PROTEINS: Structure, Function, and Genetics, Suppl. 1:50–58 (1997)

Evaluation of Comparative Protein Structure Modeling by MODELLER-3

Roberto Sánchez and Andrej Šali*
The Rockefeller University, New York, New York

ABSTRACT **We evaluate homology-derived 3D models of dihydrofolate reductase (DFR_1), phosphotransferase enzyme IIA domain ($PTE2A_3$), and mouse/human UBC9 protein ($UBC9_{24}$) which were submitted to the second Meeting on the Critical Assessment of Techniques for Protein Structure Prediction (CASP). The DFR_1 and $PTE2A_3$ models, based on alignments without large errors, were slightly closer to their corresponding X-ray structures than the closest template structures. By contrast, the $UBC9_{24}$ model was slightly worse than the best template due to a misalignment of the N-terminal helix. Although the current models appear to be more accurate than the models submitted to the CASP meeting in 1994, the four major types of errors in side chain packing, position, and conformation of aligned segments, position and conformation of inserted segments, and in alignment still occur to almost the same degree. The modest improvement probably originates from the careful manual selection of the templates and editing of the alignment, as well as from the iterative realignment and model building guided by various model evaluation techniques. This iterative approach to comparative modeling is likely to overcome at least some initial alignment errors, as demonstrated by the correct final alignment of the C terminus of DFR_1. Proteins, Suppl. 1:50–58, 1997.**

Key words: evaluation; comparative protein modeling; Modeller

INTRODUCTION

Protein modelers were challenged for the second time to model sequences without available 3D structures and to submit them to the CASP meeting in December 1996 (CASP; URL http://PredictionCenter.llnl.gov/). At the same time, the 3D structures were being determined by X-ray crystallography and NMR methods. Because the experimentally determined structures were only released at the meeting, it was possible to test the modeling methods objectively. A summary of all comparative models submitted to CASP2 can be found elsewhere in this issue (A.C.R. Martin et al.).

We submitted homology-derived models of three proteins: DFR_1, $PTE2A_3$, and $UBC9_{24}$; the subscript indicates the target sequence number assigned by the organizers of CASP2. All three structures have been determined by X-ray crystallography: DFR_1 at 2.6 Å resolution and R factor of 18% (U. Pieper and O. Herzberg, in preparation), $PTE2A_3$ at 2.4 Å resolution (K. Huang and O. Herzberg, in preparation) and $UBC9_{24}$ at 2.0 Å resolution and R factor of 16% (H. Tong and T. Sixma, in preparation). These three target sequences were chosen because they have a relatively low, <43% sequence identity with their templates. In this range of sequence similarity, the largest errors in comparative modeling due to misalignments begin to appear.[1,2] It is important to concentrate on this range of sequence similarity because most of the detectable related sequence–structure pairs are related at less than 40% sequence identity level,[3] despite earlier indications to the contrary.[4]

Our approach to comparative protein structure modeling is based on satisfaction of spatial restraints and is implemented in program Modeller.†[5] This program can be used in all stages of typical comparative modeling: Finding suitable template structures in the PDB,[6] aligning them with the sequence to be modeled, calculating the 3D model, and evaluating the model. Comparative protein modeling was recently reviewed.[7,8]

†Modeller is available at URL http://guitar.rockefeller.edu:pub/modeller and also as part of Quanta and InsightII (MSI, San Diego, CA. E-mail: blp@msi.com).

Abbreviations: DFR_1, *Haloferax volcanii* dihydrofolate reductase; $PTE2A_3$, *Mycoplasma capricolum* phosphotransferase enzyme IIA domain; $UBC9_{24}$, mouse/human UBC9 protein; NMR, nuclear magnetic resonance; PDB, Brookhaven Protein Data Bank; RMSD, root-mean-square deviation; 3D, three-dimensional; CASP, critical assessment of techniques for protein structure prediction.

Contract grant sponsor: National Institutes of Health; Contract grant number: GM 54762; Contract grant sponsor: National Science Foundation; Contract grant number: BIR-9601845.
*Correspondence to: A. Šali, The Rockefeller University, 1230 York Avenue, New York, NY 10021.
E-mail: sali@rockvax.rockefeller.edu
Received 8 May 1997; Accepted 26 August 1997

In this article, we briefly describe the modeling method and then concentrate on evaluation of the three submitted models. In particular, we discuss the question of whether or not the models are generally closer to the X-ray structure being modeled than the template structures.

METHODS

The first step in comparative modeling of the three target proteins was identification of potential template structures. This was followed by several cycles of template selection, target-template alignment, model building, and model evaluation. The aim of the iteration was to minimize the errors in the model reported by various model evaluation techniques. This iterative process, including careful manual selection of the templates and editing of the alignments, is the main difference between the current approach and that followed two years ago for the CASP1 meeting.[2] The final alignments and 3D models are available from the CASP2 Web site at URL http://PredictionCenter.llnl.gov/CASP/CM_results/.

Template Selection

Proteins that have known 3D structure and are similar to the sequences being modeled had to be identified. This was achieved by searching a set of sequences representative of the whole PDB (July 1, 1996) [6], using the SEQUENCE_SEARCH command of Modeller.[9] The representative set of protein structures included 916 chains whose sequence identity was less than 30% to any other chain in the set. The final templates were as follows: For DFR_1, 4DFR-B (30%, 1.4 Å, 91%), 3DFR (24%, 1.5 Å, 93%), and 8DFR (22%, 1.6 Å, 94%); for $PTE2A_3$, 1GPR (43%, 1.3 Å , 94%) and 1F3G (36%, 1.1 Å , 94%); and for $UBC9_{24}$, 1AAK (35%, 1.1 Å, 90%) and 2UCE (30%, 1.2 Å, 90%). The numbers in the parentheses are the percentage sequence identity, RMSD for Cα atoms, and the fraction of the equivalent Cα atoms. These were all obtained from pairwise template-target least-squares superpositions with a 3.5 Å cutoff.

Target-Template Alignment

Initial multiple template-target alignments were obtained by aligning the target sequences with the prealigned template structures, using the ALIGN2D command of Modeller.[9] This command implements a global dynamic programming[10] algorithm with a variable gap-penalty function that depends on the structural context of an insertion or a deletion (R. Sánchez and A. Šali, in preparation). The gap penalty is constructed such that insertions and deletions are less preferred within helices and sheets, buried regions, straight segments, and also between two residues that are distant in space. The alignments also depended on a 20 × 20 amino acid residue substitution matrix that was derived from 105 structure–structure alignments.[4] The initial calculated alignments were edited by hand as appropriate (see below).

Model Building

The 3D models containing all nonhydrogen atoms were obtained automatically by satisfying restraints on many distances, angles, and dihedral angles.[4,5] Spatial restraints were extracted from the alignment of the target sequence with the template structures[4,5] and from the Charmm-22 force field.[11] The whole model, including backbone, side chains, loops, and insertions, was build in one optimization. Conformation of the regions aligned with the templates was based mostly on the template structures, while the insertions were restrained mostly by the preferences of the different residue types for the different areas of the Ramachandran plot.

Model Evaluation

The models had to satisfy most restraints used to calculate them, especially the stereochemical restraints. These tests were done by the Modeller ENERGY command,[9] the Procheck program,[12] and the WhatCheck program.[13] The most important evaluation was done by "energy" profiles calculated by ProsaII, which relies on statistical potentials involving single residues and pairs of residues.[14] Additional evaluation was done by "energy" profiles calculated from a new set of statistical potentials involving pairs of atoms.[15] Side chain packing was checked by calculating cavities in the core of a protein, using the Quanta Protein Health module (MSI, San Diego, CA). If any of the model evaluation tools indicated an error in the model, the model was changed manually. For example, side chains were manually repositioned to eliminate a cavity in the core. Another example is a selection of different templates and editing of the alignment around the region with a bad ProsaII profile, followed by another round of the automated model building.

RESULTS AND DISCUSSION

Although the DFR_1, $UBC9_{24}$, and $PTE2A_3$ models have good stereochemistry, they have errors in four other categories: Distortions or shifts of a region that is aligned correctly with the templates (e.g., loops, helices, strands); errors in side chain packing; distortions or shifts of a region that does not have an equivalent segment in any of the templates (e.g., inserted loops); and distortions or shifts of a region that is aligned incorrectly with the templates (e.g., loops and larger segments with low sequence identity to the templates). Examples of these errors are described in the following sections. We also discuss the lessons learned from this experiment with respect to automated template mimicking in different regions of a model; the cycle of template selection, alignment, model building, and model evaluation;

and the relative overall similarity of a model and the templates to the target X-ray structure.

Stereochemistry of the Models

The stereochemical features of the models, such as those evaluated by the Procheck[12] and WhatCheck[13] programs, are comparable to those in the high resolution X-ray structures. These features include bond lengths, angles, improper dihedral angles, position of residues in the Ramachandran plot, peptide bonds planarity, Cα tetrahedral distortion, nonbonded interactions, hydrogen bond energies, and closeness of side chain dihedral angles to ideal values. It is not surprising that the models are stereochemically correct since they were calculated partly by optimizing the stereochemical features as encoded in the Charmm-22 force field.[11]

Errors in Side Chain Packing

The side chain rotamers were predicted surprisingly inaccurately. For example, the percentage of $\chi 1$ angles for DFR_1, $PTE2A_3$, and $UBC9_{24}$ predicted within 30° of the target values was 42%, 48%, and 65%, respectively. Since at least the $UBC9_{24}$ X-ray structure has been refined at a high resolution of 2 Å and an R factor of 16%, the low prediction accuracy must reflect significant problems with our side chain modeling procedure in this range of backbone and side chain similarities. However, the mistakes made were not trivial because the models followed their templates for conserved and similar side chains, because the model rotamers were not distorted, and because the cavities in the models were not larger than those in the X-ray structures. It is not clear what kind of improvements are needed beyond a self-evident need for a more accurate energy function and perhaps a better optimizer.

The difficulty of the side chain modeling problem in this range of sequence similarity is illustrated by the fact that the template and target X-ray structures have different rotamers for up to 45% of the conserved residues. For example, DFR_1 has 125 residues with at least one side chain dihedral angle, 29 of which are conserved in one of the templates (PDB code 3DFR), but 12 of these occur in different rotamer states. A systematic analysis of this phenomenon, based on highly refined structures, would be useful. If the target and template X-ray structures are accurate and the finding proves to be general, this indicates that the side chains should be modeled on the basis of more general physical principles[16–19] rather than by mimicking the templates,[20,21] especially when the backbones of the target and the template have an RMSD larger than 2 Å. An additional complication for the evaluation of side chain models is that for the two targets refined at a low resolution of 2.6 Å (DFR_1) and 2.4 Å ($PTE2A_3$), it is not clear that all the differences between the models and the X-ray structures are due to the mistakes in the modeling procedure.[22]

Distortions or Shifts in Correctly Aligned Regions: Template Mimicking in Different Regions of a Model

For all three models, at least two template structures were used. Thus, it was possible to determine how frequently the automated model building selected the best template for a given segment where the templates shared different degrees of structural similarity with the target structure. The ability to pick locally optimal templates is important because it allows the model to be overall closer to the correct structure than any of the individual templates.

The distances of the positions of the Cα atoms of the model and the templates from the equivalent atoms in the superposed target X-ray structure are shown for DFR_1 and $UBC9_{24}$ in Figure 1. For the correctly aligned regions, the model always follows one of the templates. When two templates differ in a given correctly aligned region, the model generally follows the template that is structurally closer to the experimental structure: Six such segments of at least three residues with distances between the templates of at least 1 Å occur in the DFR_1 and $UBC9_{24}$ models. For the correctly aligned regions, there are no examples of the model following a suboptimal template. As a consequence, the model is generally closer overall to the experimental structure than any of the templates (see also Fig. 4). However, for a given region, model building does not result in a model that is better than the best template in that region (Fig. 1).

These observations are a direct consequence of the form of the homology-derived distance restraints.[4,5] The restraints are expressed as probability density functions. When several templates are aligned with a given segment in the target sequence, a restraint on an inter- or intrasegment distance has a multimodal shape with the peaks corresponding to the equivalent distances in the templates, not to the average distance. The heights and the widths of the peaks are determined by the overall and local sequence similarities between the templates and the target sequence, such that the model is most likely to resemble the template with the most similar sequence. This means that the model is generally closer to one or the other template by construction. In order to allow for the modeling of distortions or shifts relative to the template structures, a scoring function that guides the model in the correct direction from the template to the target structure is necessary. A combination of homology-derived restraints with atom based statistical potentials[15,23–25] is perhaps one way of achieving this aim.

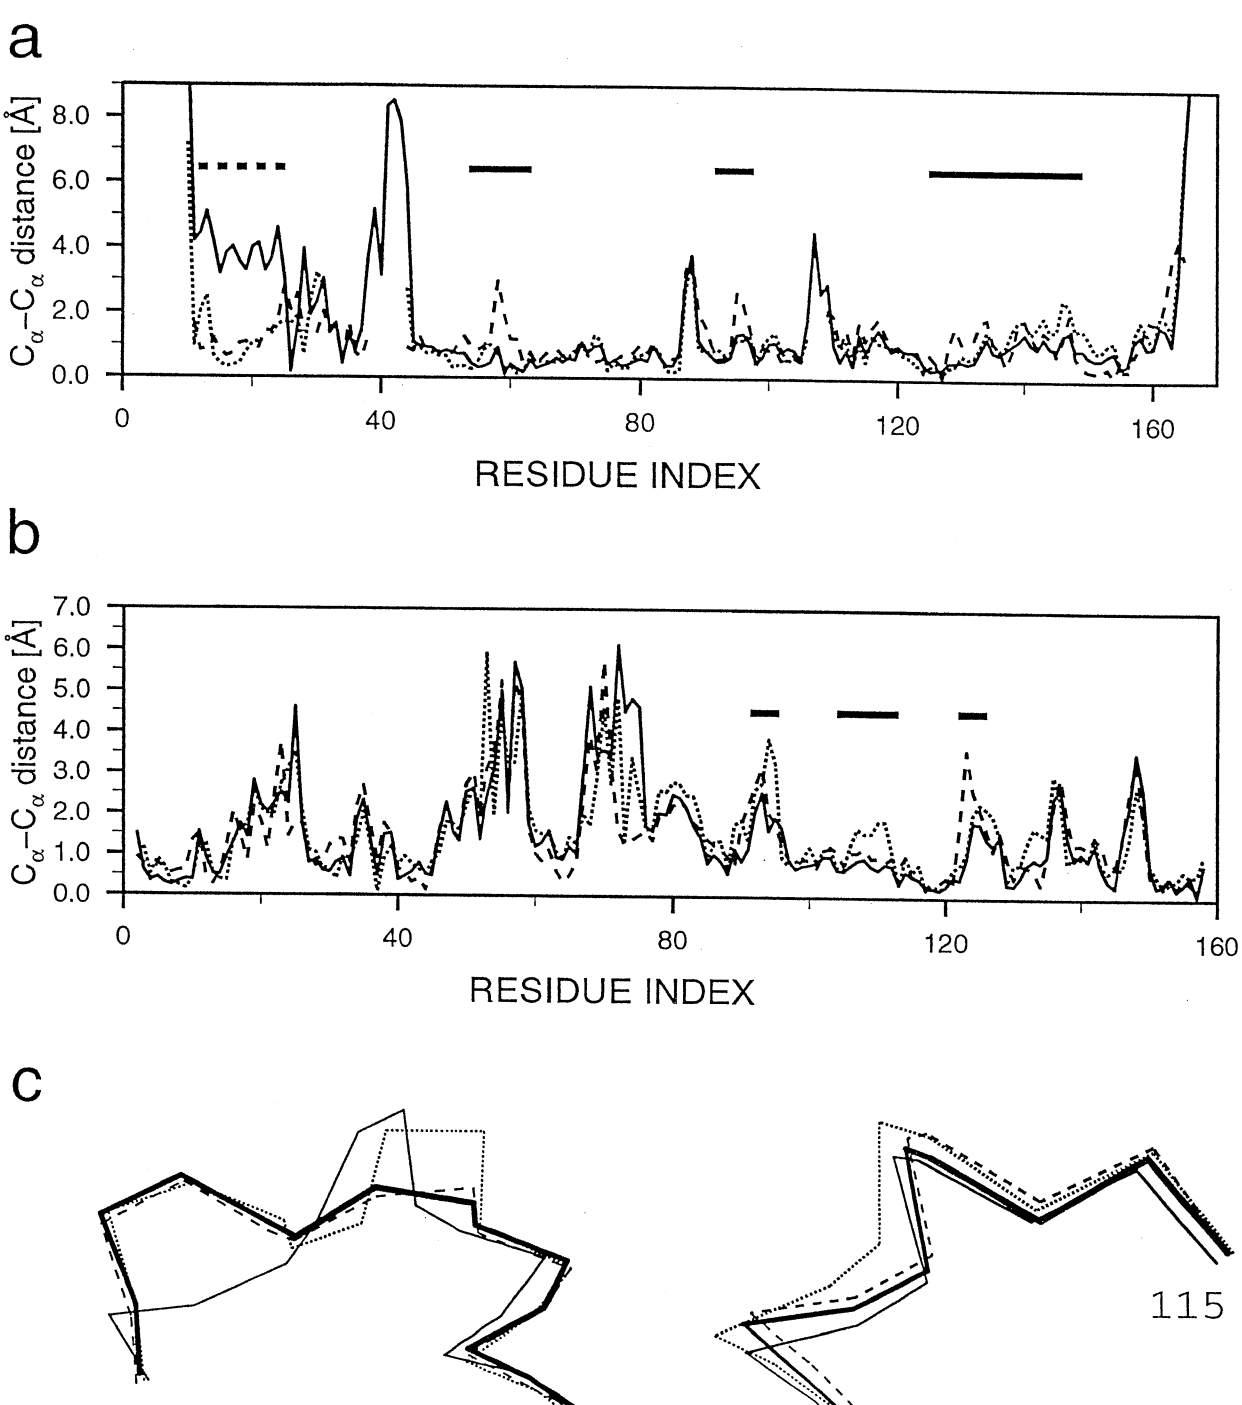

Fig. 1. Backbone errors in the $UBC9_{24}$ **(a)** and DFR_1 **(b,c)** models. The models and the templates are superposed as rigid bodies on the corresponding target structures using a cutoff of 3.5 Å for the equivalent Cα atoms. A curve in (a) and (b) shows the distances of the Cα atoms in the model and the templates from the equivalent atoms in the target. (a) $UBC9_{24}$ model - target, continuous line; template 2UCE–target, *dashed line*; template 1AAK–target, *dotted line*. (b) DFR_1 model–target, *continuous line*; template 4DFR-B–target, *dashed line*; template 3DFR–target, *dotted line*. The horizontal continuous lines above the curves indicate the correctly aligned segments of at least three residues where the best template was chosen for the model. The horizontal dashed line at the N terminus in (a) indicates the 11 misaligned residues of $UBC9_{24}$. (c) Superposition of residues 45–60 and 105–115 of the DFR_1 model with the corresponding regions in the templates and the X-ray structure. The model, *thick continuous line*; X-ray structure, *thin continuous line;* template 3DFR, *dotted line*; template 4DFR-B, *dashed line*.

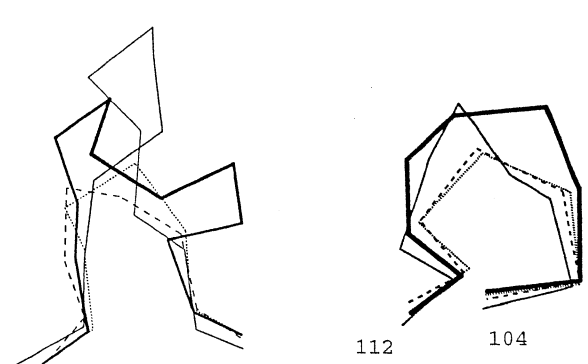


```
1aak APQDN-----NIM        1aak ILQN--QWS
2uce GPVGD-----DLY        2uce ILKD--QWS
UBC9_24 VPTKNPDGTMNLM     UBC9_24 ILEEDKDWR
```

Fig. 2. Errors in the two $UBC9_{24}$ loop models. The loops corresponding to the two insertions in the $UBC9_{24}$ model (*continuous thick line*) are shown superposed with the target X-ray structure (*continuous thin line*), the templates 1AAK (*dotted line*), and 2UCE (*dashed line*). The numbers indicate the beginning and ending residues of each segment in $UBC9_{24}$. The corresponding regions of the modeling alignment are shown below each set of the structures.

Errors in Loops

There were only two insertions in the three models, both of them in $UBC9_{24}$ (Fig. 2). The longest insertion was only five residues long (residues 40–44), and the second insertion was two residues long (residues 108–109). When the whole model was superposed on the X-ray structure, the RMSD between the backbones for the five-residue loop was 6.7 Å; when the backbones of only the two loops were superposed locally, the RMSD was 1.7 Å. Thus, both the orientation and conformation of the predicted loop were incorrect. The large difference between the two numbers shows that the positioning of the loop relative to the rest of the protein can be a very important contributor to the total error even in the case of relatively short loops. The alignment in the neighborhood of the loop was correct, except perhaps for the alignment of residue 39, which probably should not have been aligned with any residue in the templates (Fig. 2). The RMSD for the backbones of the three residues preceding (37–39) and the three residues following the loop (45–47) was 2.3 Å and 1.5 Å for the global and local superposition, respectively. The average backbone isotropic temperature factors for the five- and two-residue insertions were 24.4 $Å^2$ and 22.2 $Å^2$, respectively, compared to the slightly

a

```
          1aak  MSTPARKRLMRDF-KRLQQDPPA
          2uce  --MSSSKRIAKEL-SDLERDPPT
modeling  UBC9_24  MSGIALSRLAQER-KAWRKDHPF
3D        UBC9_24  MS GIALSRLAQERKAWRKDHPF
```

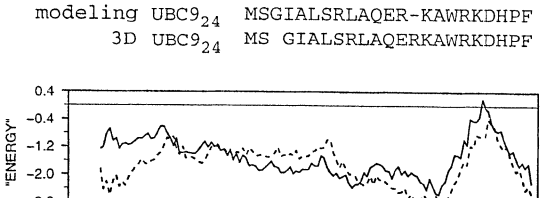

b

```
          3dfr  TKVSSRTVEDT---NPALTHTYEVWQKKA
         4dfrB  ESVFSEFHDADAQNS--HSYCFKILERR
DFR_1 model1  ELDAETDHEG---------FTLQEWVRSASSR
DFR_1 model2  ELDAETDHEG-------FTLQEWVRSASSR
```

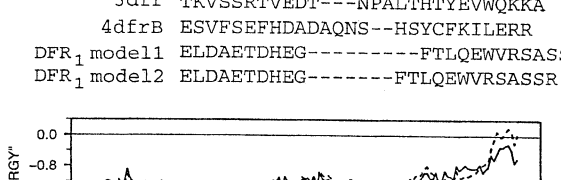

c

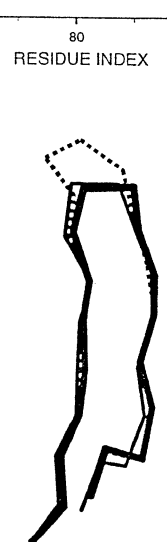

Fig. 3. Alignment problems and solutions. **(a)** Alignment of the N-terminal region of $UBC9_{24}$. The alignment used for model building (modeling) and the correct alignment derived from the superposition of the experimental structures of the templates and the target (3D) are shown. The Prosall energy profiles for the model (*continuous line*) and the target X-ray structure (*dashed line*) are shown below the alignment. Note the lower energy of the X-ray structure in the misaligned region. **(b)** The correct and alternative alignments for the C-terminal region of DFR_1. The Prosall energy profiles for the corresponding 3D models are shown below the alignment. The model based on the correct DFR_1 alignment, *continuous line*; the model based on the alternative alignment, *dashed line*. Note the positive energy for the alternative model in the C-terminal region. **(c)** Superposition of the C-terminal region of the correct (continuous line) and alternative model of DFR_1 (*dashed line*) with the X-ray structure (*thin line*).

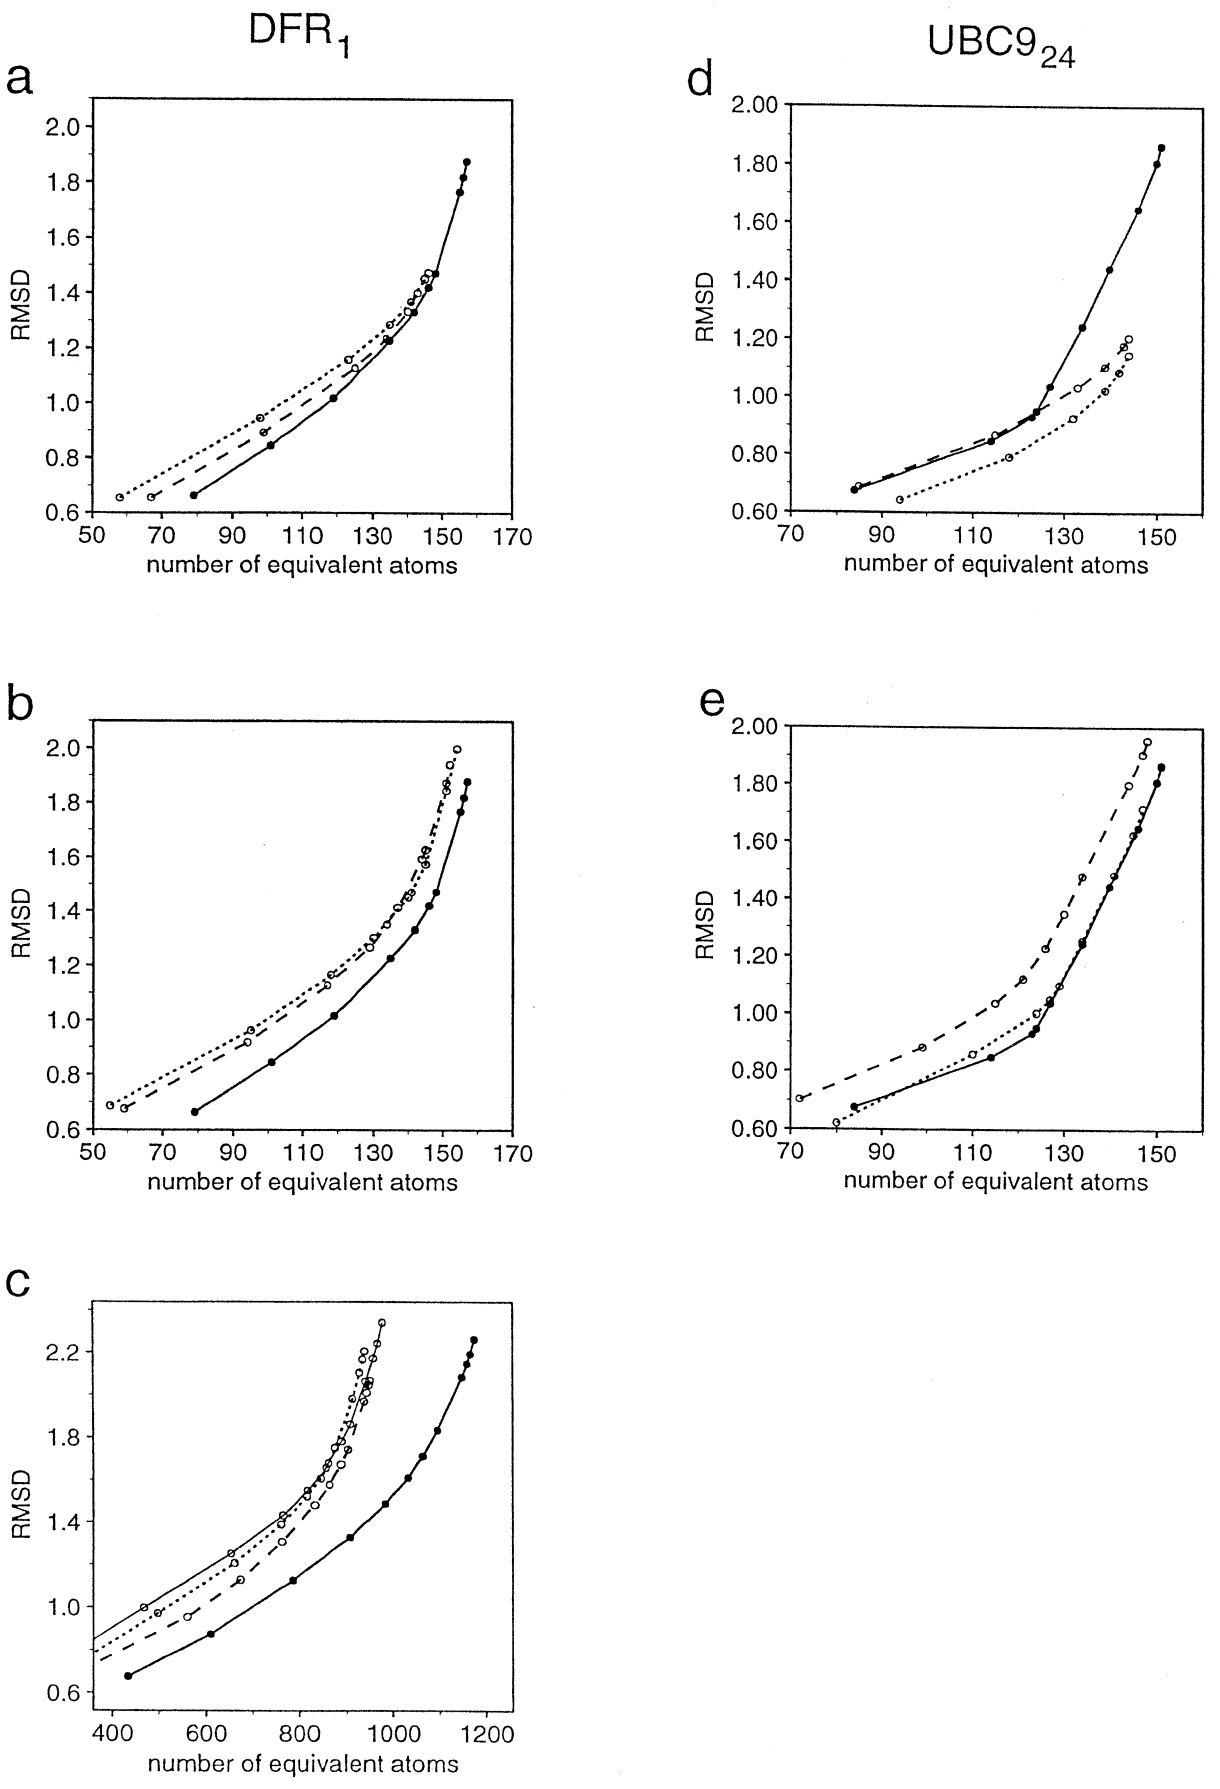

Figure 4 (legend on following page).

lower average of 16.4 Å^2 for the backbone of the whole protein. Thus, if the loops are not in contact with other protein molecules in the crystal, it is likely that the differences between the insertions in the crystal structure and the model reflect errors in the model.

Distortions or Shifts in Incorrectly Aligned Regions: The Cycle of Alignment, Model Building, and Model Evaluation

In the three models, there was only one secondary structure segment that was misaligned, the N-terminal helix of UBC9$_{24}$. In addition, there were three, zero, and one gaps in the modeling alignments for DFR$_1$, UBC9$_{24}$, and PTE2A$_3$, respectively, where one or a few residues were misaligned.

In the UBC9$_{24}$ model, the N-terminal segment of 11 residues was misaligned by one position, which resulted in large errors in the model (Fig. 1a). This misalignment was unexpected because the correct alignment corresponded to a significantly lower sequence similarity between the target and the template (Fig. 3a). For example, the number of matches between hydrophobic residues is decreased and the number of matches between hydrophobic and polar residues is increased when the incorrect alignment is corrected. The misalignment was not detected by the ProsaII profile of the model (Fig. 3a). However, the comparison of the profiles for the X-ray structure and the model shows that the X-ray structure has a lower ProsaII score in that region (Fig. 3a). This suggests that the search for the alignment with the lowest ProsaII profiles of the implied model could conceivably result in the correct alignment and thus a significantly better model in this case.

Another interesting observation is that the overall sequence identity between the target sequence and the more similar of the two templates dropped from 39% to 35% for the correct alignment. This makes the point that optimizing only sequence similarity is not always best in comparative modeling.

In the DFR$_1$ model, it was obviously difficult to align the last 13 residues, corresponding to the last strand of the last β hairpin (Fig. 3). Two plausible alternative alignments were generated manually by taking into account local sequence similarity, secondary structure predictions for DFR$_1$,[26–28] and the structures of the template proteins. The alignments were evaluated by comparing the ProsaII profiles of the models based on those alignments (Fig. 3b). One of the models had a positive profile, and the other one had a negative profile at the C terminus. A comparison of the two models with the X-ray structure showed that the model with the negative profile was indeed correctly aligned with the template (Fig. 3c).

As illustrated above, alignment errors are a major source of large errors in comparative models. We attempted to overcome this limitation by iterating through several cycles of careful manual template selection and alignment, followed by automated model building and model evaluation. This process was guided by a reduction in the errors predicted by a number of model evaluation techniques, most importantly the "energy" profiles calculated by the ProsaII program and a program of Melo and Feytmans.[15] Despite our limited experience, we believe that evaluation of an alignment at the level of the implied model is likely to overcome a significant fraction of initial alignment errors, especially when better potential functions for model evaluation become available and when the iterative procedure is automated so that a larger number of alternative alignments can be explored.[29]

Overall Accuracy of the Models: Relative Overall Similarity of a Model and the Templates to the Target X-ray Structure

We now wish to answer the question of whether the predicted structures are a better model of the experimental structures than the templates used in the calculation of the models. In other words, how much closer is a comparative model of the target sequence to the target X-ray structure than the closest template structure?

Although a single RMSD value is useful for measuring a difference between two relatively similar structures, RMSD depends on the number of equivalent atom pairs that are compared, which in turn depends on the maximal allowed distance between two equivalent atoms. This makes a single RMSD value inconvenient for comparing differences between pairs of different proteins. One solution to this problem is to define a *similarity curve* for a pairwise structure-structure comparison by plotting RMSD as a function of the number of equivalent atoms. The similarity curve is obtained by calculating RMSD at different cutoff values for equivalencing intermolecular pairs of Cα atoms and plotting the resulting RMSD values against the number of equivalent positions obtained at each cutoff. Two similarity curves, instead of two single RMSD numbers, can then be inspected for a comparison of two protein–protein matches.

The similarity curves for the three pairwise comparisons of the DFR$_1$ model and the two templates with the target structure are plotted in Figure 4a. The curves show that over a large range of the

Fig. 4. Similarity curves for the DFR$_1$ **(a, b, c)** and UBC9$_{24}$ **(d, e)** models and templates. See the Methods section for the definition of the similarity curves. (a and d) The optimal superposition of the templates and the X-ray structure was used to define the equivalent residues. (b, c, and e) The modeling alignment was used to define the equivalences between the templates and the target. (a), (b), (d), and (e), only the Cα atoms are used to calculate RMSD. (c) All atoms are used to calculate RMSD. Model–target, *thick continuous line*; template 4DFR-B–target and template 2UCE–target, *dashed line*; template–target and template 1AAK–target, *dotted line*; template 8DFR–target, *thin continuous line*.

number of equivalent atoms, the model is slightly closer to the experimental structure (lower RMSD value) than either of the two templates. In other words, at a fixed number of atoms compared, the model atoms have a lower RMSD from the X-ray structure than the template atoms; conversely, at a fixed RMSD, the model has more atoms equivalent to the X-ray structure than either of the templates. However, the differences are small, <10% over most of the similarity range.

Errors in the positioning of three gaps in the DFR_1 modeling alignment contributed to the similarity curve for the model-target comparison, but not to the template-target similarity curves in Figure 4a, which were obtained from the superposition of the crystallographic structures. In order to eliminate the contribution of the alignment errors and evaluate the model building procedure on its own, the similarity curves were recalculated using the modeling alignment for comparison of the templates with the target structure (Fig. 4b). Since the template-target comparisons now include the alignment errors, the templates are less similar to the target X-ray structure than in Figure 4a. However, the difference in how representative of the target structure are the model and the templates is still small, on the order of 10% of RMSD.

When side chain atoms were included in the calculation of the similarity curves, the DFR_1 model became an even better representation of the target structure relative to the templates (Fig. 4c). For example, the model had approximately 95% of its atoms superposed with an RMSD from the target structure of 2 Å, while the closest template only had 78% of the atoms at that level of similarity (Fig. 4c). This was expected because the templates do not share all the side chain atoms with the target structure while the model does.

In contrast to DFR_1, the $UBC9_{24}$ model is worse than the best template because of the alignment errors, primarily the shift for one position of the N-terminal 11 residues (Fig. 4d). The $PTE2A_3$ model is as close to the target structure as the best template (data not shown).

All comparative modeling methods start with an alignment of the target sequence with the template structures, followed by model building that is decoupled from the alignment procedure. Therefore, when evaluating comparative modeling methods, it is important for method developers to distinguish between errors due to misalignments and errors due to the model building procedure. This distinction is also important for the method users because the modeling alignment, not the correct alignment, would be used to extract information from the template structure in the absence of any model building. When the modeling alignment is used to compare both the model and the templates with the target structure, all three models are a better representation of the experimental structure than the templates used in their derivation (Fig. 4b,e; data not shown for $PTE2A_3$). This is especially true when the side chain as well as backbone atoms are compared (e.g., Fig. 4c). These comparisons suggest that it is better to use a comparative model of the target than homologous structures, unless only coarse predictions are made.

CONCLUSIONS

The modest improvement in our models relative to CASP1 probably originates from the careful manual selection of the templates and editing of the alignment, as well as from the iterative re alignment and model building. This suggests directions for future development of the algorithms that will, it is hoped, result in larger increases in the model accuracy.[8,29–32]

ACKNOWLEDGMENTS

We thank crystallographers U. Pieper, O. Herzberg, K. Huang, H. Tong and T. Sixma for providing the structures before their release to the PDB, and F. Melo and E. Feytmans for evaluating our models with their evaluation program. R.S. is a Howard Hughes Medical Institute predoctoral fellow. A.S. is a Sinsheimer Scholar.

REFERENCES

1. Delbaere, L.T.J., Brayer, G.D., James, M.N.G. Comparison of the predicted model of ff-lytic protease with the x-ray structure. Nature 279:165–168, 1979.
2. Sali, A., Potterton, L., Yuan, F., van Vlijmen, H., Karplus, M. Evaluation of comparative protein modeling by MODELLER. Proteins 23:318–326, 1995.
3. Rost, B. Protein structures sustain evolutionary drift. Folding Design 2:S19–S24, 1997.
4. Šali, A. and Overington, J. Derivation of rules for comparative protein modeling from a database of protein structure alignments. Protein Sci. 3:1582–1596, 1994.
5. Šali, A., Blundell, T.L. Comparative protein modelling by satisfaction of spatial restraints. J. Mol. Biol. 234:779–815, 1993.
6. Abola, E.E., Bernstein, F.C., Bryant, S.H., Koetzle, T., Weng, J. Protein data bank. In: "Crystallographic Databases: Information, Content, Software Systems, Scientific Applications." Allen, F.H., Bergerhoff, G., Sievers, R., eds. Bonn: Data Commission of the International Union of Crystallography, 1987:107–132.
7. Johnson, M.S., Srinivasan, N., Sowdhamini, R., Blundell, T.L. Knowledge-based protein modelling. CRC Crit. Rev. Biochem. Mol. Biol. 29:1–68, 1994.
8. Sánchez, R., Šali, A. Advances in comparative protein-structure modeling. Curr. Opin. Str. Biol. 7:206–214, 1997.
9. Sánchez, R., Šali, A. Modeller, A Protein Structure Modeling Program, Release 3. URL http://guitar.rockefeller.edu/, 1995.
10. Needleman, S.B., Wunsch, C.D. A general method applicable to the search for similarities in the amino acid sequence of two proteins. J. Mol. Biol. 48:443–453, 1970.
11. Brooks, B.R., Bruccoleri, R.E., Olafson, B.D., States, D.J., Swaminathan, S., Karplus, M. CHARMM: A program for macromolecular energy minimization and dynamics calculations. J. Comp. Chem. 4:187–217, 1983.
12. Laskowski, R.A., McArthur, M.W., Moss, D.S., Thornton, J.M. PROCHECK: A program to check the stereochemical quality of protein structures. J. Appl. Crystallogr. 26:283–291, 1993.

13. Hooft, R., Vriend, G., Sander, C., Abola, E. Errors in protein structures. Nature 381:272, 1996.
14. Sippl, M.J. Recognition of errors in three-dimensional structures of proteins. Proteins 17:355–362, 1993.
15. Melo, F., Feytmans, E. Novel knowledge-based mean force potential at atomic level. J. Mol. Biol. 267:207–222, 1997.
16. Koehl, P., Delarue, M. Application of a self-consistent mean field theory to predict protein side-chains conformation and estimate their conformational entropy. J. Mol. Biol. 239: 249–275, 1994.
17. Lee, C. Predicting protein mutant energetics by self consistent ensemble optimisation. J. Mol. Biol. 236:918–939, 1994.
18. Maeyer, M.D., Desmet, J., Lasters, I. All in one: A highly detailed rotamer library improves both accuracy and speed in the modelling of side chains by dead-end elimination. Folding Design 2:53–66, 1997.
19. Chung, S.Y., Subbiah, S. A structural explanation for the twilight zone of protein sequence homology. Structure 4:1123–1127, 1996.
20. Sutcliffe, M.J., Hayes, F.R.F., Blundell, T.L. Knowledge based modeling of homologous proteins, Part II: Rules for the conformation of substituted side-chains. Protein Eng. 1:385–392, 1987.
21. Summers, N.L., Karplus, M. Construction of side-chains in homology modelling: Application to the C-terminal lobe of rhizopuspepsin. J. Mol. Biol. 210:785–811, 1989.
22. Bower, M.J., Cohen, F.E., Dunbrack, R.L. Prediction of protein side-chain rotamers from a backbone-dependent rotamer library: A new homology modeling tool. J. Mol. Biol. 267:1268–1282, 1997.
23. Miyazawa, S., Jernigan, R.L. Estimation of effective interresidue contact energies from protein crystal structures: Quasi-chemical approximation. Macromolecules 18:534–552, 1985.
24. Sippl, M.J. Calculation of conformational ensembles from potentials of mean force: An approach to the knowledge-based prediction of local structures in globular proteins. J. Mol. Biol. 213:859–883, 1990.
25. DeBolt, S.E., Skolnick, J. Evaluation of atomic level mean force potentials via inverse folding and inverse refinement of protein structures: Atomic burial position and pairwise nonbonded interactions. Protein Eng. 9:937–955, 1996.
26. Rost, B., Sander, C. Prediction of protein structure at better than 70% accuracy. J. Mol. Biol. 232:584–599, 1993.
27. Biou, V., Gibrat, J.-F., Levin, J., Garnier, J. Secondary structure prediction: Combination of three different methods. Protein Eng. 2:185–191, 1988.
28. Chandonia, J.M., Karplus, M. The importance of larger data sets for protein secondary structure prediction with neural networks. Protein Sci. 5:768–774, 1996.
29. Guenther, B., Onrust, R., Šali, A., O'Donnell, M., Kuriyan, J. Crystal structure of the δ′ subunit of the clamp-loader complex of *E. coli* DNA polymerase III. Cell 91:335–345, 1997.
30. Šali, A. Modelling mutations and homologous proteins. Curr. Opin. Biotech. 6:437–451, 1995.
31. Šali, A. Protein modeling by satisfaction of spatial restraints. Mol. Med. Today 1:270–277, 1995.
32. Sánchez, R., Šali, A. Comparative protein modeling as an optimization problem. J. Mol. Struct. (Theochem) 398:489–496, 1997.