Protein structure prediction in 2002

Jack Schonbrun, William J Wedemeyer and David Baker*

Central issues concerning protein structure prediction have been highlighted by the recently published summary of the fourth community-wide protein structure prediction experiment (CASP4). Although sequence/structure alignment remains the bottleneck in comparative modeling, there has been substantial progress in fully automated remote homolog detection and in *de novo* structure prediction. Significant further progress will probably require improvements in high-resolution modeling.

Addresses
Howard Hughes Medical Institute and Department of Biochemistry, Box 357350, University of Washington, Seattle, Washington 98165, USA
*e-mail: dabaker@u.washington.edu

Current Opinion in Structural Biology 2002, 12:348–354

0959-440X/02/$ – see front matter

Abbreviations

CASP	Critical Assessment of Protein Structure Prediction
HMM	hidden Markov model
MD	molecular dynamics
PDB	Protein Data Bank
rmsd	root mean square deviation

Introduction

Progress in protein structure prediction is assessed by the community-wide experiment Critical Assessment of Protein Structure Prediction (CASP) [1]. In CASP, sequences of proteins whose experimental structures are soon to be released are made public; computational research groups are then invited to predict those structures from the target sequence and any other publicly available information. Typically, the groups are given a few months to make their predictions using the method(s) of their choice. These methods range from fully automated servers to methods that require significant human intervention.

The results of the fourth protein structure prediction experiment, CASP4, as well as those of several parallel assessments of automated prediction methods (CAFASP [2•], LiveBench [3] and EVA [4]), have been published recently in a supplementary issue of *Proteins: Structure, Function and Genetics* [5]. We focus this review on the results presented in that issue, as they provide an impartial view of progress in the field. By contrast, the performance of current methods is often difficult to assess from traditional research papers, as knowledge of the native structure can influence the prediction in ways not intended by the investigator and negative results are rarely published. The unbiased evaluation of protein structure prediction methods by the CASP experiments organized by John Moult and colleagues has made a very important contribution to the progress of such methods in recent years.

Traditionally, protein structure prediction is divided into three areas, depending on the similarity of the target to proteins of known structure. First, in comparative modeling, one or more template proteins of known structure with high sequence homology to the target sequence are identified. The target and template sequences are aligned, and a three-dimensional structure of the target protein is generated from the coordinates of the aligned residues of the template protein, combined with models for loop regions and other unaligned segments. Ideally, this three-dimensional model would then be refined to bring it closer to the true structure of the target protein. A key outstanding question is whether such refinement actually improves the predicted structure, that is, whether the refined structure is closer to the native structure than the unrefined template structure. Second, if no reliable template protein can be identified from sequence homology alone, the prediction problem is denoted as a fold recognition problem. Here, the primary goal is to identify one or more template protein structures that are consistent with the target sequence, that is, template folds that the target sequence might plausibly adopt. The subsequent protocol is similar to that of comparative modeling: align the sequences and compose a three-dimensional model from the alignment. An outstanding question in fold recognition is whether structural information can improve upon sequence-based methods, that is, whether it enables the identification of remote homologs not detectable by sequence-based methods. Third, if no template structure can be identified with confidence, the target sequence may be modeled using *de novo* (or new fold) prediction methods. An outstanding question in *de novo* prediction is whether such methods can predict structures to a resolution useful for biochemical applications.

Other questions are common to all three areas of protein structure prediction. First, how do the predictions made by fully automated methods compare to the predictions curated by human experts? This question is obviously relevant to genome-scale structure/function predictions, which require high-throughput methods. Second, have protein structure prediction methods improved over the years during which the CASP experiments have been carried out? Third, and most difficult, can we predict the future trajectory of the field, that is, guess the rate of improvement of prediction methods? In this review, we also try to identify the obstacles that are hindering progress at present.

Comparative modeling

The performance of the top eight comparative modeling groups at CASP4 was roughly similar [6•,7]. Obtaining good alignments appears to be the key element of success; loop modeling and further refinement are futile without a

reasonably accurate initial alignment. However, there is a limit to the alignment accuracy that can be achieved for distantly related proteins because the residual sequence similarity that guides the alignment becomes weaker with greater evolutionary distance. Furthermore, because of deletions and insertions, there is not necessarily a strict one-to-one correspondence between sequence positions in the two structures; in fact, different structure/structure comparison methods frequently produce quite different structure-based alignments.

Most alignment methods are based on similarity between the sequences of the target and the template proteins. Alignment accuracy is often improved by using sequence profiles constructed from multiple sequence alignments for the target, the template or both. Sternberg and co-workers [8] generated sequence profiles from structural alignments of remote homologs using the 3D-PSSM program and then compared the target sequence to these profiles. Our group used a new alignment method that rewards matches of the query sequence to residues in the template that are almost always present in the structures of homologs and penalizes insertions or deletions in regions that are consistently ungapped in homologs. In some cases, using multiple proteins as templates for different portions of the sequence yielded improved structures. For example, Venclovas [9] used multiple parent structures obtained by an iterative PSI-BLAST procedure, in which sequences found in a search based on the target sequence were used to initiate new searches to extend the range of matches. Because different segments of a target sequence may align better to different parents, it can be helpful to combine these parents into an overall template; how best to recombine the parents is an open problem.

The results of the top eight groups confirmed that (when there is significant sequence similarity between the target and template sequences) reasonable comparative models can be built consistently. However, as in previous CASP experiments, the assessors found that almost every predicted structure submitted for comparative modeling targets had a higher rmsd to the true native structure than that given by a structural alignment to the best available template [6•]. This highlights the importance and the difficulty of the alignment problem, as well as of refining initial template structures to become closer to the true native structure.

After building a three-dimensional model from an alignment, full-atom refinement methods should (in principle) be able to correct for some alignment errors and to guide models closer to the true native structure. For example, molecular dynamics (MD) protocols that simulate proteins with molecular mechanics potential functions would seem to be well suited to this task [10]. However, none of the top comparative modeling groups used these methods, probably because previous experience indicated that such methods generally make the predicted structures worse rather than better. The refinement of comparative models would seem to provide an excellent test of the application of such simulation methods, because the initial models are generally quite close to the true structure and, hence, sampling is not as difficult as in, for example, *de novo* structure prediction. A systematic comparison of experimental structures with models refined using current MD/molecular mechanics methods should be encouraged, as this would probably identify the reasons why such protocols perform poorly at present and how they should be improved. Better refinement methods would, in turn, lead to better alignments, as they might identify (and possibly correct) alignment errors.

Table 1

Current snapshot of the ranking of prediction servers conducted by LiveBench.

		Sensitivity		Specificity		Added value	
Server	Type	Easy	Hard	All	Hard	Easy	Hard
Pcons2	Consensus	6	4	2	2	3	3
ShotGun on 5	Consensus	1	2	4	4	7	5
ShotGun on 3	Consensus	2	1	1	1	2	2
Shotgun-INBGU	Threading	3	3	3	3	4	1
INBGU	Threading	7	5	6	9	5	6
Fugue3	Threading	14	8	9	8	15	9
Fugue2	Threading	12	7	8	7	10	8
Fugue1	Threading	17	14	14	11	16	15
mGenTHREADER	Threading	8	11	16	13	6	11
GenTHREADER	Threading	13	12	17	15	8	13
3D-PSSM	Threading	5	10	12	12	12	10
ORFeus	Sequence	4	6	7	6	1	4
FFAS	Sequence	9	9	5	5	9	7
Sam-T99	Sequence	10	15	13	16	11	16
Superfamily	Sequence	15	13	11	10	17	12
ORF-BLAST	BLAST	11	16	10	14	14	14
PDB-BLAST	BLAST	16	17	15	17	13	17
BLAST	BLAST	18	18	18	18	18	18

The LiveBench program conducts an automated weekly evaluation of protein structure prediction servers. There are a large number of possibilities for ranking the servers and this table presents a compilation of just a few of them. The results change every week after new proteins are added to the pool of targets. The table shows the results for eighteen servers grouped by type. The top three servers are consensus servers, which create only jury predictions based on the results obtained by other prediction servers. The next group of eight servers (actually various versions of servers developed by four research groups) represents more traditional 'threading' approaches, which utilize the structure of the template protein in their scoring function. The next group of four servers utilizes only sequence information on the template protein (and its derivatives) and could thus also be used to find homologs of a query protein with unknown structure (they can be used not only for structure prediction). The last group of three represents simple servers, which utilize exclusively the components of the BLAST family of tools. The presented ranking is divided into three main categories: sensitivity of the server, specificity of the server (reliability of the reported score) and added value (reflects the value of a server when compared to the rest). The ranking is also divided into hard and easy (found by PDB-BLAST with significant score) targets (or all in case of specificity). The details of the evaluation methods are described in the LiveBench paper [3] and on the LiveBench home page. The main results from the latest round are that the consensus methods are very valuable in judging the quality of the results produced by the community of prediction methods and the confirmation of the competitiveness of sequence-only methods with threading methods.

Despite the limitations of current methods, comparative modeling can still be very useful to biologists. The best predicted regions are often biologically important [8], because these are the most structurally conserved by evolution. Thus, comparative modeling often predicts accurately 'the parts that matter'.

Fold recognition

Fold recognition has been transformed by the advent of powerful sequence-based methods (such as PSI-BLAST [11] and hidden Markov model [HMM]-based methods [12]) that exploit evolutionary information to match sequences to known structures. In particular, PSI-BLAST is probably the most widely used tool for remote homolog detection in molecular biology today and often identifies matches missed by earlier sequence-based methods, such as BLAST. Interestingly, the CAFASP and LiveBench tests (Table 1) show that PSI-BLAST alone performs worse at fold recognition than several other sequence-based and sequence-and-structure-based methods. Nevertheless, PSI-BLAST (or a closely related method) forms the core of many of the most successful methods. For example, Koretke *et al.* [13] extended the range of PSI-BLAST, using iterative searches to expand the number of hits and then PSI-BLAST searches back to the query sequence to 'back-validate' hits with low significance.

Methods that used structural information generally did so by reducing the structure to a one-dimensional string of properties, such as secondary structure or solvent accessibility. Karplus *et al.* [12] improved the performance of their HMM-based sequence comparison method by adding a track to compare predicted and known secondary structures. One of the most successful automated servers, 3D-PSSM, evaluates the match of the query sequence to both the sequence of the template and the solvent accessibility pattern of the template structure [8]. By contrast, GenThreader evaluates alignments found through PSI-BLAST using the residue contacts in the three-dimensional structure [14]. 3D-PSSM goes beyond sequence and structural information by comparing key words in the annotations of the proteins and assigning a high score to matches with significant key word similarity.

The question of whether structural information helps significantly is still unresolved — two very successful servers, ORFeus (http://grdb.bioinfo.pl) and FFAS [15], use only sequence information (see Table 1). Although it may seem counterintuitive that structural information does not necessarily provide significant improvement, it must be kept in mind that current sequence-based methods use evolutionary information from both the target and template sequences, which helps to bridge the evolutionary distance between the two sequences. Sequence information would be expected to be insufficient to match sequences to structures of evolutionarily unrelated homologs, but even methods that utilize structural information are rarely successful in such cases [13,16•].

As in comparative modeling, the best groups achieved a similar level of performance, despite considerable differences between the methods. This is probably because they all detect distant sequence similarities at some level, augmented, to differing extents, by structural and functional similarity. There is clearly a fundamental limitation to methods that rely on conservation of structural properties across large evolutionary distances; structural changes (e.g. sidechain packing, sheet twist and helix orientations) can invalidate the implicit assumption that the target sequence has a low energy when threaded on to the template structure.

Humans versus automated servers

The performance of fully automated servers on CASP4 targets was assessed in the CAFASP2 experiment. An important conclusion from comparing the CASP/CAFASP results was that fully automated methods are approaching the accuracy of human-curated predictions. Many top fold recognition groups in CASP4 were also represented by servers and the performance of the fully automated methods was only slightly inferior to that of their human-curated counterparts. Indeed, the CAFASP authors noted that the fully automated servers outperformed ~70% of the human participants [2•], although, to be fair, many groups may have been more interested in testing the basic concepts underlying their methods than achieving optimal performance.

There is still a gap between fully automated servers and the best performing groups, even in fold recognition. The most spectacular illustration of how a human expert can outperform automated servers was the prediction of target 104 (hypothetical protein HI0065 from *Haemophilus influenzae)* by Murzin and Bateman [17], who were able not only to find a homolog but also to deduce that several secondary structure elements needed to be added/deleted to agree with the target sequence. Their remarkable model possessed a novel topology and a rmsd from the native structure of only 2.15 Å (Figure 1).

Humans have an advantage over fully automated servers in that they can integrate a wide variety of information about the target. This integration can be mimicked, to some extent, computationally by taking a consensus of the predictions of different automated servers. Such a consensus has been implemented in the Pcons2 server [18]. A key result of the CAFASP/LiveBench experiments is that such consensus servers outperform all individual automated servers (Table 1). In particular, consensus servers provide more accurate measures of prediction confidence, which are vital if predicted structures are to be used extensively by the biological community.

The LiveBench assessment experiment, organized by Rychlewski and co-workers (http://bioinfo.pl/LiveBench/), is an innovative automated ranking of automated fold recognition servers. As new structures are entered into the

Figure 1

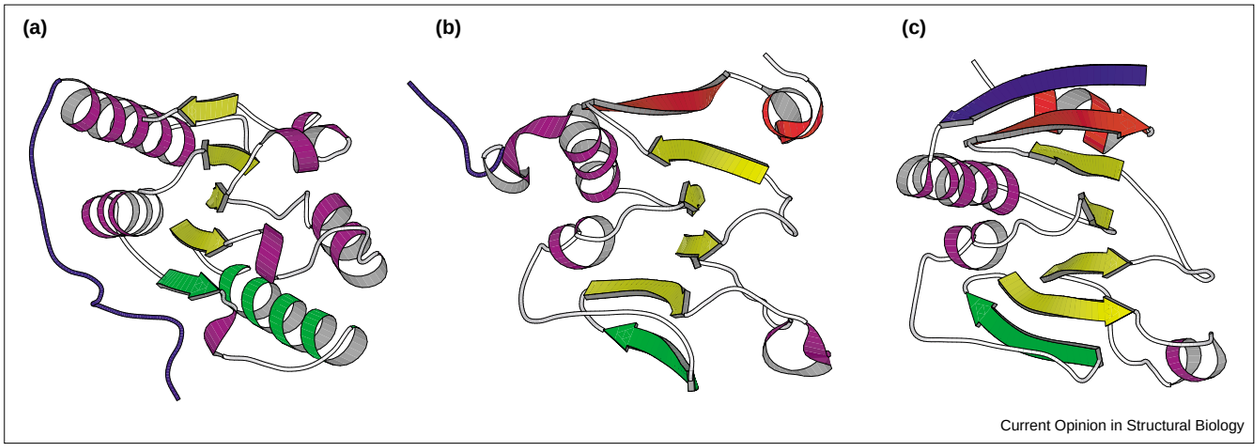

The conversion from template to prediction for target 104 (hypothetical protein HI0065 of *H. influenzae*) by Murzin and Bateman [17]. (a) The template (PDB code 1NSF) from which modeling began. Because of incompatibilities in the secondary structure predicted for target 104 and the template, several significant changes were made to the topology of the submitted model, shown in (b). A helix and strand (shown at the bottom of the figure in green) were replaced by a loop and strand going in the opposite direction. At the C terminus of the protein (shown at the top of the figure in red), an extra strand was added to the central sheet and a helix was added to the end of the chain. All of these changes are seen in the native structure (c). The coordinates of (b,c) were taken from the CASP web site (http://predictioncenter.llnl.gov/casp4); figures prepared with MOLSCRIPT [26].

PDB, their sequences are submitted to over a dozen participating servers. The resulting predictions are then compared to the native structure and the results tabulated on the web site. This continuous, large-scale evaluation of automated servers provides a statistically significant comparison of different methods (Table 1). (As an aside, LiveBench, despite its unquestioned importance for the field of protein structure prediction, is currently threatened by the uncertainties of science funding in Poland. Hopefully, alternative sources of funding can be identified, should the worst be realized.)

'*De novo*'/new fold methods

Despite sporadic reports of solutions to the *de novo* protein folding problem, the CASP2 assessor concluded that there had not been significant progress towards a solution and held out little hope for the future [19]. It was noted that large-scale structural genomics initiatives seemed likely to solve high-resolution structures for representatives of most families of naturally occurring proteins before useful models could be generated by computational methods.

Nevertheless, the CASP experiments have clearly documented significant improvements in *de novo* folding from CASP2 to CASP3, and from CASP3 to CASP4 [20•]. The Rosetta method developed in our group made particularly good *de novo* predictions in CASP4 [21•,22], as illustrated in Figure 2. The overall topology of large segments of domains or (in some cases) entire domains was predicted consistently and accurately. In some cases, the models are of sufficient quality to provide clues about protein function, which has motivated large-scale modeling of protein families using Rosetta.

We believe this improved performance stems from both the model of folding that underlies Rosetta and the way in which this model is implemented. A key element is the separation of local and nonlocal interactions. Rosetta is based on a picture of protein folding in which short segments of the chain independently sample distinct distributions of local conformations, biased by their local sequences. Folding to the native state occurs when these segments have relative orientations and conformations that allow low free energy nonlocal interactions to form throughout the protein. To implement this picture of folding as a computational algorithm, we approximate the distribution of local structures sampled by a given segment during folding by the distribution of local structures adopted by such segments in known protein structures. Thus, during a Rosetta folding simulation, each nine- and three-residue segment of the protein chain flickers between the different local structures that are consistent with its local sequence, while the nonlocal interactions are slowly optimized using a Monte Carlo search procedure. A low-resolution model of the nonlocal interactions dominated by hydrophobic burial and strand pairing is used until near the end of the simulation, when a rotamer-based explicit sidechain model with Lennard–Jones interactions is introduced. The entire procedure resembles a jigsaw puzzle in which each piece can have several alternative shapes. An advantage of Rosetta's fragment approach over traditional molecular mechanics approaches is that current potential functions are probably not accurate enough to capture all the physical interactions that determine the local conformational preferences, for example, specific sidechain–mainchain hydrogen bonds. In contrast to other prediction protocols that fix the local structure before

Figure 2

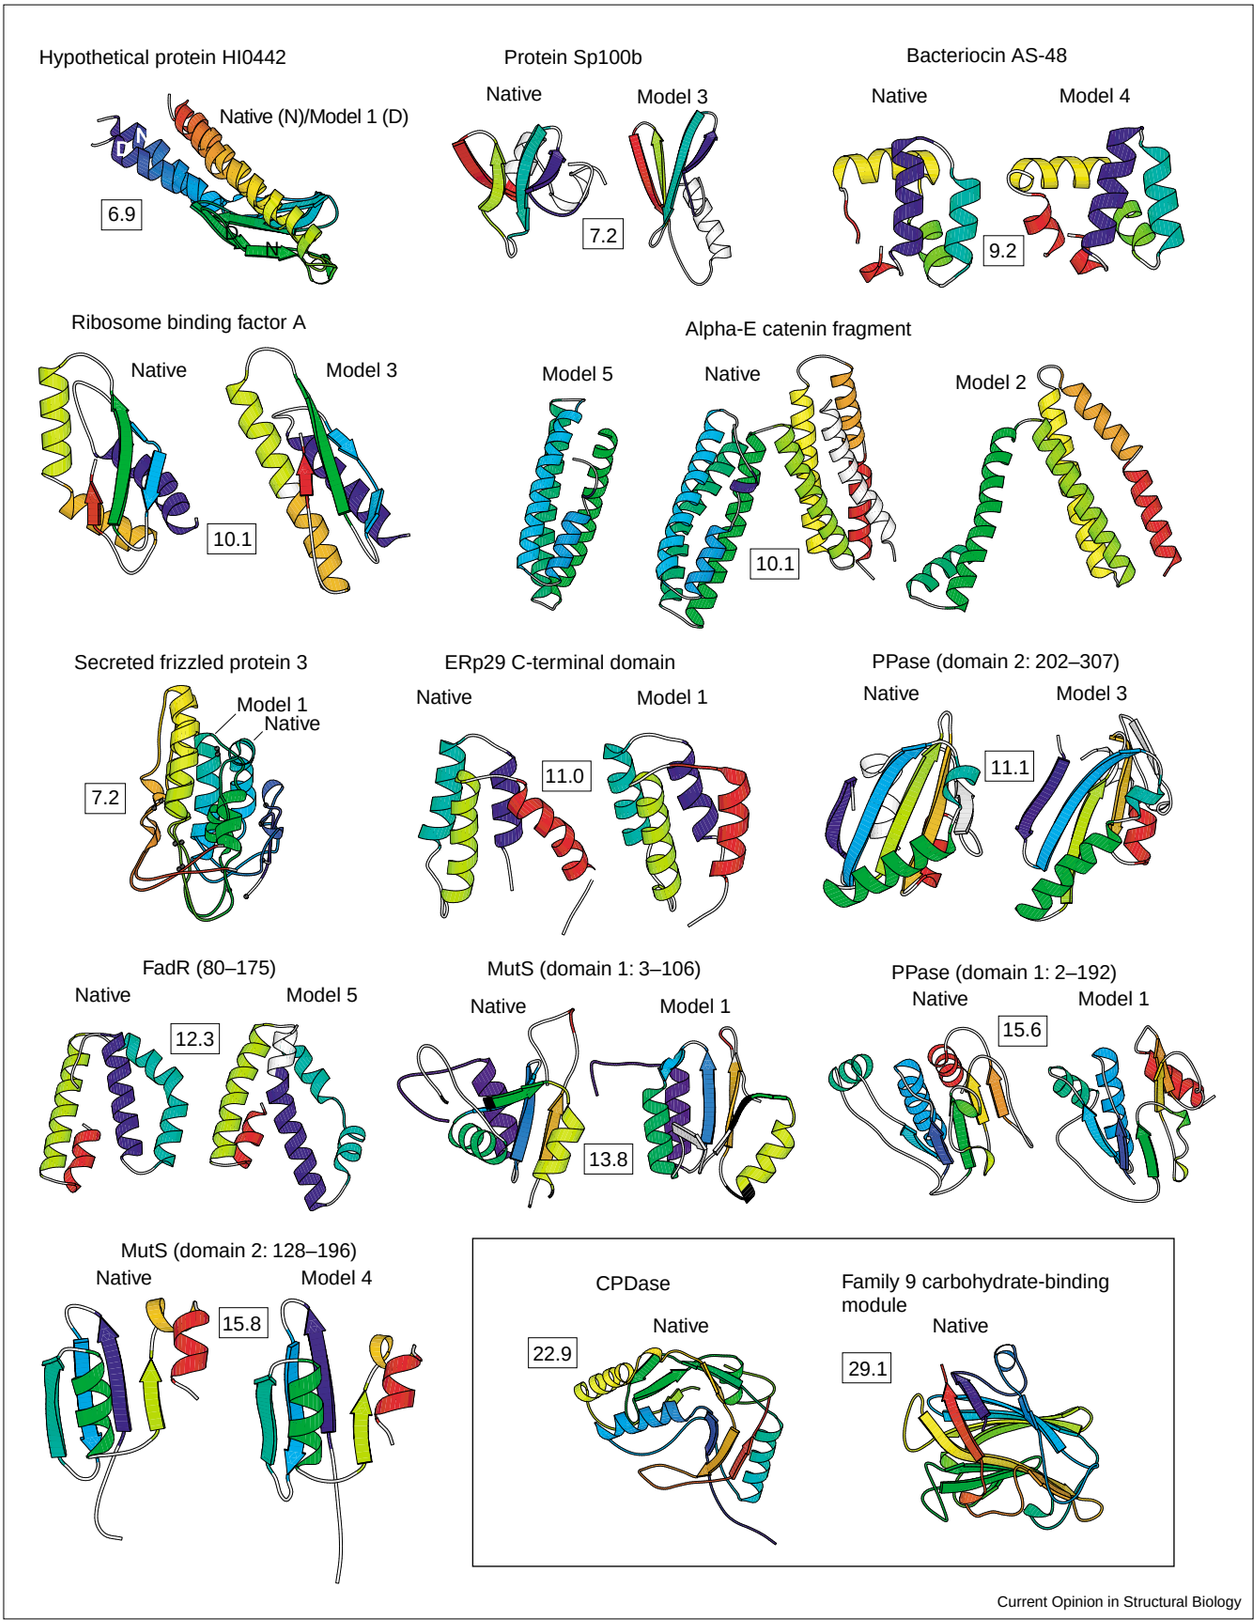

Figure 2 legend

Comparison of the best Rosetta prediction with the native structure for selected CASP4 targets. The native structure is on the left and the predicted model structure is on the right (except where they are superimposed); the chain is colored to guide visualization from blue (N terminus) to red (C terminus). The number in a box next to each protein is the native contact order. For the two structures with contact orders over 20 (shown in the box), Rosetta failed to generate any structures with the correct topology.

predicting the long-range structure, large changes in local structure can occur even quite late in the simulation in response to changes in the nonlocal interactions.

Although Rosetta consistently finds structures with topologies close to the native for many proteins, there is much room for improvement. It is still difficult to discern the native structure from other 'protein-like' decoys. The models in Figure 2 represent the best of the five structures submitted for CASP4; some of the other four models were often quite different from the native structure. The probability of Rosetta making a good prediction correlates with the average sequence separation between residues close in space in the native protein — the contact order (see Figure 2). High contact order proteins fold much more slowly than low contact order proteins (because there is a large loss in chain entropy before the formation of substantial numbers of attractive native interactions); by analogy, computational folding to the native structure may likewise take much more time. Furthermore, even when they are topologically correct, the structures generated by Rosetta are relatively low in resolution and are clearly not suitable for applications that require high-resolution detail, such as drug design and understanding enzyme action. The refinement of models using full-atom representations of the chain is an important area of current work and progress here will be necessary for both improving the accuracy of the models and increasing their reliability.

Most of the other methods deemed successful for new fold targets also built up structures from fragments of other structures in the PDB. Jones' FRAGFOLD [23] algorithm selected fragments from a library of supersecondary structure elements using a threading potential. One difference between FRAGFOLD and Rosetta is that the former does not directly utilize sequence information when picking fragments and thus may be somewhat less sensitive to turn motifs and so on. Skolnick *et al.* [24] used a threading technique to predict possible long-range contacts, which were used to guide a subsequent Monte Carlo search on a regular lattice. Contact predictions were also made by Fariselli *et al.* [25], who used a neural network applied to correlated mutation information.

Coalescence of structure prediction problems

The traditional division of structure prediction categories has broken down to some extent over the past several years. The division between fold recognition and comparative modeling has become blurred as comparative modeling methods use more and more sensitive methods to identify templates, and fold recognition techniques have improved alignments by incorporating detailed sequence information. Ultimately, it will probably be useful to build models of all plausible matches of amino acid sequences to structural templates for evaluation. With the improvement in performance of *de novo* structure prediction, predictions for proteins only very distantly related to proteins of known structure can, in some cases, be comparable to those produced by fold recognition methods [21•]. Finally, long loop modeling in comparative modeling is essentially a small *de novo* prediction problem. Over the next few years, we should see combined methods that use methodologies developed in all three areas to model all portions of a protein sequence at the highest resolution possible given the available information.

Conclusions

In closing, we return to the questions posed at the beginning of this review. Can existing refinement methods bring comparative models closer to the true structure? The answer appears to be 'no'; comparative modeling is still limited by the accuracy of the target/template alignments. Present refinement methods do not overcome misalignments, implying that progress will come through either the surveying of many possible alignments or better refinement methods. Does structural information help in identifying remote homologs for fold recognition? The answer appears to be a qualified 'yes', although the improvement is relatively minor for automated methods. Some sequence-based methods perform almost as well as methods that utilize structural information and detailed properties of the three-dimensional structure (beyond secondary structure and solvent accessibility) do not seem to contribute significantly. Furthermore, essentially all of the best methods rely primarily on sequence information; structural information is used to refine, not supplant, sequence-based methods. Lastly, *de novo* structure prediction methods have progressed to the point at which models with correct overall topology can be obtained a reasonable fraction of the time.

Regarding the relative performance of automated servers versus human experts, the CASP4/CAFASP2 experiments demonstrated clearly that fully automated fold recognition methods can produce results that are nearly as good as those produced by almost all human/expert methods. This bodes well for high-throughput genome analysis.

Finally, what of the future trajectory of structure prediction methods? Although several particular challenges remain,

all three categories of structure prediction seem to be blocked by a common obstacle, that of high-resolution modeling. The models often have the correct overall topology, but are incorrect in higher resolution details, for example, the loops in comparative models and recognized folds, and the packing and other details of *de novo* predictions. These details are clearly essential to determining the native fold of a protein and its biochemical interpretation. In order for predictions to approach the reliability of experimentally determined structures, it will be necessary to model the interactions contributing to protein stability more accurately (particularly the close complementarity of sidechain packing) and to develop improved sampling methods that can find such well-packed conformations. The ability to consistently model these higher resolution features of protein structure would contribute to all areas of protein structure prediction. Being so close to the problem ourselves, it is difficult to estimate the rate of progress, but we permit ourselves the guardedly optimistic prediction that significant progress towards high-resolution predictions of protein structures will occur within the next several years — by CASP6 in 2004, if not CASP5.

Acknowledgements

We thank Leszek Rychlewski for Table 1, Rich Bonneau for Figure 2, John Moult and his colleagues for conducting the CASP experiments, and Kevin Karplus and members of the Baker group for very helpful comments on the manuscript.

References and recommended reading

Papers of particular interest, published within the annual period of review, have been highlighted as:

- • of special interest
- •• of outstanding interest

1. Moult J, Fidelis K, Zemla A, Hubbard T: Critical assessment of methods of protein structure prediction (CASP): round IV. *Proteins* 2001, 45(suppl 5):2-7.

2. • Fischer D, Elofsson A, Rychlewski L, Pazos F, Valencia A, Rost B, Ortiz AR, Dunbrack RL Jr: CAFASP2: the second critical assessment of fully automated structure prediction methods. *Proteins* 2001, 45(suppl 5):171-183.
A comparison of the ability of automated servers to recognize folds.

3. Bujnicki JM, Elofsson A, Fischer D, Rychlewski L: LiveBench-2: large-scale automated evaluation of protein structure prediction servers. *Proteins* 2001, 45(suppl 5):184-191.

4. Rost B, Eyrich VA: EVA: large-scale analysis of secondary structure prediction. *Proteins* 2001, 45(suppl 5):192-199.

5. Lattman EE: CASP4 editorial. *Proteins* 2001, 45(suppl 5):1.

6. • Tramontano A, Leplae R, Morea V: Analysis and assessment of comparative modeling predictions in CASP4. *Proteins* 2001, 45(suppl 5):22-38.
The assessors' summary of the comparative modeling category in CASP4.

7. Marti-Renom MA, Madhusudhan MS, Fiser A, Rost B, Sali A: Reliability of assessment of protein structure prediction methods. *Structure* 2002, 10:435-440.

8. Bates PA, Kelley LA, MacCallum RM, Sternberg MJ: Enhancement of protein modeling by human intervention in applying the automatic programs 3D-JIGSAW and 3D-PSSM. *Proteins* 2001, 45(suppl 5):39-46.

9. Venclovas C: Comparative modeling of CASP4 target proteins: combining results of sequence search with three-dimensional structure assessment. *Proteins* 2001, 45(suppl 5):47-54.

10. Daggett V: Long timescale simulations. *Curr Opin Struct Biol* 2000, 10:160-164.

11. Altschul SF, Madden TL, Schaffer AA, Zhang J, Zhang Z, Miller W, Lipman DJ: Gapped BLAST and PSI-BLAST: a new generation of protein database search programs. *Nucleic Acids Res* 1997, 25:3389-3402.

12. Karplus K, Karchin R, Barrett C, Tu S, Cline M, Diekhans M, Grate L, Casper J, Hughey R: What is the value added by human intervention in protein structure prediction? *Proteins* 2001, 45(suppl 5):86-91.

13. Koretke KK, Russell RB, Lupas AN: Fold recognition from sequence comparisons. *Proteins* 2001, 45(suppl 5):68-75.

14. Jones DT: GenTHREADER: an efficient and reliable protein fold recognition method for genomic sequences. *J Mol Biol* 1999, 287:797-815.

15. Rychlewski L, Jaroszewski L, Li W, Godzik A: Comparison of sequence profiles. Strategies for structural predictions using sequence information. *Protein Sci* 2000, 9:232-241.

16. • Sippl MJ, Lackner P, Domingues FS, Prlic A, Malik R, Andreeva A, Wiederstein M: Assessment of the CASP4 fold recognition category. *Proteins* 2001, 45(suppl 5):55-67.
The assessors' summary of the fold recognition category in CASP4.

17. Murzin AG, Bateman A: CASP2 knowledge-based approach to distant homology recognition and fold prediction in CASP4. *Proteins* 2001, 45(suppl 5):76-85.

18. Lundstrom J, Rychlewski L, Bujnicki J, Elofsson A: Pcons: a neural-network-based consensus predictor that improves fold recognition. *Protein Sci* 2001, 10:2354-2362.

19. Lesk AM: CASP2: report on ***ab initio*** predictions. *Proteins* 1997, 41(suppl 1):151-166.

20. • Venclovas C, Zemla A, Fidelis K, Moult J: Comparison of performance in successive CASP experiments. *Proteins* 2001, 45(suppl 5):163-170.
An evaluation of progress in successive CASP experiments.

21. • Lesk AM, Lo Conte L, Hubbard TJ: Assessment of novel fold targets in CASP4: Predictions of three-dimensional structures, secondary structures, and interresidue contacts. *Proteins* 2001, 45(suppl 5):98-118.
The assessors' summary of new fold prediction in CASP4.

22. Bonneau R, Tsai J, Ruczinski I, Chivian D, Rohl C, Strauss CE, Baker D: Rosetta in CASP4: progress in ***ab initio*** protein structure prediction. *Proteins* 2001, 45(suppl 5):119-126.

23. Jones DT: Predicting novel protein folds by using FRAGFOLD. *Proteins* 2001, 45(suppl 5):127-132.

24. Skolnick J, Kolinski A, Kihara D, Betancourt M, Rotkiewicz P, Boniecki M: ***Ab initio*** protein structure prediction via a combination of threading, lattice folding, clustering, and structure refinement. *Proteins* 2001, 45(suppl 5):149-156.

25. Fariselli P, Olmea O, Valencia A, Casadio R: Progress in predicting inter-residue contacts of proteins with neural networks and correlated mutations. *Proteins* 2001, 45(suppl 5):157-162.

26. Kraulis PJ: MOLSCRIPT: a program to produce both detailed and schematic plots of protein structures. *J Appl Crystallogr* 1991, 24:946-950.