), possibly due to the enhanced prediction accuracy of CCMpred with the accumulation of sequence data in the past years. However, RDb₂C still outperforms bbcontacts by ~6 percentage points at the residue level, in terms of F1-scores. At the strand level, RDb₂C and bbcontacts have different preferences of Precision and Recall, but comprehensively RDb₂C achieves a higher level of F1-scores (~76%) and outperforms bbcontacts by ~4 percentage points.

Table 5. Performance comparison of RDb₂C and bbcontacts at residue level and strand level.

Evaluation	Methods	BetaSheet916			BetaSheet1452		
		Precision	Recall	F1-score	Precision	Recall	F1-score
Residue level	RDb ₂ C	68.00%	55.62%	61.19%	67.91%	57.69%	62.38%

	bbcontacts	58.12%	53.87%	55.91%	58.43%	55.16%	56.75%
Strand level	RDb ₂ C	85.01%	67.74%	75.40%	85.69%	69.17%	76.55%
	bbcontacts	86.68%	60.99%	71.60%	88.26%	61.01%	72.14%

Winner in each category is highlighted in bold.

Subsequently, we systematically compared the F1-scores of RDb₂C and bbcontacts for individual proteins in the BetaSheet916 and BetaSheet1452 sets (Fig 4). At the residue level, RDb₂C outperforms bbcontacts on 69.32% targets of the BetaSheet916 set and 72.56% targets of the BetaSheet1452 set, respectively, in terms of F1-scores. The superiority of RDb₂C over bbcontacts is statistically significant (p-value < 10⁻¹⁰) in both test sets. At the strand level, RDb₂C exhibits better performance on 61.57% and 63.36% targets of the BetaSheet916 and BetaSheet1452 sets, respectively, and this advantage is also statistically significant with p-values < 10⁻¹⁰.




Fig 4. Comparison of RDb₂C and bbcontacts for individual proteins of the BetaSheet916 and BetaSheet1452 sets. Each individual protein is represented as a dot. The green dots and blue dots represent targets that are better predicted by RDb₂C and by bbcontacts, respectively, in terms of

F1-scores. The cases are bisected to two methods. In both test sets and at both residue and strand levels, RDb₂C outperforms bbcontacts significantly (p-value < 10⁻¹⁰).

To compare with other previous methods that have reported results only for DSSP-based predictions, we evaluated the DSSP-based models for RDb₂C and bbcontacts at the residue level. As shown in Table 6, RDb₂C outperforms bbcontacts by 2-3 percentage points with the knowledge of native secondary structures, while both RDb₂C and bbcontacts remarkably outperform previous methods by large margins.

Table 6. Performance comparison of DSSP-based RDb₂C, bbcontacts and other methods at the residue level.

Methods	BetaSheet916			BetaSheet1452		
	Precision	Recall	F-measure	Precision	Recall	F-measure
RDb ₂ C	76.28%	65.94%	70.74%	76.56%	67.86%	71.95%
bbcontacts	72.39%	65.10%	68.55%	73.17%	65.39%	69.06%
BCov6*	42.40%	43.90%	43.10%	42%	45%	43%
BCov*	40.90%	42.40%	41.60%			
MLN-2S*	47.30%	42.70%	44.90%			
MLN*	46.10%	39.30%	42.40%			
BetaPro*	38.00%	44.10%	40.80%			

Data for BCov6/BCov and MLN-2S/MLN/BetaPro are taken from [29] and [27], respectively. Winner in each category is highlighted in bold.

The advantage of RDb₂C over bbcontacts in models constructed with predicted secondary structures may arise from two facets of differences: 1) different programs adopted for secondary structure prediction (DeepCNF in RDb₂C vs. PSIPRED pipelined with HHSuite in bbcontacts); 2) difference in program design. To test the former point, we first compared the prediction power of DeepCNF and the PSIPRED pipeline used in bbcontacts (Table 7). In all categories, DeepCNF has comparable or slightly weaker prediction power than the PSIPRED pipeline. Furthermore, we tested the bbcontacts model constructed with DeepCNF prediction as input. The DeepCNF-based bbcontacts model achieves residue-level F1-scores of 55.17% and 56.19% in the BetaSheet916 and BetaSheet1452 sets, respectively, nearly indistinguishable with the original PSIPRED-based model (55.91% and 56.75%, respectively). Therefore, the superiority of RDb₂C over bbcontacts is mainly attributed to the unique design of our method, for instance, the application of ridge detection and the novel multi-stage framework.

Table 7. Performance comparison of DeepCNF and PSIPRED in the BetaSheet916 and BetaSheet1452 sets.

Test Set	Method	Secondary structure category	Precision	Recall	F1-score
BetaSheet916	PSIPRED	H	90.3%	85.9%	88.1%
		E	86.8%	78.9%	82.6%
		C	79.3%	86.8%	82.9%
	DeepCNF	H	92.6%	78.8%	85.2%

		E	86.4%	76.9%	81.4%
		C	75.1%	88.9%	81.4%
		H	90.4%	87.2%	88.8%
	PSIPRED	E	87.3%	79.1%	83.0%
		C	79.2%	86.3%	82.6%
BetaSheet1452		H	92.6%	80.4%	86.0%
	DeepCNF	E	87.4%	76.5%	81.6%
		C	74.5%	88.9%	81.0%

In Fig 5, we include three protein cases as examples to show the improvement in the prediction of β - β contacts using RDb₂C and bbcontacts. In these examples, the raw CCMpred maps are dominated by noises, which hinders visual identification of β - β interactions. Although both RDb₂C and bbcontacts are capable of finding signals from the noises, the native β - β contacts could be more successfully identified by RDb₂C, at both residue and strand levels.

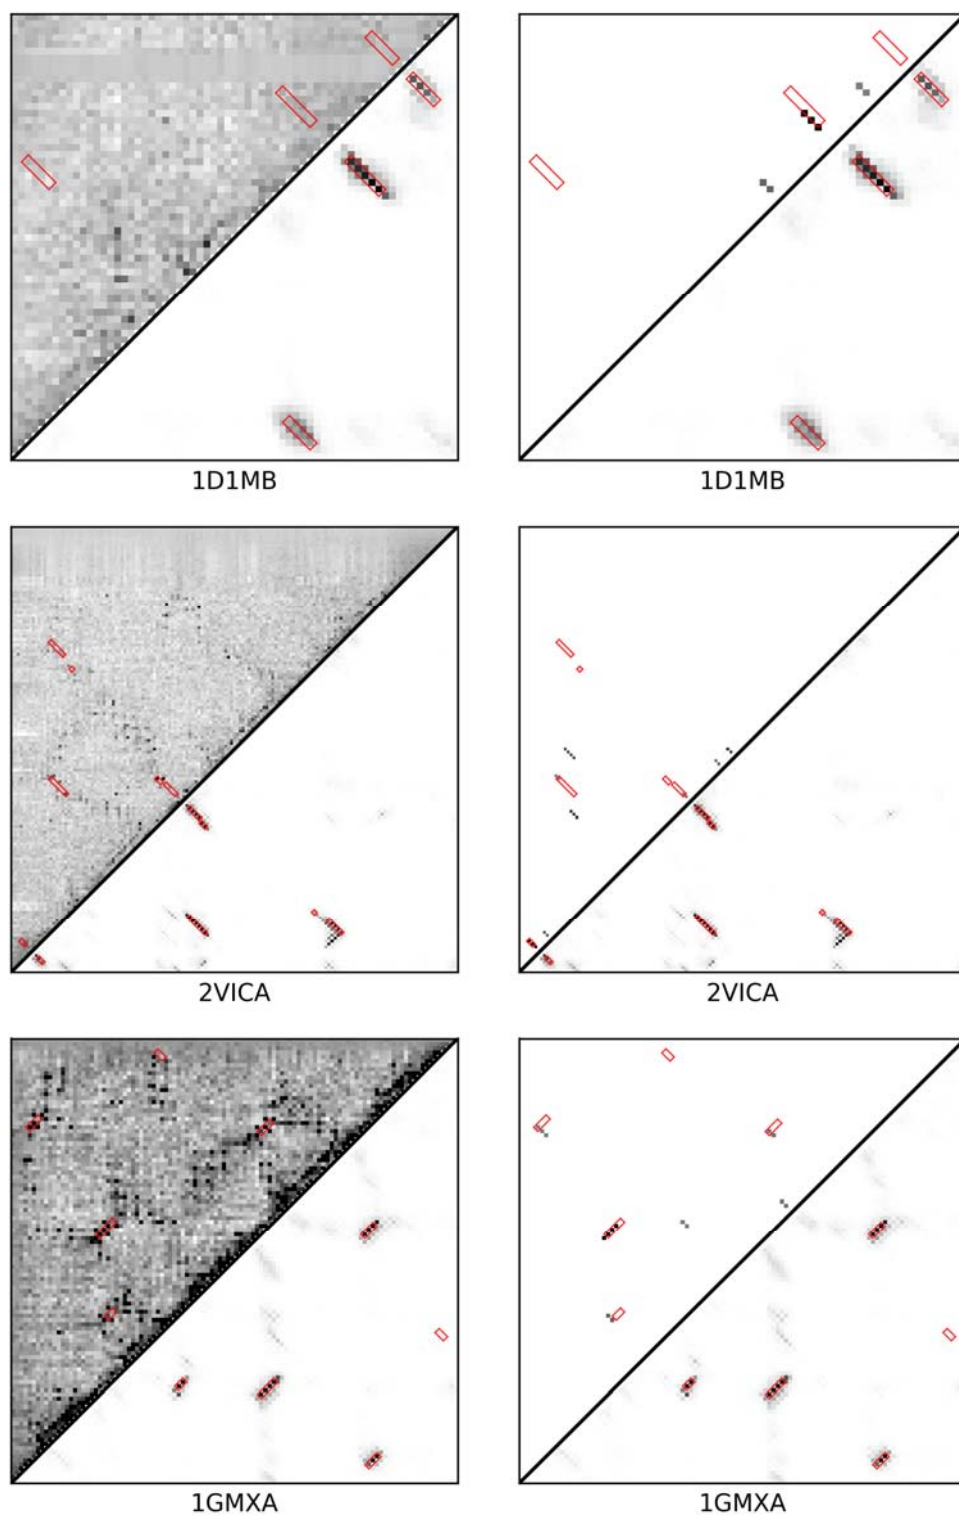


Fig 5. Case studies for CCMpred-based predictions. We illustrate three CCMpred-based case studies. In the left-handed panel, the upper left triangle is the raw CCMpred map, while the lower right triangle is the prediction by RDb₂C. In the right-handed panel, the upper left triangle is replaced by results of bbcontacts to facilitate direct comparison with RDb₂C (i.e. the lower right

triangle). The native β - β contact regions are highlighted by red boxes.

Pipelined with RaptorX-Contact

RDb₂C is developed to refine the prediction of β - β contacts from any predicted contact maps. To verify the general applicability, we tested the performance of our method on contact maps predicted by RaptorX-Contact, one of the most successful residue contact predictors in the latest CASP12 competition. The whole framework was optimized in the same training set, except that the raw maps were obtained from the RaptorX-Contact server. Due to the failure in processing a few protein targets by the server, available proteins in the training set reduces to 383 CATH domains (S1 Table). Considering the time consumption in server submission, this test was conducted only on the BetaSheet916 set. Similarly, the number of available proteins in the BetaSheet916 set was shrunk to 858.

To evaluate the prediction powers of RaptorX-Contact and CCMpred in the β regions, we collected the prediction scores of all pairs of β residues as referred by DSSP assignment. These scores were then sorted and an adjustable cutoff value was used to identify the positive predictions. In this manner, Precision and Recall values at various cutoff values could be collected, which enables the plotting of PR curve as well as the calculation of optimal F1-score. Noticeably, the F1-scores derived in this way may be overestimated, because knowledge of native secondary structures is utilized and because the cutoff is self-optimized rather than estimated independently. Results suggest that RaptorX-Contact provides significantly more accurate residue contact prediction than CCMpred. As for β - β contacts, CCMpred only achieves an F1-score of 20.28%, while RaptorX-Contact attains 60.23%. However, even starting from the poor contact maps of CCMpred, RDb₂C could improve the prediction of β - β contacts to a level comparable to RaptorX-Contact (~61%, see Table 3).

The evaluation of our models optimized on the RaptorX-Contact maps is summarized in Table 8. Unlike previous results (see Table 1), the model performance shows negligible improvement in sequential stages, which indicates that prediction could terminate in early stages when the input residue contact maps are of high quality. Nevertheless, RDb₂C finally reaches impressively high F1-scores of 76.17% and 85.65% at the residue and strand levels, respectively. Notably, performance of these levels could ensure both prediction accuracy (Precision) and coverage of native β - β contacts (Recall) at sufficiently high values (>70%), which thus would greatly benefit the tertiary structure prediction of mainly β proteins.

Table 8. Performance of RDb₂C at residue level and strand level on the 5-fold cross-validation and shrunk BetaSheet916 set.

Level	Stage	Cross-validation	BetaSheet916(858)		
		F1-score	Precision	Recall	F1-score
Residue Level	1 st stage	71.70%	81.02%	71.01%	75.69%
	2 nd stage	72.18%	79.48%	73.47%	76.36%
	3 rd stage	71.89%	78.84%	73.67%	76.17%
Strand Level	1 st stage	82.28%	93.96%	77.94%	85.20%
	2 nd stage	86.80%	95.40%	78.61%	86.20%

3rd stage 88.10% 95.59% 77.57% 85.65%

Performances of the models with the window size of 5 are listed here as the representatives of the first-stage models. Winner in each category is highlighted in bold.

In comparison to CCMpred-based results (see Table 5), F1-scores are improved by ~15 percentage points, which is mainly attributed to the greatly enhanced quality of residue contact map predicted by RaptorX-Contact. As suggested by the evaluation of feature importance (Table 9), ridge features and raw RaptorX-Contact scores in combination still provide major contribution to the prediction power. However, with the remarkable improvement in the quality of the input map, contribution of the individual ridge features becomes less important, when compared with CCMpred-based predictions (see Table 2).

Table 9. The feature importance in the first-stage models starting with RaptorX-Contact predictions.

Window size	1 st stage	-Ridge	-RaptorX	-Ridge -RaptorX	-DeepCNF	-Map Features	-Position Features
3×3	71.51%	71.02%	66.48%	13.04%	70.14%	71.34%	71.30%
5×5	71.70%	71.58%	66.77%	15.75%	70.50%	71.37%	71.37%
7×7	71.50%	71.47%	66.93%	17.95%	70.59%	71.31%	71.21%
9×9	71.43%	71.44%	66.70%	19.80%	70.39%	71.03%	71.08%

The table lists F1-scores of the re-conducted cross-validation without the corresponding features.

On the other hand, RDb₂C is capable of further improving the high-quality contact prediction of RaptorX-Contact. In specific, the F1-score of β - β contacts increases from an estimated number of ~60% to 76.17%. The great improvement by RDb₂C is also illustrated in the PR curves (Fig 6). Considering that knowledge of native secondary structures is required in the generation of RaptorX-Contact curve, we also included the PR curve of the DSSP-based RDb₂C model for a fair comparison. The DSSP-based RDb₂C model could further improve F1-score to 85.30%. Fig 7 shows the comparison of RDb₂C over RaptorX-Contact on two protein cases, where the raw RaptorX-Contact maps are noisy but native β - β contacts could be successfully recognized after refinement using RDb₂C.

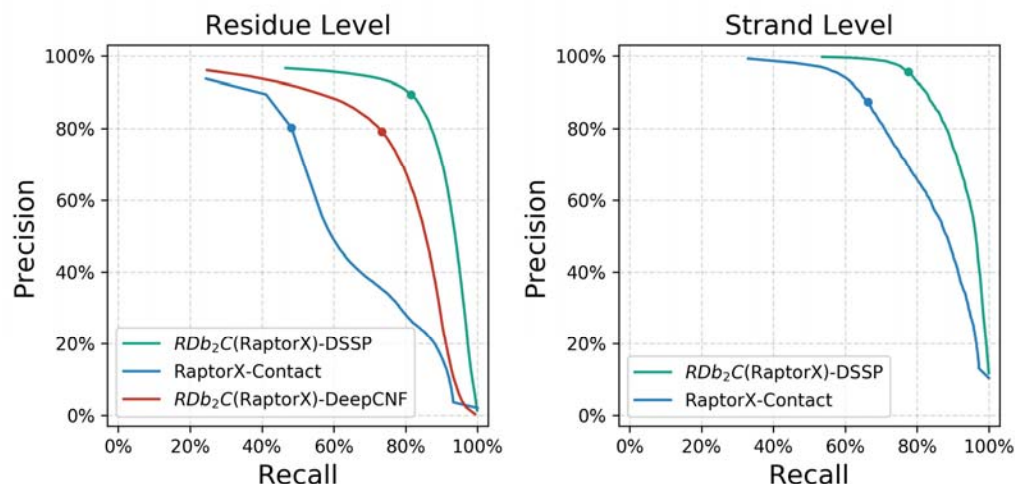


Fig 6. The PR curves in the shrunk BetaSheet916 set. *RDb₂C* (green for DSSP-based model and red for DeepCNF-based model) exhibits significant improvement over the raw RaptorX-Contact prediction (blue). The dots on the PR curve illustrate model performance at the suggested *RDb₂C* cutoffs and the optimized RaptorX-Contact cutoffs.

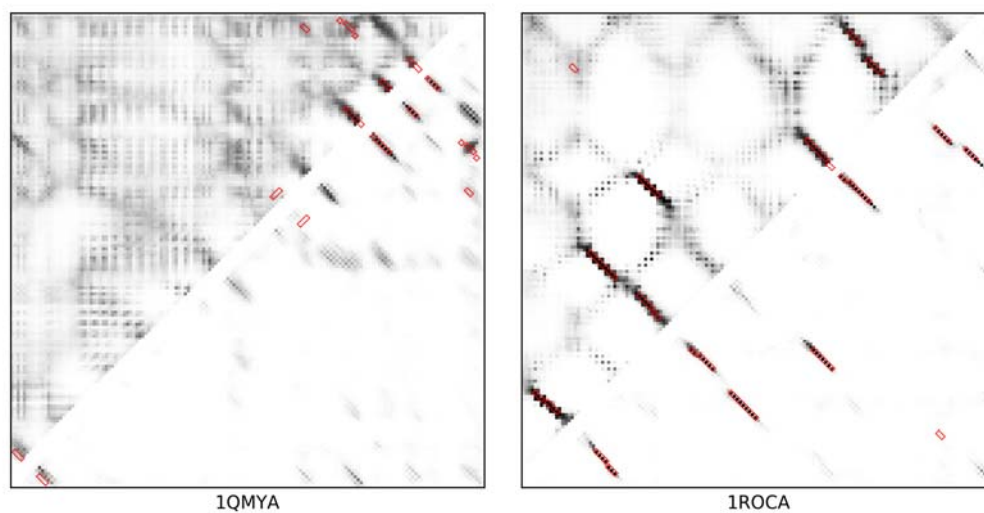


Fig 7. Case studies for RaptorX-Contact-based predictions. We illustrate two RaptorX-Contact-based case studies: 1QMYA (left) and 1ROCA (right). In each plot, the upper left triangle is the raw RaptorX-Contact map, while the lower right triangle is the prediction by *RDb₂C*. The native β - β contact regions are highlighted by red boxes.

Evaluation for the contribution in tertiary structure prediction

In order to justify the effectiveness of our method in the practical structure prediction, we chose 61 mainly β proteins (with $\geq 50\%$ of β residues) from the shrunk BetaSheet916 set (S2 Table) and constructed the tertiary structure models of them with predicted contacts taken as constraints, following the standard CONFOLD protocol [37]. As the numbers of predicted and native β - β

contact pairs are always less than $0.5L$ ($S2$ Table; L is the protein length), which is not sufficient for structural modeling, we retained all β - β contacts predicted by the RDb₂C model in pipeline with RaptorX-Contact at the suggested cutoff as the highly reliable contact pairs, and then enriched the list of contact pairs to $1L$ by collecting the high-ranked and non-redundant RaptorX-Contact predictions. These top $1L$ residue contacts were used as distance restraints to fold the protein. Specifically, a strict restraint of 3.5 - 6\AA was applied to constrain the C_{β} atoms of residue pairs from the more reliable RDb₂C prediction, whereas a loose restraint of 3.5 - 10\AA was adopted for the non-redundant residue pairs enriched from RaptorX-Contact results because of their lower confidence level. As a control, the top $1L$ residue contacts were directly chosen from the RaptorX-Contact prediction and a uniform standard restraint of 3.5 - 8\AA was engaged to constrain the C_{β} atoms of these residue pairs.

For each tested protein, the model with the best TM-score [38] within the top 5 models reported by CONFOLD was chosen for evaluation. According to our results, models constructed with the top $1L$ RaptorX-Contact predictions reach an average TM-score of 0.442 . In contrast, when supplemented with the refined top $1L$ contacts by RDb₂C, the average TM-score markedly increases to 0.506 . Specifically, among the 61 mainly β proteins, prediction using RDb₂C refinement outperforms that using RaptorX-Contact raw scores in 83.61% and 85.25% of cases when evaluated by TM-score and RMSD, respectively (Fig 8 and S2 Table). The superiority of RDb₂C over RaptorX-Contact is statistically significant (p -value $< 10^{-8}$) for both RMSD and TM-score.

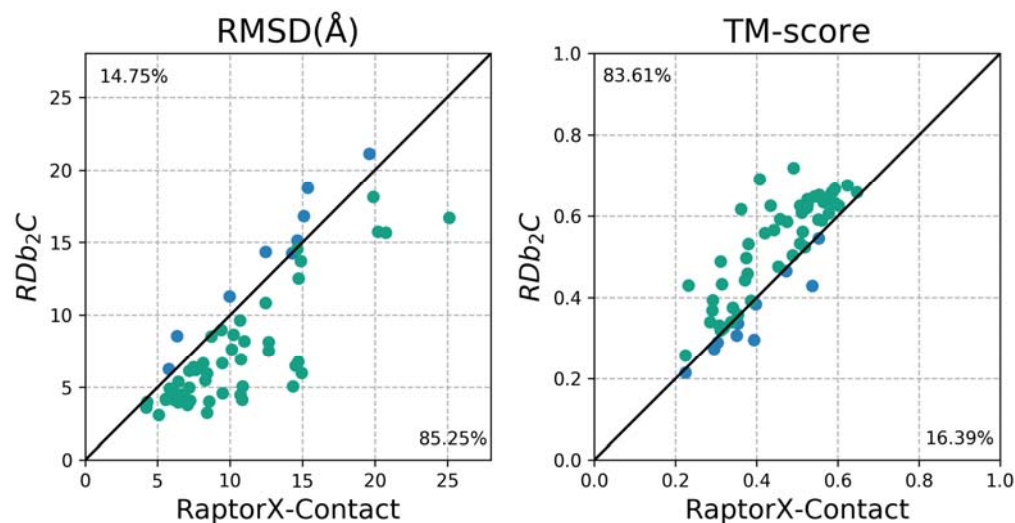


Fig 8. Comparison of the best of the top 5 models generated using the RaptorX-Contact prediction and the RDb₂C refinement for individual targets of the 61 mainly β proteins. The green dots and blue dots represent targets that are better predicted by RDb₂C and by RaptorX-Contact respectively. Detailed results are listed in S2 Table. For both RMSD and TM-score, RDb₂C outperforms RaptorX-Contact significantly (p -value $< 10^{-8}$).

Fig 9 shows the comparison of one protein case, where the RDb₂C results successfully correct the topology mismatch in the RaptorX-Contact model. Because our predictions focus on the more detailed hydrogen bonding interactions, instead of direct use as the distance restraints for

residue C_{β} atoms, it is possible to further improve the structure prediction by utilizing our prediction more delicately, for instance, to restrain the respective hydrogen bonding donors and acceptors of two paired β residues.




Fig 9. Case study for structure prediction. We illustrate the predicted structures of 1OUSB based on the refined predictions by RDb₂C (left) and the raw RaptorX-Contact predictions (right), respectively. Comparing to the native structure (blue), the predicted structure based on RDb₂C (orange) has a higher TM-score (0.6172 vs. 0.3612) and smaller RMSD (4.13Å vs. 10.84Å) than the predicted structure based on the raw RaptorX-Contact prediction (red).

Runtime and memory consumption

We evaluated the running time of RDb₂C on a Dell 5810 workstation (Intel Xeon E5-1620 v3 3.50 GHz CPU, 4 cores, 8 threads and 32 GB RAM) with 8 threads, based on the BetaSheet916 set. Time consumption increases with the size of target protein in a quadratic manner (Fig 10). A typical 400-residue protein needs 20 seconds to complete the prediction. The general memory usage is about 6.3GB. Generally speaking, the runtime and memory usage of RDb₂C are acceptable for practical protein structure prediction.




Fig 10. The relationship between runtime and the number of residues. The time consumed increases steadily with the rise of the number of residues (the I/O time is not included).

In conclusion, we developed a ridge-detection-based algorithm with a multi-stage random forest framework to refine the prediction of β - β contacts from a predicted residue contact map. The ridge information could effectively capture the pattern of consecutive residue contacts in interacting β strands. Our method could be pipelined with any residue contact predictors. Tests on CCMpred and RaptorX-Contact suggest that RDb₂C could improve the prediction of β - β contacts for residue contact predictors of various levels of accuracy. Furthermore, improvement in β - β contact prediction could facilitate the structural prediction of mainly β proteins. The runtime and memory of our method are acceptable for practical use.

Materials and Methods

Dataset

We used two well-established datasets for testing: BetaSheet916 [26] and BetaSheet1452 [29]. These two datasets have been widely accepted, thus allowing performance comparison to previous methods. Both datasets were filtered for redundancy. The β residues were defined using DSSP [36], and both β -bridge and extended β -strand residues (B and E in DSSP) were considered as β residues.

We built our training set from the CATH database of protein domain, version 4.1 [39]. Since our work focused on contacts in β strands, only β and α/β domains were considered. In order to eliminate the redundancy between the training set and test sets, we removed all domains from the

training set that belongs to the same CATH fold groups as proteins in the two test sets. The fragmented and overly short (< 30 residues) domains were also discarded. Finally, only domains in the CATH S35 set [40] (a subset of CATH with pairwise sequence identity <35%) were kept to reduce the redundancy inside the training set. Thus, there were 493 domains in our training set (Table 10 and S1 Table).

Table 10. General information of the training and test sets.

Numbers	Training set	BetaSheet916	BetaSheet1452
Proteins	493	916	1452
Residues	73580	187516	361668
β residues	22283	48996	88352
β - β contact residue pairs	13278	31638	56552
β strands	4633	10745	19186
β strand pairs	2678	8172	14241

In the training set, true β contacts were calculated following the DSSP definition with isolated β -bridge pairs ignored. The DSSP assignment was simplified into 3 categories: H, E and C. The secondary structure probabilities were predicted by DeepCNF [35]. The MSAs were built by HHblits [41] against the UniProt20 database [42], from which residue contact maps were then predicted by CCMpred. ProDy [43] was adopted as a package in Python for dealing with PDB files and analyzing protein structures.

Ridge Features

We employed the ridge as a proxy to capture consecutively distributed regions of relatively strong signals. The ridge is an extended concept of a local maximum. In an N dimensional space, a local maximum point should be maximal in all N dimensions, while a ridge describes a continuous curve each point of which is the local maximum in the N-1 dimensional subspace orthogonal to the curve. Fig 11A demonstrates a ridge on a 2D image, where the vertical axis stands for the signal strength. Ridge is a good measure to characterize the central axis of an elongated object, i.e. consecutive residue contacts in interacting β strands on a residue contact map.

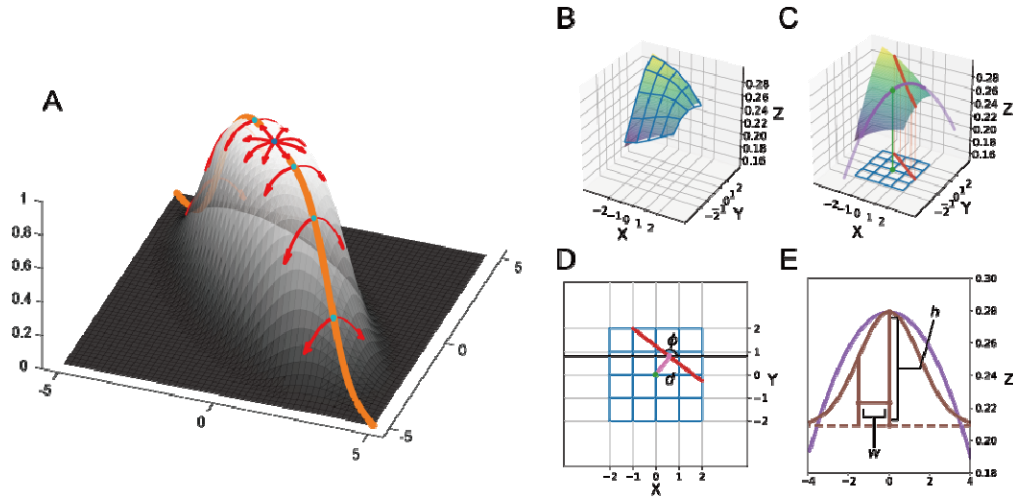


Fig 11. Ridge features from the original map. (A) The orange line indicates the ridge on the 2D function surface. All ridge points on the ridge line are the maxima in the directions perpendicular to the line (red arrows). The local maximum point (dark blue) is also a ridge point based on the definition. (B) For each given point on the contact map, we select local region (i.e. the grid points) to approximate a quadratic function. (C) On the quadratic function surface, we could identify the linear ridge and project it to the XY plane. (D) Direction of the ridge ϕ and distance from the original given point to the ridge d could be obtained from the projection. (E) We could also identify the principal curvature direction on the ridge and approximate the cross section curve with a Gaussian ridge. The height h and width w are defined as the height and the standard deviation of the Gaussian function. Details are given in the S1 Text.

For any given point on the 2D map, we firstly estimated the local 1st order and 2nd order derivatives to build the local gradient $\nabla \mathbf{f}$ and the Hessian matrix \mathbf{H} via an ordinary least squares on the extended surrounding region with the size of 5×5 . Then we calculated the two principal curvatures (λ_p, λ_q) by performing eigendecomposition to the Hessian matrix:

$$\mathbf{H} = \begin{bmatrix} \mathbf{v}_p & \mathbf{v}_q \end{bmatrix} \begin{bmatrix} \lambda_p & 0 \\ 0 & \lambda_q \end{bmatrix} \begin{bmatrix} \mathbf{v}_p & \mathbf{v}_q \end{bmatrix}^{-1}, \quad (1)$$

where $\lambda_p \leq \lambda_q$.

We required at least one principal curvature is negative (i.e. concave) and the directional derivative along the corresponding direction is zero to guarantee the property of ridge points:

$$\begin{aligned} \lambda_p &< 0 \\ \nabla \mathbf{f} \cdot \mathbf{v}_p &= 0 \end{aligned} \quad (2)$$

By locating such points on the contact map, we could identify the axis of the elongated region with relatively strong signals.

However, straightforward ridge detection described as above is not practical on discrete maps for several reasons. Firstly, the ridge could not always locate exactly on a discrete point. Secondly, straightforward method will include all ridges without considering the ridge height or strength. For the first issue, we could roughly locate the ridge position by approximating the neighboring region

with a quadratic function according to the estimated gradient and Hessian matrix (Fig 11B). Under the approximation, the ridge is a straight line (Fig 11C), from which we could identify the direction (\square) and the distance from the original given point (d) in the XY plane (Fig 11D). To solve the second issue, we introduced the γ -normalized scale method developed by Lindeberg [34]. In specific, we utilized the square principal curvature difference (NL), a measure introduced in Lindeberg's work, to quantify the ridge strength:

$$NL = (\lambda_p^2 - \lambda_q^2)^2. \quad (3)$$

Here, we describe the procedure briefly. We smoothed the map with a Gaussian filter at a series of scale σ . However, NL is not guaranteed to reach maxima at the scale of the ridge width. Lindeberg introduced γ -normalized NL to solve this problem. By multiplying σ^λ with a carefully-selected γ , the γ -normalized NL could reach maxima at desired ridge width:

$$NL_\gamma = \sigma^6 (\lambda_p^2 - \lambda_q^2)^2. \quad (4)$$

The γ -normalized scale method could provide an unbiased estimate of the ridge width (w). We further estimate the ridge height (h) via a similar process (Fig 11E). More details of the γ -normalized scale method and the corresponding calculation protocol in processing contact maps could be found in the S1 Text.

Model Features

For a given point on the predicted residue contact map, we calculated the ridge features (including ridge direction \square , distance to the ridge d , ridge height h and ridge width w (see Fig 11B-E)). These features and scores of the input map jointly constitute 5 N×N matrices (Fig 12). We also incorporated the predicted secondary structure probabilities (for H, E and C) from DeepCNF. Furthermore, to describe positions of the target residue pair, we included the difference in indices of the two residues as well as distances of each residue to both ends of the protein in the amino acid sequence as position features. To characterize the quality of the original contact map, we employed the number of homologous sequences in the MSA per residue as well as the standard deviation of prediction scores as map features.




Fig 12. Summary of features adopted in our model. For each target protein with N residues, we have the original CCMpred map with the size of $N \times N$. We calculate the ridge features for each point on the map to get 4 $N \times N$ matrices (2 $N \times N$ matrices after feature selection). In total, we have $N \times N \times 5$ ($N \times N \times 3$ after feature selection) 2D features. The secondary structure prediction from DeepCNF provides an $N \times 3$ 1D feature matrix. In addition, we have 2 map features (the sequence/residue ratio and CCMpred standard deviation) and 5 position features (1 residue index difference and 4 distances to protein ends). The data in this figure were generated from the protein 1AHQA.

Model Training and Feature Selection

We applied a 3-stage random forest framework to predict the β - β residue contacts using all features described previously. All random forest models in all stages were set up with 500 decision trees and were optimized by 5-fold cross-validation using the scikit-learn package [44]. The cross-validation was applied in a protein-wise manner, by which the training set proteins were randomly partitioned into 5 mutually exclusive subsets with roughly the same size. Combinations of four subsets were then iteratively used to train the model and to predict the unselected subset. Since all proteins in the training set were predicted independently, the suggested cutoffs were optimized in the cross-validation. Finally, the whole training set was utilized to train a separate model as the final model for evaluation in the test sets.

At the first stage, in addition to features of the target residue pair, we adopted an adjustable window to consider the effect of neighboring residues. Specifically, 2D features (ridge features and the original contact map) of all residue pairs falling within the square window centered at the focus point were included. Secondary structure features of all residues falling within the 1D windows centered at the two target residues were also extracted. Map property features and

position features were extracted for the target residue pair only, because they were invariant for the target and neighboring residue pairs. We employed various values of the window size (w_s), including 1, 3, 5, 7 and 9, to train multiple random forest models at the first stage. Because of the scarcity of β - β residue contacts, the negative (Neg) samples greatly outnumbered the positive (Pos) ones with a Pos/Neg ratio of about 1:600. To simplify the model training, we under-sampled negative samples at different Pos/Neg ratios from 1:1 to 1:40. The under-sampling was implemented in a protein-wise manner. That is, for each protein, the number of negative samples was specifically set based on the number of positive samples. Based on the cross-validation results (Table 11), improvement in model performance becomes saturated at Pos/Neg ratios of 1:40. Therefore, each random forest model was trained at 1:40 Pos/Neg ratios. At the same time, we noticed that the model with the window size of 1 significantly underperforms models of the other window sizes. Therefore, we selected window sizes of 3, 5, 7 and 9 for further model optimization.

Table 11. The cross-validation F1-scores for different window sizes and Pos/Neg ratios.

Pos/Neg	1×1	3×3	5×5	7×7	9×9
1:1	32.55%	37.41%	37.61%	35.93%	34.96%
1:5	35.31%	41.39%	41.03%	39.42%	38.51%
1:10	36.31%	42.65%	42.00%	40.76%	39.39%
1:20	36.90%	43.41%	43.22%	41.92%	40.60%
1:30	37.20%	44.15%	43.61%	42.67%	41.23%
1:40	37.33%	44.16%	44.02%	42.90%	41.71%

Winner in each category is highlighted in bold. The row of the selected Pos/Neg ratio is shown in shadow.

We performed the feature selection by removing features group by group and re-conducting the 5-fold cross-validation. We found that the ridge width w and the distance from the ridge d are not essential for the model. After removing these two sets of features, only the ridge height h and the direction of the ridge \square were kept as ridge features. Thus, we obtained the optimized feature combination as indicated in Table 2. We further optimized the shape of the window. Because β - β pattern depends on the signals on diagonal and off-diagonal directions, we used the cross-shaped masks with different diagonal width (d_w) besides the square window mask for 2D features (Fig 13). For all window sizes, the best masks were the ones with the diagonal width of 3 (Table 12). Eventually, we chose the models with the diagonal width of 3 as the final ones.

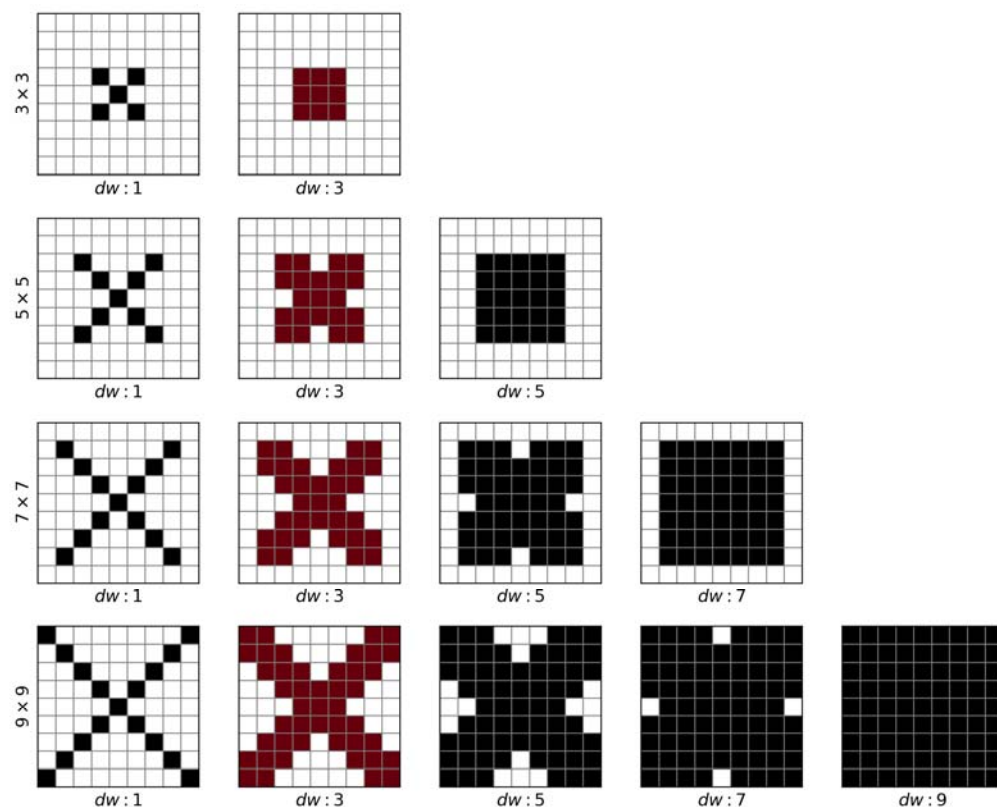


Fig 13. An illustration of the window mask. The selected features are labeled in dark colors. The final window masks that were selected are marked in red.

Table 12. The cross-validation F1-scores for different window sizes and diagonal widths.

Window size	Diagonal width				
	1	3	5	7	9
3×3	41.20%	44.40%			
5×5	41.94%	45.44%	45.27%		
7×7	41.12%	44.80%	44.58%	44.64%	
9×9	40.32%	44.30%	43.95%	43.55%	43.36%

Winner for each window size is highlighted in bold. The column of the selected diagonal width is shown in shadow.

Predictions from the first-stage models were then fed to models at the second stage. In specific, we retained the output scores of the first-stage models as additional 2D features. Unlike the strong constraints applied by bbcontacts that artificially restricts each residue to form no more than two β - β contacts, we included the ranks of each point among the output scores of each column and row and allowed the random forest model to automatically learn the geometry constraints. Hence, output map from each first-stage model provided $N \times N \times 3$ features (1 $N \times N$ raw output and 2 $N \times N$ rankings). Subsequently, we performed the feature selection again as the first stage. The first-stage raw scores, the first-stage rankings, ridge features (ridge height h and ridge direction \square) and predicted secondary structure information by DeepCNF were finally retained

after feature selection (Fig 14). The window size and the diagonal width were both optimized at 3 (3×3 square). Then, we combined features from the 4 first-stage models of various window sizes to construct a comprehensive second-stage random forest model. At the third stage, we carried out a similar protocol as the second stage and obtained a final third-stage random forest model.




Fig 14. An illustration of the multi-stage framework. In our 3-stage framework, we firstly construct models with different window sizes. We then integrate four models to get the second-stage results. The final result is obtained from the third-stage model. The data in this figure were generated from the protein 1AHQA.

The overall framework was constructed for two different types of secondary structure information, prediction from DeepCNF and assignment from DSSP, respectively. For DSSP-based models, the secondary structure probability is set to 1 for the native category and 0 for the others.

Evaluation

Performance was evaluated at both residue and strand levels, using measures including Precision, Recall as well as F1-score. Precision and Recall quantify proportions of true positives within all predicted and all native β - β contacts, respectively, while F1-score is the harmonic mean

of Precision and Recall:

$$\begin{aligned}\text{Precision} &= \frac{TP}{TP + FP} \\ \text{Recall} &= \frac{TP}{TP + FN} \\ \text{F1-score} &= \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}},\end{aligned}\tag{5}$$

where TP , FP and FN denote true positives, false positives and false negatives, respectively.

Although our method was developed with predicted secondary structure information for practical protein structure prediction, we performed evaluation for models fed with predicted and DSSP-assigned secondary structures respectively to simplify comparison with previous methods. Since *bbcontacts* is the best method so far and exhibits significantly superior performance to all previous ones, we mainly compared our method with *bbcontacts*. Results of *bbcontacts* were obtained following the protocol of the original paper, with secondary structures predicted by PSIPRED [45]. The residue-level evaluation is straightforward, while the strand-level evaluation, however, could only be conducted with the knowledge of clearly defined secondary structures. Thus, we only provide the strand-level results for DSSP-based models. As for the definition of strand pairing, we regard a pair of β strands as interacting if at least one pair of residues on the two strands is predicted as contacting.

Structure modeling using predicted contacts

All 61 mainly β proteins (with $\geq 50\%$ of β residues) were chosen from the shrunk BetaSheet916 set (S2 Table), and tertiary structure models of them were constructed with predicted contacts taken as constraints, using the downloadable programs of Crystallography & NMR System (CNS) [46] suite and CONFOLD package [37]. We retained all β - β contacts predicted by the RDb₂C model in pipeline with RaptorX-Contact at the suggested cutoff as the highly reliable contact pairs, and then enriched the list of contact pairs to 1*L* by collecting the high-ranked and non-redundant RaptorX-Contact predictions from the region outside the predicted β - β region of RDb₂C (All contacts falling within the square window covering the RDb₂C prediction points or lines are considered as redundant). These top 1*L* residue contacts were used as distance restraints to fold the protein following the standard CONFOLD protocol, with the DeepCNF results supplemented as predicted secondary structures [37]. A strict restraint of 3.5-6Å was applied to constrain the C _{β} atoms for the more reliable contact pairs of RDb₂C prediction, whereas a loose restraint of 3.5-10Å were adopted for the non-redundant contact pairs enriched from RaptorX-Contact because these complement pairs are of lower confidence levels. In the control experiment, the top 1*L* residue contacts were directly chosen from the RaptorX-Contact results and a uniform standard restraint of 3.5-8Å was engaged to constrain all contact pairs. For each tested protein, 20 models were generated by CONFOLD, and the 5 models that fit the restraints best were retained. The model with the highest TM-score among the top 5 models was then taken as the representative one for evaluation.

Acknowledgements

We gratefully thank Prof. Jinbo Xu for his help in the job submission using the RaptorX-Contact server.

Funding

This work has been supported by the funds from the National Natural Science Foundation of China (#31670723) and from the Beijing Innovation Center of Structural Biology.

Supplementary materials

S1 Text. Technical details of the γ -normalized scale method for ridge detection and the corresponding calculation protocol in processing contact maps.

S1 Table. List of domains in the training set.

S2 Table. Results of structure prediction for 61 mainly β proteins.

References

1. Anfinsen CB. Principles that govern the folding of protein chains. *Science*. 1973;181(4096):223-30. doi: 10.1126/science.181.4096.223.
2. Kinch LN, Li W, Monastyrskyy B, Kryshtafovych A, Grishin NV. Assessment of CASP11 contact-assisted predictions. *Proteins: Structure, Function, and Bioinformatics*. 2016;84(S1):164-80.
3. Monastyrskyy B, D'Andrea D, Fidelis K, Tramontano A, Kryshtafovych A. New encouraging developments in contact prediction: Assessment of the CASP11 results. *Proteins: Structure, Function, and Bioinformatics*. 2016;84(S1):131-44.
4. Göbel U, Sander C, Schneider R, Valencia A. Correlated mutations and residue contacts in proteins. *Proteins: Structure, Function, and Bioinformatics*. 1994;18(4):309-17.
5. Kim DE, DiMaio F, Yu R, Wang R, Song Y, Baker D. One contact for every twelve residues allows robust and accurate topology-level protein structure modeling. *Proteins: Structure, Function, and Bioinformatics*. 2014;82(S2):208-18.
6. Simkovic F, Ovchinnikov S, Baker D, Rigden DJ. Applications of contact predictions to structural biology. *IUCrJ*. 2017;4(3).
7. Simkovic F, Thomas JM, Keegan RM, Winn MD, Mayans O, Rigden DJ. Residue contacts predicted by evolutionary covariance extend the application of ab initio molecular replacement to larger and more challenging protein folds. *IUCrJ*. 2016;3(4):259-70.
8. Kass I, Horovitz A. Mapping pathways of allosteric communication in GroEL by analysis of correlated mutations. *Proteins: Structure, Function, and Bioinformatics*. 2002;48(4):611-7.
9. Gloor GB, Martin LC, Wahl LM, Dunn SD. Mutual information in protein multiple sequence alignments reveals two classes of coevolving positions. *Biochemistry*. 2005;44(19):7156-65.
10. Dunn SD, Wahl LM, Gloor GB. Mutual information without the influence of phylogeny or entropy dramatically improves residue contact prediction. *Bioinformatics*. 2008;24(3):333-40.
11. Halabi N, Rivoire O, Leibler S, Ranganathan R. Protein sectors: evolutionary units of three-dimensional structure. *Cell*. 2009;138(4):774-86.
12. Morcos F, Pagnani A, Lunt B, Bertolino A, Marks DS, Sander C, et al. Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proceedings of the National Academy of Sciences*. 2011;108(49):E1293-E301.
13. Jones DT, Buchan DW, Cozzetto D, Pontil M. PSICOV: precise structural contact prediction using sparse inverse covariance estimation on large multiple sequence alignments. *Bioinformatics*. 2012;28(2):184-90.
14. Ekeberg M, Lövkvist C, Lan Y, Weigt M, Aurell E. Improved contact prediction in proteins: using pseudolikelihoods to infer Potts models. *Physical Review E*. 2013;87(1):012707.
15. Kamisetty H, Ovchinnikov S, Baker D. Assessing the utility of coevolution-based residue-residue contact predictions in a sequence-and structure-rich era. *Proceedings of the National Academy of Sciences*. 2013;110(39):15674-9.
16. Seemayer S, Gruber M, Söding J. CCMpred—fast and precise prediction of protein residue-residue contacts from correlated mutations. *Bioinformatics*. 2014;30(21):3128-30.
17. Kaján L, Hopf TA, Kalaš M, Marks DS, Rost B. FreeContact: fast and free software for protein contact prediction from residue co-evolution. *BMC bioinformatics*. 2014;15(1):85.
18. Skwark MJ, Raimondi D, Michel M, Elofsson A. Improved contact predictions using the recognition of protein like contact patterns. *PLoS computational biology*. 2014;10(11):e1003889.

19. Jones DT, Singh T, Kosciolk T, Tetchner S. MetaPSICOV: combining coevolution methods for accurate prediction of contacts and long range hydrogen bonding in proteins. *Bioinformatics*. 2015;31(7):999-1006. doi: 10.1093/bioinformatics/btu791.
20. Du T, Liao L, Wu CH, Sun B. Prediction of residue-residue contact matrix for protein-protein interaction with Fisher score features and deep learning. *Methods*. 2016;110:97-105.
21. Xiong D, Zeng J, Gong H. A deep learning framework for improving long-range residue-residue contact prediction using a hierarchical strategy. *Bioinformatics*. 2017;33(17):2675–83.
22. He B, Mortuza S, Wang Y, Shen H-B, Zhang Y. NeBcon: protein contact map prediction using neural network training coupled with naïve Bayes classifiers. *Bioinformatics*. 2017;33(15):2296–306.
23. Wang S, Sun S, Li Z, Zhang R, Xu J. Accurate De Novo Prediction of Protein Contact Map by Ultra-Deep Learning Model. *PLoS Computational Biology*. 2017;13(1):e1005324.
24. Michel M, Hayat S, Skwark MJ, Sander C, Marks DS, Elofsson A. PconsFold: improved contact predictions improve protein models. *Bioinformatics*. 2014;30(17):i482-i8.
25. Hubbard TJ, editor Use of β -strand Interaction Pseudo-Potentials in Protein Structure Prediction and Modeling. Proceedings of the Hawaii International Conference on System Sciences; 1994: IEEE INSTITUTE OF ELECTRICAL AND ELECTRONICS.
26. Cheng J, Baldi P. Three-stage prediction of protein β -sheets by neural networks, alignments and graph algorithms. *Bioinformatics*. 2005;21(suppl 1):i75-i84.
27. Lippi M, Frasconi P. Prediction of protein β -residue contacts by Markov logic networks with grounding-specific weights. *Bioinformatics*. 2009;25(18):2326-33.
28. Burkoff NS, Várnai C, Wild DL. Predicting protein β -sheet contacts using a maximum entropy-based correlated mutation measure. *Bioinformatics*. 2013;29(5):580-7.
29. Savojardo C, Fariselli P, Martelli PL, Casadio R. BCov: a method for predicting β -sheet topology using sparse inverse covariance estimation and integer programming. *Bioinformatics*. 2013;29(24):3151–7.
30. Andreani J, Söding J. bbcontacts: prediction of β -strand pairing from direct coupling patterns. *Bioinformatics*. 2015;31(11):1729-37.
31. Haralick RM. Ridges and valleys on digital images. *Computer Vision, Graphics, and Image Processing*. 1983;22(1):28-38.
32. Gauch JM, Pizer SM. Multiresolution analysis of ridges and valleys in grey-scale images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 1993;15(6):635-46.
33. Eberly D, Gardner R, Morse B, Pizer S, Scharlach C. Ridges for image analysis. *Journal of Mathematical Imaging and Vision*. 1994;4(4):353-73.
34. Lindeberg T, editor Edge detection and ridge detection with automatic scale selection. *Computer Vision and Pattern Recognition, 1996 Proceedings CVPR'96, 1996 IEEE Computer Society Conference on*; 1996: IEEE.
35. Wang S, Weng S, Ma J, Tang Q. DeepCNF-D: predicting protein order/disorder regions by weighted deep convolutional neural fields. *International journal of molecular sciences*. 2015;16(8):17315-30.
36. Kabsch W, Sander C. Dictionary of protein secondary structure: pattern recognition of hydrogen-bonded and geometrical features. *Biopolymers*. 1983;22(12):2577-637.
37. Adhikari B, Bhattacharya D, Cao R, Cheng J. CONFOLD: residue-residue contact-guided ab initio protein folding. *Proteins: Structure, Function, and Bioinformatics*. 2015;83(8):1436-49.
38. Zhang Y, Skolnick J. Scoring function for automated assessment of protein structure template

- quality. *Proteins: Structure, Function, and Bioinformatics*. 2004;57(4):702-10.
39. Sillitoe I, Lewis TE, Cuff A, Das S, Ashford P, Dawson NL, et al. CATH: comprehensive structural and functional annotations for genome sequences. *Nucleic acids research*. 2015;43(D1):D376-D81.
 40. Dawson NL, Lewis TE, Das S, Lees JG, Lee D, Ashford P, et al. CATH: an expanded resource to predict protein function through structure and sequence. *Nucleic acids research*. 2016;45(D1):D289-D95.
 41. Remmert M, Biegert A, Hauser A, Söding J. HHblits: lightning-fast iterative protein sequence searching by HMM-HMM alignment. *Nature methods*. 2012;9(2):173-5.
 42. Consortium U. UniProt: the universal protein knowledgebase. *Nucleic acids research*. 2017;45(D1):D158-D69.
 43. Bakan A, Dutta A, Mao W, Liu Y, Chennubhotla C, Lezon TR, et al. Evol and ProDy for bridging protein sequence evolution and structural dynamics. *Bioinformatics*. 2014;30(18):2681-3.
 44. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*. 2011;12(Oct):2825-30.
 45. Jones DT. Protein secondary structure prediction based on position-specific scoring matrices. *Journal of molecular biology*. 1999;292(2):195-202.
 46. Brünger AT, Adams PD, Clore GM, DeLano WL, Gros P, Grosse-Kunstleve RW, et al. Crystallography & NMR system: a new software suite for macromolecular structure determination. *Acta Crystallographica Section D: Biological Crystallography*. 1998;54(5):905-21.