

Constructing Structure Ensembles of Intrinsically Disordered Proteins from Chemical Shift Data[†]

Huichao Gong¹, Sai Zhang¹, Jiangdian Wang², and Haipeng Gong^{3,4},
and Jianyang Zeng^{1,4,*}

¹ Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, 100084, China

² Biostatistics and Research Decision Sciences — Asia Pacific, Merck Research Laboratory, Beijing, 100015, China

³ School of Life Sciences, Tsinghua University, Beijing, 100084, China

⁴ MOE Key Laboratory of Bioinformatics, Tsinghua University, Beijing, 100084, China

Abstract. Modeling the structural ensemble of intrinsically disordered proteins (IDPs), which lack fixed structures, is essential in understanding their cellular functions and revealing their regulation mechanisms in signaling pathways of related diseases (e.g., cancers and neurodegenerative disorders). Though the ensemble concept has been widely believed to be the most accurate way to depict 3D structures of IDPs, few of the traditional ensemble-based approaches effectively address the *degeneracy problem* which occurs when multiple solutions are consistent with experimental data and is the main challenge in the IDP ensemble construction task. In this paper, based on a predefined conformational library, we formalize the structure ensemble construction problem into a least squares framework, which provides the optimal solution when the data constraints outnumber unknown variables. To deal with the degeneracy problem, we further propose a regularized regression approach based on the elastic net technique with the assumption that the weights to be estimated for individual structures in the ensemble are sparse. We have validated our methods through a reference ensemble approach as well as by testing the real biological data of three proteins, including alpha-synuclein, the translocation domain of Colocin N and the K18 domain of Tau protein.

1 Introduction

Unlike traditional ordered proteins, which generally take a well-defined 3D structure to perform their functions, *intrinsically disordered proteins* (IDPs) lack ordered or fixed 3D structures [20, 8, 9]. However, IDPs usually play important roles in essential biological processes and are generally associated with many diseases, such as cancers [20] and neurodegenerative disorders [22]. Therefore, modeling the atomic structural details of IDPs is critical for understanding their cellular functions and revealing their regulation mechanisms in signaling pathways of related diseases.

Ensemble modeling has been believed to be the most accurate way to describe the 3D structures of IDPs [6, 12, 13]. In general, a structure ensemble consists of two parts: a set of 3D structures and the corresponding weights that describe the likelihoods of

[†]This work was supported in part by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the National Natural Science Foundation of China Grant 61033001, 61361136003 and 61472205, and China's Youth 1000-Talent Program.

*Corresponding author: Jianyang Zeng, zengjy321@tsinghua.edu.cn.

individual conformations. When experimental data (e.g., chemical shifts, nuclear Overhauser effect (NOE) distances and residual dipolar couplings) are available, the 3D structures and their weights are often constructed to optimally match these experimental observations. In the literature, numerous approaches have been proposed to construct 3D structure ensembles of IDPs from experimental data [21, 25, 24, 17, 12]. These approaches can be roughly divided into two categories [13]. The first one is called the *ensemble-restrained molecular dynamics* (MD) simulation, which explicitly incorporates available data restraints in the potential function and enforces the whole ensemble to satisfy these experimental constraints during the simulation [1]. In principle, these ensemble-restrained approaches require a large number of constraints to drive accurate MD simulation. On the other hand, in practice, such a large number of constraints are rarely available from experimental data for an IDP. As discussed in [17], such ensemble construction approaches usually fail when there is an insufficient number of constraints available for supporting the restrained MD simulation. The second group of ensemble construction approaches [6, 25, 12, 5] applies a different strategy. They first generate an initial structure pool (which is also called the *predefined conformational library*) using MD simulation, and then select a subset of these structures to represent conformational space. After that, the corresponding weights of individual selected structures are computed to match experimental observations. In this strategy, when constructing the IDP ensemble, stochastic sampling algorithms, such as Monte Carlo, simulated annealing and evolutionary algorithms, are often applied to stochastically assign random weights to the selected structures in order to achieve the best possible solutions [13]. Unfortunately, these stochastic approaches cannot provide any theoretical guarantee to find the global optimal solutions within limited simulation time.

One main challenge in the IDP ensemble construction task lies in addressing the *degeneracy problem* [12, 35]: due to the large degrees of freedom in this problem, there can exist multiple ensemble solutions to satisfy experimental data. This problem remains a challenge even after the initial conformational library has been defined. The degeneracy problem in IDP ensemble construction has been well discussed in [12] and [13].

In this paper, we propose two effective methods to accurately construct the structure ensembles of IDPs. Our approaches follow the strategy of constructing a predefined conformational library, but use more elegant algorithms to compute the optimal weights of the selected structures to satisfy experimental constraints. In particular, when the constraints derived from experimental data outnumber the unknown variables (i.e., the weights of individual structures that need to be estimated), we use a least squares method to compute the optimal solution. Note that a similar least squares method has been proposed to characterize RNA ensembles in a recent study [14]. However, [14] did not handle the situation in which the unknown variables outnumber the constraints derived from experimental data. In the second case (i.e., the number of experimental constraints is smaller than the number of the selected structures), the problem becomes ill-posed and underdetermined, which is the main cause of the degeneracy problem. To address this problem, we assume that the weights to be estimated are sparse, that is, most of them are with zero values. With this reasonable assumption, we propose a regularized regression approach based on the elastic net technique [39] to derive the optimal

weights that best interpret experimental data. In this paper, we mainly use chemical shifts as experimental data to drive the ensemble construction process. Chemical shifts provide strong indicators about local chemical environment in a protein structure, and have been widely used for protein structure modeling [4, 31]. Notably, our methods can be easily extended to incorporate other nuclear magnetic resonance (NMR) data, such as NOEs and RDCs.

Unlike other existing ensemble construction approaches, which stochastically compute the approximate weights of the structures in the predefined conformational library, our least squares method can find the global optimal solution in closed-form. When only a small number of experimental restraints are available, our elastic net based regression approach provides a new framework with reasonable sparsity assumption for addressing the degeneracy problem in the IDP ensemble construction process. To our knowledge, our work provides the first framework to estimate the best solution by applying the least squares method and the elastic net regression approach respectively, depending on whether the experimental data are sufficient or not. We have validated the performance of our algorithms using a reference ensemble approach [23], in which simulated data are used to verify the ensemble construction results. In addition, we have tested our methods on real biological data of three proteins, including alpha-synuclein, the translocation domain of Colicin N and the K18 domain of Tau protein. The test results have demonstrated that our methods can be effectively used to construct accurate structure ensembles of IDPs using chemical shift data.

2 Methods

2.1 Overview

Our goal is to construct an ensemble of representative structures for a given IDP, and to derive the accurate weights associated with individual conformations in this ensemble. As shown in Fig. 1, our pipeline for constructing an IDP structure ensemble consists of four steps. In Step 1, we use MD simulation to generate the lifetime trajectories of the target IDP, which yields a large number of different conformations. We call this initial set of different structures the *initial structure pool*. In Step 2, we select a subset of conformations from the initial structure pool by applying a clustering algorithm [19] to group structurally-similar conformations, and picking the conformation with the lowest average root-mean-square deviation (RMSD) to all other conformations within the same cluster to represent each cluster. In Step 3, we predict the backbone chemical shifts of the selected structures using software SHIFTX2 [18]. In Step 4, we compute the weights of the selected structures such that the ensemble average best fits the chemical shift data. In particular, this step is divided into two scenarios. When the number of constraints derived from chemical shift data is equal to or larger than the number of unknown variables that need to be estimated (which

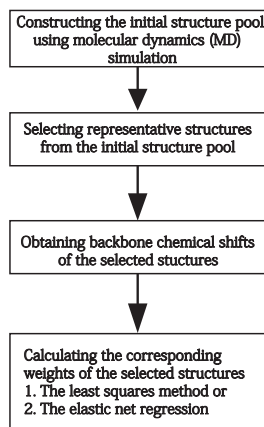


Fig. 1. Our pipeline of IDP ensemble construction.

is equal to or larger than the number of unknown variables that need to be estimated (which

is equivalent to the number of the selected structures from the initial pool), we compute the weights of all representative structures using a least squares method to minimize the discrepancy between experimentally-measured and computationally-predicted chemical shifts. When the unknown variables outnumber the constraints derived from experimental observations, the problem becomes ill-posed and underdetermined such that there are an infinite number of solutions. In this case, we add a combination of L_1 - and L_2 -norm penalties to the objective function, and apply an elastic net technique to solve the regression problem.

2.2 Constructing the Initial Structure Pool

We construct the initial structure pool of an IDP using the molecular dynamics (MD) simulation package NAMD2.8 [27]. In particular, we perform a 10ns MD simulation at the temperature 300K using the CHARMM27 force field [3] and the Generalized Born Implicit Solvent model [34]. We use either available structures from the Protein Data Bank (PDB) [2] or random structures generated by Xplor-NIH [30, 29] as the starting templates. The trajectory coordinates are saved per picosecond during the simulation, and in total 10,000 final structures are output as the initial structure pool, which is similar to procedure used in [12]. Here, we assume that the structures generated by the MD simulation can represent the diverse conformational space of the IDP. In principle, other softwares, such as Flexible-Meccano [26] and TraDES [10, 11], can also be used to increase the efficiency of conformational sampling.

2.3 Selecting Representative Structures from the Initial Pool

Considering the difficulty to directly process a large number of conformations in the initial pool, we select a subset of representative structures using the following clustering scheme: (1) Pick a structure S_i from the initial pool, and create a cluster G_i which only contains S_i at the beginning. (2) Calculate the RMSD of backbone atoms between S_i and every other structure S_j in the pool. If their RMSD is less than a cutoff (e.g., 2.1 Å, which was also used in [12]), move S_j from the structure pool to cluster G_i . (3) When all structures in the initial pool have been compared with S_i , remove S_i from the initial pool. The above process is repeated until all structures in the initial pool have been clustered. After that, for each cluster, the structure with the lowest average RMSD to all other structures in the same cluster is chosen as a representative conformation.

2.4 Obtaining the Chemical Shifts of Backbone Atoms

In our approach, we use the chemical shifts of backbone atoms, such as HN, N, CA and HA, to drive the construction of an IDP ensemble. The experimental values of chemical shifts are obtained from the Biological Magnetic Resonance Bank (BMRB) [7], while the predicted chemical shifts of the selected conformations are produced by SHIFTX2 [18] based on available structural information.

2.5 Calculating the Corresponding Weights of the Selected Structures

After selecting the representative conformations from the initial structure pool, we need to determine their corresponding weights. Depending on whether the number of constraints is less than the number of the representative structures in the ensemble, we use two different strategies to compute the optimal weights for the experimental data. Below we will describe the details of these two strategies.

The Least Squares Method When the number of unknown weights that need to be estimated is equal to or less than the number of constraints derived from chemical shifts, we apply a least squares approach to solve the regression problem. We first introduce notation before describing the details of this algorithm. Let S_j represent the j^{th} conformation in the set of the selected structures. Let a_{ij} denote the predicted backbone chemical shift of the i^{th} residue in structure S_j , and let b_i denote the experimental chemical shift of the corresponding backbone atom in the i^{th} residue in the protein, which is derived from the BMRB [7] (here, for simplicity, we assume each residue has only one backbone atom with available chemical shift). The difference between predicted and experimental backbone chemical shifts of the i^{th} residue, denoted by ε_i , can be defined as $\varepsilon_i = \left| \sum_{j=1}^n a_{ij} w_j - b_i \right|$, where w_j represents the weight of structure S_j and n stands for the total number of representative structures. Then the overall difference between predicted chemical shifts and experimental observations over all residues is defined as $\varepsilon = \sum_{i=1}^m \varepsilon_i^2$, where m stands for the number of residues whose backbone chemical shifts are available from experimental data. Then, our goal is to find the optimal weights for all representative structures that minimize ε .

We use \mathbf{A} to denote a matrix that contains all backbone chemical shifts a_{ij} predicted from SHIFTX2, \mathbf{b} to denote a vector that includes all experimental chemical shifts b_i , and \mathbf{w} to denote a vector that includes the weights of all representative structures, i.e., $\mathbf{A} = (a_{ij})_{m \times n}$, $\mathbf{b} = (b_1, b_2, \dots, b_m)^{\text{T}}$, and $\mathbf{w} = (w_1, w_2, \dots, w_n)^{\text{T}}$. Then the difference between predicted and experimental chemical shifts can be expressed as:

$$\varepsilon = \|\mathbf{A}\mathbf{w} - \mathbf{b}\|^2.$$

In addition, we need to consider the following two constraints:

$$0 \leq w_j \leq 1, \forall 1 \leq j \leq n; \text{ and } \sum_{j=1}^n w_j = 1.$$

Overall, the regression problem can be described as:

$$\min_{\mathbf{w}} \varepsilon = \|\mathbf{A}\mathbf{w} - \mathbf{b}\|^2 \quad (1)$$

$$\text{s.t. } 0 \leq w_j \leq 1, \forall 1 \leq j \leq n; \quad (2)$$

$$\sum_{j=1}^n w_j = 1. \quad (3)$$

When $m \geq n$, the above form is a typical data fitting problem with an overdetermined linear system. Thus we can solve this problem using a least squares approach that guarantees to find the optimal weights that minimize ε . In principle, the restricted quadratic programming [15] can be used to solve the above least squares problem.

The Elastic Net Method The least squares method described previously can be used to efficiently solve the regression problem in overdetermined form. However, when the unknown weights associated with representative conformations outnumber the constraints derived from experimental data, the regression model described in Section ‘‘The Least

Squares Method” becomes underdetermined and cannot be solved by the least squares method. In this case, we need to resort to other techniques to effectively solve the problem. Here, we assume that only a small number of structures are truly present in the ensemble, i.e., the weights are sparse. According to this sparsity assumption, we introduce an elastic net method (i.e., adding a combination of L_1 - and L_2 -norm penalty terms) to solve the regression problem which can address the deficiency caused by the insufficient number of experimental constraints, and thus handle the degeneracy problem well. In the new formalization of our regression problem, we relax the constraint in Equation (3), as we mainly focus on the relative weighting factors among structures in the ensemble. We can renormalize the calculated weights and obtain the updated weights of individual structures in the ensemble later. Overall, we aim to solve the following optimization problem:

$$\min_{\mathbf{w}} \varepsilon = \|\mathbf{A}\mathbf{w} - \mathbf{b}\|^2 + \lambda P_{\alpha}(\mathbf{w}) \quad (4)$$

$$s.t. 0 \leq w_j \leq 1, \forall 1 \leq j \leq n, \quad (5)$$

where $P_{\alpha}(\mathbf{w}) = \frac{1-\alpha}{2} \|\mathbf{w}\|_2^2 + \alpha \|\mathbf{w}\|_1 = \sum_{j=1}^n (\frac{1-\alpha}{2} w_j^2 + \alpha |w_j|)$, while λ and α are regularization parameters. To enforce the sparsity on the solution, we set specific values of α and λ . Alternatively, we can use a cross-validation procedure [16, 32] to determine their values.

2.6 Implementation

We implement the above two methods in MATLAB. For the least square algorithm, we call the function *lsqin* in MATLAB to calculate the optimal weights. For the elastic net based method, we first call the corresponding elastic net function *cvglmnet* [16, 32] in MATLAB to solve the regression problem through a cross-validation procedure. After that, we normalize all calculated weights and use these updated weights as our final solution.

3 Results

3.1 Validation through A Reference Ensemble Approach

We first evaluated the performance of our algorithms using the *reference ensemble method* [23]. More specifically, this method first constructs a set of “true” conformations and their corresponding weights, which are called the *reference ensemble*. Then “experimental” data are synthesized based on “true” information of this reference ensemble. The algorithm being evaluated takes these synthetic “experimental” data as input and computes a structure ensemble. By comparing the computed ensemble with the reference one, we can properly assess performance of the proposed algorithm.

Here, we validated our least squares method using alpha-synuclein protein, a 140-residue IDP that had been previously studied in [33, 28, 36], and the elastic net method using the K18 domain of Tau protein. We illustrated the validation of these two approaches individually.

For alpha-synuclein, we first constructed a reference ensemble and synthesized their “experimental” data. Then, we used these data to back-compute the optimal weights with the least squares method. Finally, we compared the performance of our method




Fig. 2. The validation results of the least squares method through the reference ensemble method. In each scenario, we compared the performance of our method with that of the Monte Carlo approach, MC stands for the Monte Carlo sampling method and LS stands for our least squares approach. In Panels (a) and (c), the weights of individual conformations in the reference ensemble followed the Gaussian distribution, while in Panels (b) and (d), the weights followed the uniform distribution. Both $\Omega(w^C, w^T)$ and the RMSD of HN chemical shift were calculated for each case. Each test was conducted 50 times with different random seeds.




Fig. 3. Results on using different chemical shift data in our validation test. Here, Gaussian noise with standard deviation 0.02ppm was used to synthesize the chemical shift data. Each test was repeated 50 times with different random seeds. A factor 0.1 was multiplied to the chemical shifts of heavy atoms (i.e., CA and N) to combine with those of hydrogen atoms (i.e., HN and HA).

with the Monte Carlo approach. As shown in Fig. 2, our least squares method significantly outperformed the Monte Carlo approach. More details of the validation of the least squares method can be found in Supplementary Material [40] Section S1.

In addition, to investigate the robustness of our method upon the types of experimental data, we tested different combinations of input chemical shift data. As shown in Fig. 3, we tested the following four combinations of chemical shift data: (1) HN chemical shifts only; (2) HN and HA chemical shifts; (3) HN, HA and CA chemical shifts; and (4) HN, HA, CA and N chemical shifts. As Fig. 2 had shown that the distribution of the synthesized weights did not affect the results, here we set w^T to follow the Gaussian distribution and set standard deviation as 0.02 for noise in synthesized data. As shown in Fig. 3, incorporating more chemical shift data yielded better results, i.e., smaller $\Omega(w^C, w^T)$ values between computed and “true” weights as well as RMSD values between back-computed and experimental data. This trend was expected, as considering more data restraints usually alleviated the difficulty of constructing an IDP en-




Fig. 4. The validation results of the elastic net method through the reference ensemble method. We compared the performance of our method with that of the Monte Carlo approach, MC stands for the Monte Carlo sampling method and EN stands for our elastic net approach. Both $O(w^C, w^T)$ and the RMSD of HN chemical shift were calculated. The test was conducted 50 times with different random seeds.

semble with the larger degrees of freedom and led to better modeling results. Though the accuracy seemed to only increase slightly when we added CA chemical shifts, we found that the interquartile range in the box plots became less, which implies the results became more stable.

Next, we validated the elastic net method with the similar strategy. As shown in Fig. 4, we found that the elastic net method significantly outperformed the Monte Carlo approach in this sparsity condition. In addition, compared to the previous non-sparsity situation (Fig.2), a larger improvement over the Monte Carlo approach was observed for the elastic net method. More details of our validation of the elastic net method can be found in Supplementary Material [40] Section S2.

3.2 Tests on Real Data

In this section, we tested our ensemble construction methods on real data of three IDPs, including alpha-synuclein, the translocation domain of Colicin N and the K18 domain of Tau protein, which contain 140, 90 and 130 residues, respectively. For the first two proteins (i.e., alpha-synuclein and Colicin N), in which the number of constraints derived from chemical shift data is larger than the number of representative structures, we applied the least squares algorithm to compute the optimal weights. While for the K18 domain of Tau protein, we used the elastic net method. For alpha-synuclein and the translocation domain of Colicin N, we ran a 10ns MD simulation and obtained a structure pool of 10,000 structures for each IDP. For the K18 domain of Tau protein, we ran a 30ns MD simulation and had a structure pool of 10,000 structures. After that, we clustered each structural pool using the procedure described in Section 2.3 with the clustering cutoff 2.1 Å. In total, we obtained 133, 233 and 882 conformations in the initial pool for alpha-synuclein, the translocation domain of Colicin N and the K18 domain of Tau protein, respectively.

The experimental chemical shift data were obtained from the Biological Magnetic Resonance Bank (BMRB) [7] and used as input data to our ensemble construction. In particular, we used HN, HA, CA and N chemical shifts for alpha-synuclein, and HN, CA and N chemical shifts for the translocation domain of Colicin N and the K18 domain of Tau protein (chemical shifts of HA for these two proteins are not available from the

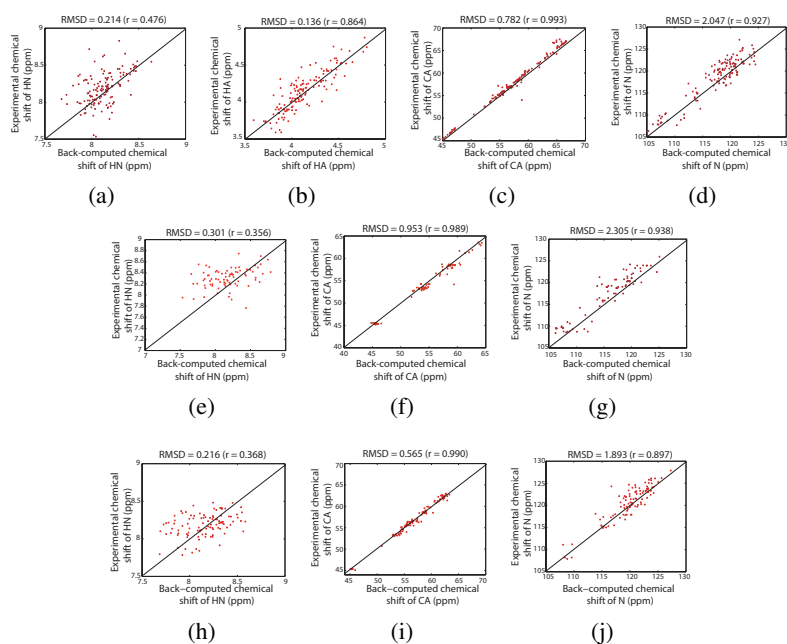


Fig. 5. Correlation between back-computed and experimental chemical shifts for alpha-synuclein, the translocation domain of Colicin N and K18 domain of Tau protein. Panels (a) to (d) are for alpha-synuclein, Panels (e) to (g) are for the translocation domain of Colicin N, and Panels (h) to (j) are for K18 domain of Tau protein. The x label stands for (ensemble-averaged) back-computed chemical shifts, while y label represents experimental chemical shifts. The symbol r represents the Pearson correlation.

BMRB database). When chemical shifts of both hydrogen and heavy atoms (e.g., CA and N) were combined, we set 0.1 as the weighting factor for the chemical shifts of heavy atoms.

To evaluate the performance of our methods, we mainly focused on the RMSD and Pearson correlation between back-computed and experimental chemical shifts of the tested proteins. The back-computed chemical shifts were calculated based on the ensemble-averaged values, which were shown in Fig. 5. We found that the back-computed chemical shifts of heavy atoms based on the constructed IDP ensembles agreed well with experimental values, with Pearson correlation above 0.89. On the other hand, the results on the back-computed chemical shifts of HN and HA (especially HN atoms) became worse. As stated in [38], this phenomenon is probably because the chemical shifts of hydrogen can be predicted much less reliably than the heavy atoms using SHIFTX2. Theoretically, the hydrogen chemical shifts are often affected more significantly by the electric field effect, ring currents, and other local shielding phenomena, which raises the difficulty for their accurate prediction using the current chemical shift prediction programs (e.g., SHIFTX2) [38, 18].




Fig. 6. The structure overlay of the 10 most probable structures (i.e., with the largest weights) in the ensemble of Tau protein computed by our elastic net based method. Panel (a) shows the view of these structures aligned using residues 20-44, and Panel (b) shows the view of zooming into the fragment of residues 20-44.

3.3 Case Study




Fig. 7. The contact maps for the top 10 structures in the ensemble of K18 domain of Tau protein calculated by our elastic net method. From (a) to (j), the weights of the corresponding structures are in decreasing order.

In this section, we focused on the IDP structure ensemble constructed by our elastic net method for the K18 domain of Tau Protein, and analyzed the structural details of long-range interactions between different residues. Overall, the structure ensemble of this protein constructed by our method contains only 58 structures with non-zero weights (i.e., whose weights is larger than the threshold 10^{-7}), as compared to the 882 structures selected from the initial pool. For the constructed ensemble, the RMSD and correlation between back-computed and experimental chemical shifts of CA were 0.953 \AA and 0.989, respectively indicating that our elastic net algorithm computed a reasonably good structure ensemble for Tau protein that agreed with experimental data. In particular, we picked the 10 most probable structures (i.e., with the largest weights) from the ensemble and looked into the details of these conformations. Fig. 6 shows the overlay of these 10 structures, where the fragments between residues 20-44 were used

in structure alignment. The region in residues 20-44 has been believed to participate interactions between two different regions divided by a turn motif “PGGG” [12]. The visualization confirmed this finding, though our results are less obvious compared to those in [12].

Next, we examined the contact maps of the top 10 structures with the largest weights which were produced based on the distance cutoff 25 Å between CA atoms. Though some contact maps may look similar, these structures are actually quite different. As shown in Fig. 7, for most structures, the regions near residues 33-38 and 64-69 exhibited long-range interactions with the N-terminal residues that are at least 5 residues away along the sequence.

These observations were consistent with the previous studies in [24]. These two regions (i.e., residues 33-38 and 64-69) form the *paired helical filament* (PHF), and are supposed to play core functions in the transition state of the Tau aggregation process from normal/unfolded form to pathological states [37]. Thus, the structures modeled by our computational method can provide useful molecular basis for further investigating the functional roles of Tau protein in the related diseases.

4 Conclusions

Constructing structure ensembles for IDPs is a challenging but important task for understanding their cellular functions. In this paper, we proposed two novel approaches based on least squares and elastic net techniques to construct the structure ensemble of an IDP from chemical shift data. Validation via the reference ensemble approach and tests on real data has demonstrated the superiority of our methods over traditional stochastic sampling based approaches. Our least squares method can find the global optimal solution, and the constructed ensemble can depict the structures of IDPs more accurately. Furthermore, our elastic net based method can successfully address the degeneracy problem, which is the current main challenge of the IDP ensemble construction task.

There are several possible extensions to our IDP structure ensemble. In the current version of our ensemble construction pipeline, we assume that the life trajectories generated by the MD simulation can represent the whole conformational space. In the future, we will improve the current procedure of constructing the initial structure pool, such that it can be better integrated with the computation of optimal weights. In addition, for future work, we will incorporate more information to guide our ensemble construction, such as additional experimental data and available prior knowledge about IDP structures.

References

1. Jane R. Allison *et al.* *Journal of the American Chemical Society*, 131(51):18314–18326, 2009.
2. Helen M. Berman *et al.* *Nucleic Acids Research*, 28(1):235–242, 2000.
3. Bernard R. Brooks *et al.* *Journal of Computational Chemistry*, 4(2):187–217, 1983.
4. Andrea Cavalli *et al.* *Proceedings of the National Academy of Sciences*, 104(23):9615–9620, 2007.
5. Yiwen Chen *et al.* *Biophysical Journal*, 93(7):2300 – 2306, 2007.
6. Wing-Yiu Choy and Julie D. Forman-Kay. *Journal of Molecular Biology*, 308(5):1011 – 1032, 2001.
7. Jurgen F. Doreleijers *et al.* *Journal of Biomolecular NMR*, 26(2):139–146, 2003.

8. A.Keith Dunker *et al.* *Journal of Molecular Graphics and Modelling*, 19(1):26 – 59, 2001.
9. H. Jane Dyson and Peter E. Wright. *Nat Rev Mol Cell Biol*, 6(3):197–208, March 2005.
10. Howard J. Feldman and Christopher W.V. Hogue. *Proteins: Structure, Function, and Bioinformatics*, 39(2):112–131, 2000.
11. Howard J. Feldman and Christopher W.V. Hogue. *Proteins: Structure, Function, and Bioinformatics*, 46(1):8–23, 2002.
12. Charles K. Fisher *et al.* *Journal of the American Chemical Society*, 132(42):14919–14927, 2010.
13. Charles K Fisher and Collin M Stultz. *Current Opinion in Structural Biology*, 21(3):426 – 431, 2011.
14. Rasmus Fonseca *et al.* *Nucleic Acids Research*, 42(15):9562–72, August 2014.
15. Marguerite Frank and Philip Wolfe. *Naval Research Logistics Quarterly*, 3(1-2):95–110, 1956.
16. Jerome H. Friedman *et al.* *Journal of Statistical Software*, 33(1):1–22, 2 2010.
17. Debabani Ganguly and Jianhan Chen. *Journal of Molecular Biology*, 390(3):467 – 477, 2009.
18. Beomsoo Han *et al.* *Journal of Biomolecular NMR*, 50(1):43–57, 2011.
19. Trevor Hastie *et al.* Springer Series in Statistics. Springer New York Inc., New York, NY, USA, 2001.
20. Lilia M. Iakoucheva *et al.* *Journal of Molecular Biology*, 323(3):573 – 584, 2002.
21. Malene Ringkjøbing Jensen *et al.* *Journal of the American Chemical Society*, 132(4):1270–1272, 2010.
22. Janina S. Karres *et al.* *Cell*, 131(1):136 – 145, 2007.
23. John Kuriyan *et al.* *Journal of Molecular Biology*, 190(2):227 – 254, 1986.
24. Joseph A. Marsh and Julie D. Forman-Kay. *Journal of Molecular Biology*, 391(2):359 – 374, 2009.
25. Gabrielle Nodet *et al.* *Journal of the American Chemical Society*, 131(49):17908–17918, 2009.
26. Valéry Ozenne *et al.* *Bioinformatics*, 28(11):1463–1470, 2012.
27. James C. Phillips *et al.* *Journal of Computational Chemistry*, 26(16):1781–1802, 2005.
28. Alessandra Recchia *et al.* *The FASEB Journal*, 18(6):617–626, 2004.
29. Charles D. Schwieters *et al.* *Progress in Nuclear Magnetic Resonance Spectroscopy*, 48(1):47 – 62, 2006.
30. Charles D Schwieters *et al.* *Journal of Magnetic Resonance*, 160(1):65 – 73, 2003.
31. Yang Shen *et al.* *Proceedings of the National Academy of Sciences*, 105(12):4685–4690, 2008.
32. Noah Simon *et al.* *Journal of Statistical Software*, 39(5):1–13, 3 2011.
33. Kamil Tamiola *et al.* *Journal of the American Chemical Society*, 132(51):18000–18003, 2010.
34. David E. Tanner *et al.* *Journal of Chemical Theory and Computation*, 7(11):3635–3642, 2011.
35. Tsuyoshi Terakawa and Shoji Takada. *Biophysical Journal*, 101(6):1450 – 1458, 2011.
36. Kostas Vekrellis and Leonidas Stefanis. *Expert Opinion on Therapeutic Targets*, 16(4):421–432, 2012.
37. Martin von Bergen *et al.* *Biochimica et Biophysica Acta (BBA) - Molecular Basis of Disease*, 1739(2C3):158 – 166, 2005.
38. David S. Wishart and Alex M. Nip. *Biochemistry and Cell Biology*, 76(2):153–163, 1998.
39. Hui Zou and Trevor Hastie. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.
40. H. Gong *et al.* Constructing Structure Ensembles of Intrinsically Disordered Proteins from Chemical Shift Data–Supplementary Material. Available at <http://iiis.tsinghua.edu.cn/%7ecompbio/papers/recomb15IDPSM.pdf>, 2015