

Protein model refinement using an optimized physics-based all-atom force field

Anna Jagielska, Liliana Wroblewska, and Jeffrey Skolnick*

Center for the Study of Systems Biology, School of Biology, Georgia Institute of Technology, 250 14th Street NW, Atlanta, GA 30318

Edited by Harold A. Scheraga, Cornell University, Ithaca, NY, and approved March 14, 2008 (received for review January 3, 2008)

One of the greatest challenges in protein structure prediction is the refinement of low-resolution predicted models to high-resolution structures that are close to the native state. Although contemporary structure prediction methods can assemble the correct topology for a large fraction of protein domains, such approximate models are often not of the resolution required for many important applications, including studies of reaction mechanisms and virtual ligand screening. Thus, the development of a method that could bring those structures closer to the native state is of great importance. We recently optimized the relative weights of the components of the Amber ff03 potential on a large set of decoy structures to create a funnel-shaped energy landscape with the native structure at the global minimum. Such an energy function might be able to drive proteins toward their native structure. In this work, for a test set of 47 proteins, with 100 decoy structures per protein that have a range of structural similarities to the native state, we demonstrate that our optimized potential can drive protein models closer to their native structure. Comparing the lowest-energy structure from each trajectory with the starting decoy, structural improvement is seen for 70% of the models on average. The ability to do such systematic structural refinements by using a physics-based all-atom potential represents a promising approach to high-resolution structure prediction.

Amber force field | force field optimization | protein structure prediction | all-atom potential

The past several years have witnessed significant progress in the field of protein structure prediction (1–10), with contemporary methods being able to assemble the correct topology for a large fraction of protein domains. Such approximately correct models typically vary in their structural similarity to the native state, with a rmsd (root mean square deviation) from native that ranges from 1 Å to ≈ 6 Å. Models with a rmsd to native of 1–2 Å are comparable to experimentally obtained structures and can be used in a broad range of applications, including studies of reaction mechanisms and virtual ligand screening (2). In contrast, the range of applicability of lower-resolution models (with a 3- to 6-Å rmsd from native) is smaller (2). Structure prediction methods often use a coarse-grained protein representation to enhance the conformational search efficiency. To further improve model quality, it is possible that additional structural details need to be included. A tempting approach is to use an all-atom detailed protein representation for the final stages of structure prediction, but despite considerable effort, all-atom refinement has seen little success (11–13). Over the years, there have been individual examples of successful refinement (11–16), with the best improvement of ≈ 2 Å (12, 16) and the largest refinement benchmarks consisting of a small set of proteins (12, 13, 17–19). Although isolated examples of refinement have been reported, the methods are far from routine; in reality, most models deteriorate instead of improve.

In protein structure prediction and refinement, the challenge is twofold: one needs an effective conformational search scheme and an energy function whose global minimum is in the protein's native state. Moreover, the energy surface should be funnel-like so that the potential can drive the structure toward lower energy, more native-like conformations. For this to occur, the energy

function must have a correlation with native structure similarity (20). Previously (21), we explored the possibility of creating a funnel-like shape for the Amber ff03-based potential (22) by global optimization of the weights of the individual energy components. The optimized force field had a significant correlation with native similarity and, for decoy evaluation, could recognize the native conformation among decoy structures for a large fraction of proteins examined. Here, we test the refinement ability of the newly derived potential on a representative benchmark set of 47 proteins (among them, eight proteins were a part of the training set used in the optimization of the force field, and the remaining 39 proteins composed the testing set), each having a diverse set of compact all-atom decoy structures. In the following, we first describe the correlation of the energy with native-likeness of the decoys obtained in the refinement conformational search. Next, we discuss the improvement of decoy structures within each individual refinement trajectory. Finally, we present the results of structural refinement over the best initial models in the entire decoy ensemble and discuss the ability of the force field to select structures close to the native state from the entire decoy set.

Results and Discussion

Correlation of the Energy with Native Similarity as Measured by the Template Modeling Score (TM-Score). The ability of a force field to refine a model and select the native or close to native structure by using energy as the selection criterion is related to the correlation of its energy with native structure similarity. The force field used in this study had an average correlation coefficient of energy with TM-score [a measure of structural similarity (23) whose range is (0, 1], with a TM-score of 0.3 for the best structural alignment of a pair of randomly related structures and 1.0 for identical structures] of 0.59 after global optimization as calculated on a large decoy set of 58 proteins (21). Moreover, for 47% of the tested proteins, the correlation coefficient was significant, above 0.60. However, the decoys were generated by using a different force field, the ff03 Amber potential (22). This high correlation coefficient of energy with native-likeness might be an artifact of optimization and could be lost during a thorough conformational search that is driven by the new potential. To explore this issue, we calculated the correlation coefficient (CC) between the energy and TM-score for the structures generated during conformational refinement by using the optimized force field for the 39 protein test set [see supporting information (SI) Table S1]. The resulting average value of CC was 0.59, with 46% of the proteins having the CC above 0.60; this corresponds well to previously obtained values after force field optimization and

Author contributions: A.J., L.W., and J.S. designed research; A.J. and L.W. performed research; A.J. and L.W. analyzed data; and A.J., L.W., and J.S. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

*To whom correspondence should be addressed. E-mail: skolnick@gatech.edu.

This article contains supporting information online at www.pnas.org/cgi/content/full/0800054105/DCSupplemental.

© 2008 by The National Academy of Sciences of the USA

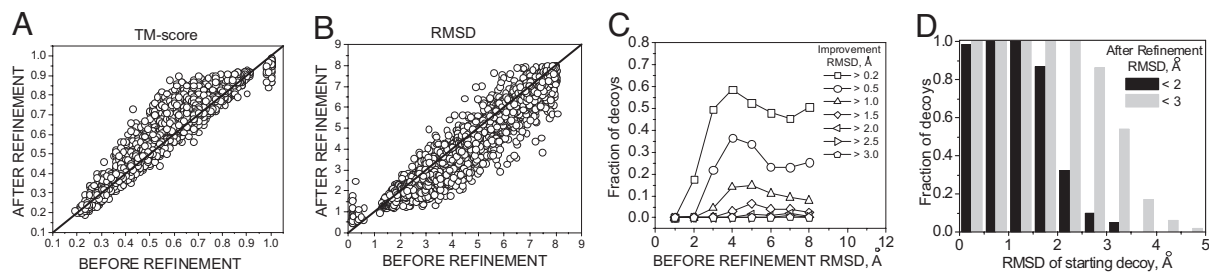


Fig. 1. For all refinement trajectories, results of decoy structure refinement within each trajectory (100 decoys per protein, 39 proteins). (A and B) TM-score (A) and C^{α} rmsd (B) to the native structure of the lowest-energy decoy after refinement versus decoy's initial TM-score (A) or C^{α} rmsd (B). (C) Fraction of decoys that refined by more than a given C^{α} rmsd threshold (for 0.2-, 0.5-, 1.0-, 1.5-, and 2.0-Å rmsd thresholds) with respect to their initial native similarity. (D) Fraction of decoys that refined to (or remained within) the accuracy of 2-Å (black) and 3-Å (gray) C^{α} rmsd to the native state with respect to their initial native similarity. Fractions of decoys were calculated in 1-Å bins.

decoy ranking (21). Therefore, during the conformational search, the characteristics of our optimized energy landscape are preserved. Fig. S1 shows examples of the energy-TM-score clouds for the unoptimized original ff03 force field and decoys generated with the original ff03 potential (Fig. S1A), the optimized force field and decoys generated with the original ff03 potential (Fig. S1B), and the optimized force field and decoys generated with the optimized force field during the refinement runs (Fig. S1C).

For selection of the correct structure on the basis of energy, the most important quantity is the correlation of the minimum energy structure at a given TM-score with native structural similarity. This correlation of energy vs. TM-score for decoys obtained during the refinement search was 0.44, with 46% of proteins having a CC >0.6 and 51% having a CC >0.5. The corresponding correlation for the decoy set used in force field optimization and decoy ranking was 0.49, with 33% of proteins having a CC >0.6 and 53% having a CC >0.5.

Refinement of Protein Decoys. Improvement within a trajectory. For each member of the 39-protein test set, the ability to refine 100 starting structures spanning a range of TM-scores was examined. During refinement, the TM-score and C^{α} rmsd improve for the majority of the decoys. In Fig. 1, the TM-score (A) and C^{α} rmsd (B) to the native structure of the lowest-energy decoy from each refinement trajectory are compared with the initial decoy TM-score and rmsd. For TM-score compared to rmsd, we observe a more pronounced improvement than decoy deterioration. This reflects the fact that the force field was optimized as a function of the TM-score, not the rmsd. Furthermore, the TM-score is generally more sensitive to native-like features. Sometimes improvements in TM-score may cause an increase of rmsd from the native structure, e.g., when the core of a protein is improved at the cost of moving a protein's termini farther from the native state. Below rmsd values of 1 Å (or TM-scores above 0.9), the force field cannot differentiate among structures. This insensitivity is evident in Fig. 1, where the native structures (TM-score close to 1) drift away from the crystal structures on average by ≈ 0.1 TM-score unit or by a rmsd of 1 Å. This drift defines the resolution of the force field, which corresponds to a C^{α} rmsd of ≈ 1 Å from the native state.

On average, with respect to the initial structure and over the whole range of native similarity, 70% of decoys improve their TM-score, 18% get worse, and 12% are unchanged (Fig. 1A). When rmsd (Fig. 1B) is used, 70% of decoys improve, 28% deteriorate, and 2% are unchanged. In Fig. 1C, we show the fraction of decoys that improved by more than a given rmsd value with respect to the decoy's initial rmsd. For 37% of decoys with initial rmsd of 3–4 Å, we observe improvement of >0.5 Å, and for 14%, the improvement was >1.0 Å. There are 28 cases of

improvement compared to only 3 cases of deterioration larger than 2 Å (with the largest improvement being 3.63 Å and largest deterioration being 2.21 Å).

The ultimate goal, however, is to refine structures to near experimental accuracy, i.e., below 2- to 3-Å C^{α} rmsd to the native state. In our benchmark, as shown in Fig. 1B, a significant fraction of decoys improved to such accuracy; we refined 12%, 10%, and 4% of all of the decoys to an rmsd below 3, 2.5, and 2 Å, respectively, starting from structures with initial rmsd greater than the given threshold. Especially encouraging is that some of the refinements to near experimental accuracy occurred for the decoys that were 4–6 Å away from the native state; there are also improvements from above 3 Å to below 1.5 Å. The fractions of decoys that improved to an rmsd of 2 and 3 Å with respect to their initial rmsd to the native state are presented in Fig. 1D. We observed refinement below 2 Å rmsd for 34% of decoys with initial rmsd between 2 and 2.5 Å; an rmsd below 3 Å was obtained for 54% of decoys with an initial rmsd between 3 and 3.5 Å. Refinement below an rmsd of 2 Å is seen for decoys that were as far from the native state as 4 Å; improvement below 3 Å rmsd was seen even for decoys with initial rmsd 6 Å to the native structure. An improved sampling scheme may allow us to achieve larger improvement in structure and enable the refinement to experimental accuracy for decoys more structurally distant from the native state (above 4 Å).

More detailed analysis of the refinement results, demonstrating a domination of structural improvement over deterioration can be found in Fig. S2. Note that the native state is reasonably stable and deteriorates on average by only ≈ 0.05 TM-score unit, or 0.65 Å in rmsd.

We also analyzed the improvement of different structural protein elements, i.e., helices, β -sheets, and loops. In Fig. 2A–C and A'–C', the C^{α} rmsd to the native structure for helices (Fig. 2A and A'), β -sheets (Fig. 2B and B'), and loops (Fig. 2C and C') for the lowest-energy decoy from each trajectory is shown with respect to its initial rmsd (see *Materials and Methods* for secondary structure definitions). In Fig. 2A–C, we consider all of the secondary structure elements of a given type from the native structure superimposed together onto related elements in a decoy (e.g., all of the regions that are helical in the native structure were superimposed with the same regions in the decoy). This comparison, therefore, includes both the relative orientation and geometry of all secondary structure elements. In this analysis, we obtained improvement for 67%, 81%, and 66% of decoys over helical, β -sheet, and loop regions, respectively. The average improvement/deterioration was 0.41/0.27 Å for helices, 0.48/0.24 Å for β -sheets, and 0.58/0.39 Å for loop regions. The distribution of structural improvements with respect to the initial rmsd over secondary structure is shown in Fig. S3A–C. For example, improvements larger than 1.0 Å are

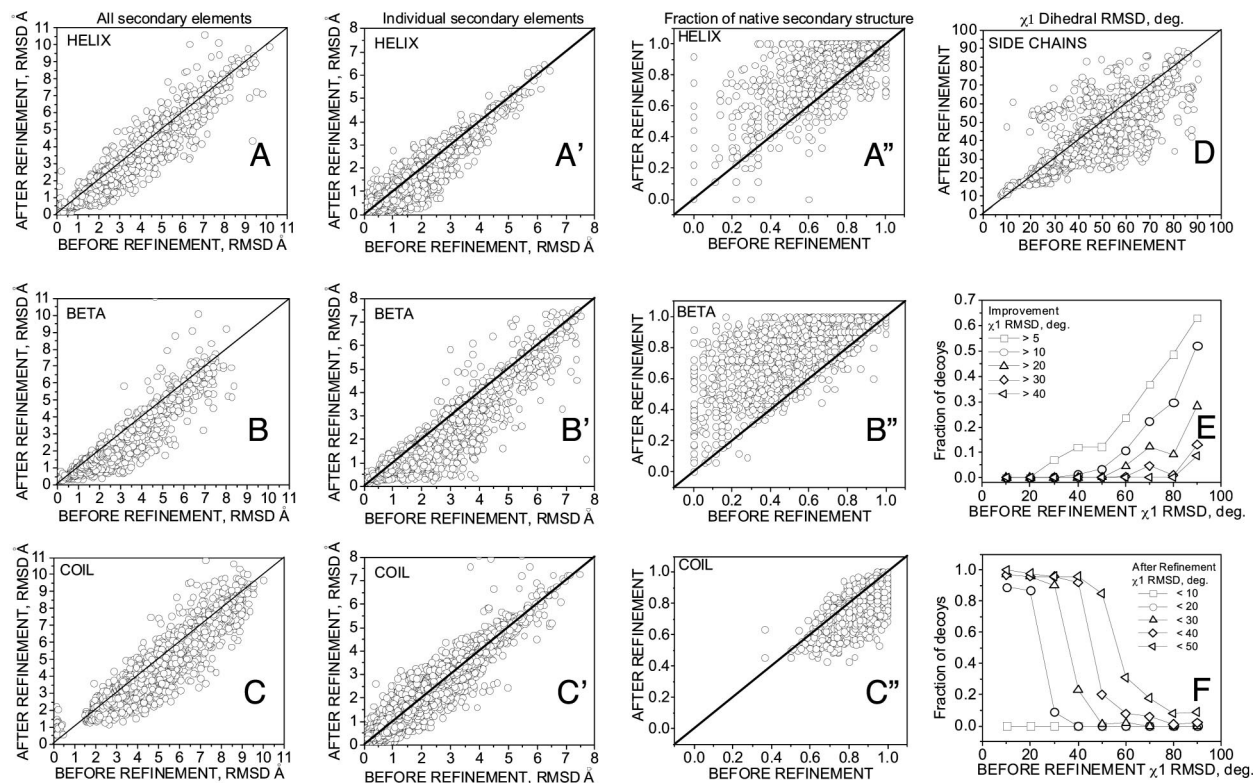


Fig. 2. Results of decoy structure refinement within each trajectory (100 decoys per protein, 39 proteins) for different structural elements. (A–C and A'–C') C^{α} rmsd to the native structure of the lowest-energy decoy after refinement versus the decoy's initial C^{α} rmsd, for helical (A), β -sheet (B), and loop (C) regions when all the native secondary structure elements of a given type are superimposed together onto the corresponding decoy regions, and for helices (A'), β -sheets (B'), and loops (C') when individual secondary structure elements from the native are superimposed onto related decoy regions. (A''–C'') Comparison of fractions of the native secondary structure in the refined (lowest-energy decoy from each trajectory) and the initial decoy, for helical (A''), β -sheet (B''), and loop (C'') regions. (D–F) Refinement of packing of buried side chains. (D) χ_1 rmsd (degrees) to the native structure of the lowest-energy decoy after refinement with respect to χ_1 rmsd of the initial structure. (E and F) Fraction of the lowest-energy decoys that improved their χ_1 rmsd by more than (E) or below (F) a given threshold with respect to χ_1 rmsd of the initial structure. Fractions of decoys were calculated in 10° bins.

observed for $\approx 10\%$ and 23% of decoys with initial rmsd between 3 and 4 Å over helical (Fig. S3A) and β -sheet (Fig. S3B) structures, respectively, and for $\approx 25\%$ of decoys with initial rmsd between 4 and 5 Å over loop regions (Fig. S3C).

In Fig. 2 A'–C', we considered each secondary structure element separately (e.g., each individual helix from the native structure was superimposed onto the related region in the decoy). In this analysis, only refinement of the secondary structure is considered, without the orientation component. We see improvement for 70%, 76%, and 50% of decoys for the helical, β -sheet, and coil regions, respectively. The distribution of structural improvements with respect to the initial rmsd over individual secondary structure elements is shown in Fig. S3 A'–C'. Improvements larger than 1.0 Å are observed for $\approx 3\%$, 21%, and 2% of decoys with initial rmsd between 3 and 4 Å over helical (Fig. S3A'), β -sheet (Fig. S3B'), and loop (Fig. S3C') regions, respectively.

In Fig. 2 A''–C'', the assignment to secondary structure class between the native structure and decoys is compared. The fractions of the native helical (Fig. 2A'') and β -sheet (Fig. 2B'') content improve for the majority of the decoys. The fraction of native coil (Fig. 2C'') decreases; we observe some tendency of our force field to turn coil residues into helical conformation.

Refinement is seen on average for all secondary structure types, including loop regions, which are recognized as the most difficult to refine. Both the orientation and geometry of individual secondary structure elements improve. However, the refinement of loops is more pronounced in the correction of their

relative orientation than in the improvement of the internal structure of individual loops.

Our force field was optimized only with respect to the correlation of the energy with the main-chain native similarity; the side-chain geometry has been neglected in the optimization process. Despite this neglect, we observe an average improvement in the χ_1 dihedral angles of the side chains in the protein interior. Fig. 2D shows the χ_1 rmsd of the buried side chains to the native structure for the lowest-energy decoy from each trajectory with respect to the initial decoy's χ_1 rmsd (see *Materials and Methods* for explanation of the χ_1 rmsd calculation and the definition of buried side chains). Side-chain packing improves for 68% and declines for 32% of the decoys, as judged by χ_1 rmsd. The extent of improvement is large for larger initial distortions of χ_1 (Fig. 2E); the largest improvements are $>40^{\circ}$. In Fig. 2F, the fraction of decoys that improved their side-chain packing below a given threshold of χ_1 rmsd is shown with respect to the initial χ_1 rmsd. For 20% of the decoys with an initial χ_1 rmsd between 40° and 50° , the χ_1 rmsd decreases below 40° .

In Fig. S4, we show the fractions of decoys that improve/deteriorate for each protein in the entire 47-protein set. For most proteins, more than 50% of the decoys improve. Only 4 proteins, 1a19A, 1b9wA, 1c1yB, and 1dt4A, have less than a 50% improvement (with only 1b9wA and 1dt4A having more deteriorations than improvements). There are 9 proteins in the set that contain disulfide bonds (1a43, 1aazA, 1bunB, 1bvnT, 1cc7A, 1dtdB, 1f94A, 1b9wA, and 1cbp). Their refinement results are on average worse (61% improvement and 26% deterioration, as measured by the TM-score) than those for the remaining

(Fig. S6, black bars). Structures coming from the native trajectory are usually better packed than decoys and therefore favored by energy. The discrepancy between results when the native trajectory is included/excluded points out the need for better conformational sampling.

We additionally consider the quality of the lowest-energy conformation (results not shown in Fig. S6). When native decoys are excluded, 59% of proteins have the lowest-energy structure with a TM-score above 0.70, and when native decoys are included, 79% have the lowest-energy structure with a TM-score above 0.70; 66% (native decoys excluded) and 82% (native decoys included) of proteins have the lowest-energy structure with an rmsd below 3.5 Å. Based on the above results, we conclude that the ability of the force field to find the native-like structures among decoys is quite good.

Analysis of Possible Reasons for the Observed Refinement. We performed numerous analyses to establish the significance of our refinement results. One of the common problems in the training of potential functions is the decoy generation procedure (20). Some decoy sets contain certain characteristics that are easy to memorize during training, but are not transferable to other decoy sets. A classic example is the effect of “swollen” decoys, where the decoy set has a large correlation between native similarity and radius of gyration (CCRG) (21). With such a correlation, the recognition of the nearest-native structure is trivial: just pick the best-packed model. In contrast, the average CCRG for decoys in our set of 47 proteins is small (0.4). To ensure that the improvements are not dominated by those proteins with highest CCRG, for each protein we compared the correlation between native-similarity and the radius of gyration with the fraction of decoys that improved during the refinement. Although the fraction of improvements is on average a little higher for proteins with high CCRG, such a cross-correlation is only 0.4. There are proteins whose decoy set has a very low CCRG (e.g., 0.15) and yet >70% of their models improve during refinement. Improvement in structure quality dominates over deterioration also for decoys that are better packed than the native structure (smaller radius of gyration). Monitoring the extent of decoy improvement with respect to the contact order (24) of the decoys, instead of the radius of gyration, gives similar results. Although significant improvements are observed more often for decoys with lower contact order, the overall correlation between the rate of improvement and the initial contact order with respect to the native is low. Finally, we also found that structure improvement does not depend on the number of hydrogen bonds (normalized by the protein length). Decoys of β -type structures tend to refine slightly better than α -helix-containing protein structures, but a larger protein set is needed to confirm this tentative observation.

Conclusions. In our previous work (21), we optimized a physics-based, all-atom energy function derived from the Amber ff03 potential to improve the correlation between native similarity (represented by the TM-score) and energy. Here, we tested the ability of such a funnel-shaped potential to refine decoys of 47 single-domain, nonhomologous proteins with different folds, 39 of which were not a part of the optimization protein set. When the lowest-energy structure from the particular refinement conformational search trajectory is compared with the starting decoy, we observe structural improvements for 70% of the models on average; 10% of decoys refined to near experimental accuracy, below 2.5 Å. Such systematic refinement results suggest a promising approach to high-resolution structure prediction. In a more stringent test, when the best (of five lowest-energy) refined structures selected by their energies are compared against the best available starting decoy within a given native similarity range, we see improvement for the majority of

proteins that have a significant correlation of energy with TM-score.

As we discussed previously (21), the optimized force field used here does not include the electrostatic energy and generalized Born solvation (25) terms present in the original Amber ff03 force field (22). For the set of compact decoys used for force field optimization, these energy terms were uncorrelated with native-likeness and their relative weights were small. The lack of these terms might cause the appearance of some low-energy, unphysical structures during the conformational search, e.g., the burial of hydrophilic residues in the protein core. However, during the course of our refinement simulations, we did not observe such low-energy unphysical structures. Their absence can be attributed partially to the rather local conformational search during refinement that starts from already packed structures. For purposes of refinement, the exclusion of these energy components carries the advantage of faster energy evaluation and concomitantly more extensive conformational sampling for a given amount of simulation time.

We do not observe a significant correlation between the successful refinement and decoy radius of gyration, contact order, or number of hydrogen bonds per residue. Decoys of β -type structures tend to refine slightly better than those of α -type structures, probably because of the relatively large contribution of the hydrogen bond energy that is more sensitive to the misfolding of β -sheets, whereas the geometry of helices is rather insensitive to such effects. Finally, based on the results for small disulfide-bonded proteins, a better treatment of S-S bridges is required.

The direct comparison with the results of others is difficult because each group uses different procedures, refinement criteria, and different sets of proteins; however, none of the previously examined sets was large enough to be statistically significant. The status of previous work can be found in the [SI Text](#).

Overall, we have demonstrated successful refinement for the majority of testing proteins over a range of lengths and with different secondary structure classes. The protein decoy structures systematically improve over all ranges of native similarity and for all major structural elements, i.e., helices, β -sheets, and loops; only for structures below the resolution of the potential, i.e., with a rmsd to native below 1 Å or a TM-score above 0.9, does this conclusion not hold. We also see improvement of side-chain packing in the interior of the protein. However, it is important to recognize that these results have been demonstrated for protein structures that satisfy the following: (i) the proteins are single-domain monomers, without cofactors; (ii) the decoys are compact, spanning the range of 0–8 Å C α rmsd to the native structure; and (iii) the conformational search typically did not explore global changes in structures. It may be possible to apply our refinement protocol to low-resolution decoys directly generated by the TASSER coarse-grained force field (1) so that we can begin to address the end game of protein structure prediction: protein structure refinement that has been a long-sought goal of hierarchical approaches to protein structure prediction.

Materials and Methods

Conformational Search. To search protein conformational space, we used the newly developed A-TASSER program, described previously (21). A-TASSER (for atomic-TASSER) represents the protein at atomic detail and employs the Replica Exchange Monte Carlo (REMC) (26, 27) search method with a Parallel Hyperbolic Sampling (PHS) acceptance criterion (28) to reduce higher energy barriers. A-TASSER uses three types of moves that change the torsional angles of the molecule: local “fixed end” moves (29), end moves, and the side-chain moves. During refinement, the bond lengths and valence angles were fixed at the values taken from the relaxed starting structures, after a gradient-based minimization with Amber ff03 force field (22). A detailed description of the

A-TASSER search scheme can be found in the [SI Text](#) and in our previous work (21). We do note that enhancements in the move set are required.

Force Field. The potential energy function used in this study to refine protein models is calculated according to Eq. 1:

$$E_{\text{ff03/HB/R}} = w_{\text{DIH}}E_{\text{DIH}} + w_{\text{VDW}}E_{\text{VDW}} + w_{\text{VDW1-4}}E_{\text{VDW1-4}} + w_{\text{SA}}E_{\text{SA}} + w_{\text{HB}}E_{\text{HB}} \quad [1]$$

In Eq. 1, the following abbreviations and symbols are used: E , total energy; w , weight of a given energy component; DIH, dihedral term; VDW, van der Waals component; VDW1–4, van der Waals energy for atom pairs separated by less than four bonds; SA, surface area-dependent term (the hydrophobic component of the solvation free energy); and HB, hydrogen bond term. The E_{DIH} , E_{VDW} , $E_{\text{VDW1-4}}$, and E_{SA} energy terms are identical with those used in the ff03 Amber force field. The E_{HB} hydrogen bond energy was implemented by following the DSSP approach (30) and was described previously (21). The weights of the energy terms, w (Table S3), were adjusted by using a global optimization method (31) for a large set of decoy structures of a representative 58-protein set (21). Optimization was aimed at maximizing the correlation of the energy with TM-score (23) and the energy gap between the native state and the decoys. The force field used in this study has an average correlation coefficient of energy with TM-score of 0.59 and ranks structures with TM-score >0.9 (native-like) as the lowest in energy for 72% of proteins, as calculated for the protein decoy set used in the optimization study (21).

Protein Set and Starting Decoy Structures. We tested our method on 47 proteins, a subset of a previously prepared (23) comprehensive benchmark set, which covers the PDB library (32) with lengths from 41 to 200 residues at 35% sequence identity. The chosen proteins span lengths from 54 to 123 residues and represent different secondary structural classes. The list of proteins can be found in Table S1. Among them, eight (marked in Table S1 and in Fig. S4) were a part of the training set used in the optimization of the force field (21). These

were excluded from most analyses to avoid any possible memorization effects. Only Fig. S4 shows results for all 47 proteins; all other results include only the 39 testing proteins. For each protein, we randomly chose 100 decoys from the force field optimization decoy set such that they span the range of C α rmsd to the native structure from 0 to 8 Å. These 100 decoys per protein and the native structures in all-atom representation were starting models in our refinement benchmark.

Refinement Protocol. For each decoy, we ran an A-TASSER search consisting of 1,000 swaps between replicas and 200 steps of Parallel Hyperbolic Sampling of each replica between swaps. From each decoy trajectory, the lowest or the best of the five lowest-energy structures were selected for analysis. No clustering was used in decoy selection.

Selection of Secondary Structure Elements. The helices, β -strands, and loop regions of protein were defined by using the DSSP program (30) as applied to the native structure. Only continuous elements longer than three residues were considered.

Selection of Buried Side Chains. The side chain was considered buried in the interior of a protein if its surface accessible area in the native structure was <50% of the surface accessible area of the free amino acid flanked by single glycine residues. The surface accessible area of each side chain was calculated by using DSSP (30).

rmsd over χ_1 Dihedral Angles. The rmsd over χ_1 dihedral angles was calculated as the rmsd between the set of χ_1 dihedral angles for buried side chains in the native structure and the χ_1 dihedral angles of the same side chains in the decoy structure.

ACKNOWLEDGMENTS. This research was supported in part by National Institutes of Health Grant RR-12255. Calculations were conducted partly by using the resources of the Terascale Computing System at the Pittsburgh Supercomputer Center.

- Zhang Y, Arakaki AK, Skolnick J (2005) TASSER: An automated method for the prediction of protein tertiary structures in CASP6. *Proteins* 61(Suppl 7):91–98.
- Baker D, Sali A (2001) Protein structure prediction and structural genomics. *Science* 294:93–96.
- Ginalski K, Grishin NV, Godzik A, Rychlewski L (2005) Practical lessons from protein structure prediction. *Nucleic Acids Res* 33:1874–1891.
- Fisher D (2003) 3D-SHOTGUN: A novel, cooperative, fold-recognition meta-predictor. *Proteins* 51:434–441.
- Kolinski A, Bujnicki J (2005) Generalized protein structure prediction based on combination of fold-recognition with de novo folding and evaluation of models. *Proteins* 61:84–90.
- Lundstrom J, Rychlewski L, Bujnicki J, Elofsson A (2001) Pcons: A neural-network-based consensus predictor that improves fold recognition. *Protein Sci* 10:2354–2362.
- Simons KT, Bonneau R, Ruczinski I, Baker D (1999) Ab initio protein structure prediction of CASP III targets using ROSETTA. *Proteins* 37(Suppl 3):171–176.
- Zhang Y, Kolinski A, Skolnick J (2003) TOUCHSTONE II: A new approach to ab initio protein structure prediction. *Biophys J* 85:1145–1164.
- Zhang Y, Skolnick J (2004) Automated structure prediction of weakly homologous proteins on a genomic scale. *Proc Natl Acad Sci USA* 101:7594–7599.
- Zhang Y, Skolnick J (2004) Tertiary structure predictions on a comprehensive benchmark of medium to large size proteins. *Biophys J* 87:2647–2655.
- Lee MR, Tsai J, Baker D, Kollman PA (2001) Molecular dynamics in the endgame of protein structure prediction. *J Mol Biol* 313:417–430.
- Chen J, Brooks CL, III (2007) Can molecular dynamics simulations provide high-resolution refinement of protein structure? *Proteins* 67:922–930.
- Fan H, Mark AE (2004) Refinement of homology-based protein structures by molecular dynamics simulation techniques. *Protein Sci* 13:211–220.
- Simmerling C, Strockbine B, Roitberg AE (2002) All-atom structure prediction and folding simulations of a stable protein. *J Am Chem Soc* 124:11258–11259.
- Lee MR, Baker D, Kollman PA (2001) 2.1 and 1.8 Å average C α RMSD structure predictions on two small proteins, HP-36 and S15. *J Am Chem Soc* 123:1040–1046.
- Vieth M, Kolinski A, Brooks CL, III, Skolnick J (1994) Prediction of the folding pathways and structure of the GCN4 “leucine zipper.” *J Mol Biol* 237:361–367.
- Bradley P, Misura KM, Baker D (2005) Toward high-resolution de novo structure prediction for small proteins. *Science* 309:1868–1871.
- Misura KM, Baker D (2005) Progress and challenges in high-resolution refinement of protein structure models. *Proteins* 59:15–29.
- Misura KM, Chivian D, Rohl CA, Kim DE, Baker D (2006) Physically realistic homology models built with ROSETTA can be more accurate than their templates. *Proc Natl Acad Sci USA* 103:5361–5366.
- Wroblewska L, Skolnick J (2007) Can a physics-based, all-atom potential find a protein's native structure among misfolded structures? I. Large scale AMBER benchmarking. *J Comput Chem* 28:2059–2066.
- Wroblewska L, Jagielska A, Skolnick J (2008) Development of a physics-based force field for the scoring and refinement of protein models. *Biophys J* 94:3227–3240.
- Duan Y, et al. (2003) A point-charge force field for molecular mechanics simulations of proteins based on condensed-phase quantum mechanical calculations. *J Comput Chem* 24:1999–2012.
- Zhang Y, Skolnick J (2004) Scoring function for automated assessment of protein structure template quality. *Proteins* 57:702–710.
- Plaxco KW, Simons KT, Baker D (1998) Contact order, transition state placement, and the refolding rates of single domain proteins. *J Mol Biol* 277:985–994.
- Onufriev A, Bashford D, Case DA (2004) Exploring protein native states and large-scale conformational changes with a modified generalized Born model. *Proteins* 55:383–394.
- Hansmann UHE (1997) Parallel tempering algorithm for conformational studies of Biological Molecules. *Chem Phys Lett* 281:140–150.
- Swendsen RH, Wang JS (1986) Replica Monte Carlo simulation of spin glasses. *Phys Rev Lett* 57:2607–2609.
- Zhang Y, Kihara D, Skolnick J (2002) Local energy landscape flattening: Parallel hyperbolic Monte Carlo sampling of protein folding. *Proteins* 48:192–201.
- Betancourt MR (2005) Efficient Monte Carlo trial moves for polypeptide simulations. *J Chem Phys* 123:174905.
- Kabsch W, Sander C (1983) Dictionary of protein secondary structure: Pattern recognition of hydrogen-bonded and geometrical features. *Biopolymers* 22:2577–2637.
- Cendes T (1988) Nonlinear parameter estimation by global optimization—efficiency and reliability. *Acta Cybernetica* 8:361–370.
- Berman HM, et al. (2000) The Protein Data Bank. *Nucleic Acids Res* 28:235–242.