Acta Cryst. (1994). D50, 24–36

Cross-Validation Tests of Time-Averaged Molecular Dynamics Refinements for Determination of Protein Structures by X-ray Crystallography

BY JAMES B. CLARAGE* AND GEORGE N. PHILLIPS JR

Keck Center for Computational Biology and Department of Biochemistry and Cell Biology, Rice University, Houston, TX 77251-1892, USA

(*Received* 5 *March* 1993; *accepted* 8 *September* 1993)

Abstract

Time-averaged structure-factor restraints have been used in two molecular dynamics refinement schemes to define ensembles of conformations for myoglobin that fit the experimentally measured Bragg scattering from *P*6 crystals. The geometries of the structures have been maintained to the same currently acceptable limits in all cases. Free *R* value analysis was used to assess the validity of the two approaches. In the first scheme, where atoms have no *B* values, the decrease in *R* value was found to be spurious as judged by a concomitant increase in the free *R* value. The other scheme, however, which retains individual *B* values, was found to yield both low *R* values and low free *R* values; thus, here the additional variables introduced by modeling the protein in terms of an evolving ensemble of states do not overfit the data. For comparison, refinements were also carried out on the system using several other techniques for isotropic and anisotropic crystallographic refinement. The time-averaged refinements with *B* values compare quite favorably with the standard methods, but yield additional information about substates of the system. Hence, correctly applied time-averaged refinements can yield accurate models for protein molecules; moreover, by essentially relaxing the harmonic approximation from the refinement process, these refinements allow a more detailed description of the motions of complex molecules, such as proteins, to be determined from X-ray crystallographic data.

Introduction

For all its triumphs, the lingering embarrassment of protein crystallography is its failure to yield a structural model for a protein molecule in the crystalline state which agrees with the X-ray scattering data to anywhere near the noise level of the measurements. Agreement is typically judged from the crystallographic *R* value:

$$R = \frac{\sum_{\mathbf{h}} ||F_o(\mathbf{h})| - |F_c(\mathbf{h})||}{\sum_{\mathbf{h}} |F_o(\mathbf{h})|}, \quad (1)$$

which is the difference between observed structure-factor amplitudes $|F_o|$ and structure-factor amplitudes calculated from the model $|F_c|$, averaged over all measured reflections with Miller index **h**. The inherent error in X-ray intensity measurements can be estimated from the variance between intensities which should be equal because of symmetry considerations. The *R* value for symmetry-related reflections is typically 5% or better. But the *R* value for a refined atomic model of the protein can rarely be made smaller than 10–15%; errors less than 20% for the high-resolution reflections are generally considered a success. The situation is even more acute for low-resolution data, which is ironic since counting statistics guarantee that these stronger reflections can be measured more accurately. Indeed, *R* values are so large for lower order reflections that these data are customarily ignored in protein crystallographic refinement.

This predicament is not inherent in the scattering theory of X-ray crystallography. Crystal structures of small molecules can routinely be determined to within the noise level in the data. So what makes protein crystals different from monatomic or small-molecule crystals? First, owing to their large size and weakly coupled atomic interactions, protein molecules can assume many different, though closely related, conformational substates at room temperature, a fact demonstrated by spectroscopy (Hong *et al.*, 1990), X-ray diffuse scattering (Boylan & Phillips, 1986; Caspar, Clarage, Salunke & Clarage, 1988; Clarage, Clarage & Caspar, 1992; Chacko & Phillips, 1992) and molecular dynamics calculations (Karplus & McCammon, 1983). Secondly, protein crystals are usually only about half protein; the other half of the crystal volume is occupied by solvent molecules. It is these two properties of protein crystals which make them inherently difficult to model.

At low resolution the X-ray intensities are strongly dependent upon the bulk properties of the crystal, *viz.* the contrast between protein and solvent. Since

* Author for correspondence: Department of Biochemistry and Cell Biology, PO Box 1892, Rice University, Houston, TX 77251, USA.

fluctuations in the protein density are local and of relatively small amplitude at the length scales probed by low-resolution data, it must be our lack of knowledge about the solvent density which is responsible for the poor agreement between models and data. Simply modelling the solvent as a constant, equal to its average density, goes a long way over ignoring the solvent's presence (Blake, Pulford & Artymiuk, 1983). When the contrast between the solvent and protein is significant (*e.g.* low salt) iterative refinement of the solvent density (Badger & Caspar, 1991) is feasible and yields dramatic improvements in the agreement between calculated and observed low-order reflections.

At high resolution, only the well ordered regions of the crystal can contribute to the diffraction, namely the atoms in the protein molecule and those in the hydration shell which are not as free to diffuse. Thus, the large R values for high-resolution data are due to inadequate modelling of the electron density of a protein molecule and ordered water. It is standard to model the electron density for a protein molecule by refining a single site for each atom in the structure. As with small-molecule crystallography, such a model overestimates the high-order scattering and therefore, following Debye, each atom's electron density is smeared out with a B value (Debye, 1914) to disorder the density. The B value assumes that all atoms vibrate harmonically about a single equilibrium position, an assumption which is presumably true for most small molecules since it leads to reasonable R values, but which evidently breaks down for macromolecules. Obviously, a single conformation with simple deviations is not an accurate model for even the more ordered regions of a protein crystal.

Schemes to account for the detailed structure of the solvent [*e.g.* refining pseudo-atoms of water density (Schoenborn, 1988)] or protein [*e.g.* anisotropic temperature factors, twin conformations (Kuriyan *et al.*, 1991), time-averaged refinement (Gros, van Gunsteren & Hol, 1990)] are somewhat suspect because they introduce extra variables into the least-squares refinement of the model and, therefore, run the risk of over-fitting the data. For this reason, such attempts have never gained wide acceptance in the structural biology community. In this paper we limit our scope to the problem of accurately modeling the high-resolution data from protein crystals. In particular, we apply Brünger's crystallographic cross-validation methods to determine the reliability of two strategies for modeling the electron density from *P*6 myoglobin crystals with an ensemble of conformations. The two strategies differed only in whether initial B values were assigned to the atoms in the protein. In both cases, conformations were generated from molecular dynamics with a time-averaged crystallographic restraint similar to that introduced by Gros, van Gunsteren & Hol (1990) for bovine pancreatic phospholipase A_2.

Methods

Experimental data

The protein studied was a myoglobin mutant expressed in *E coli* from a synthetic gene, where Val68 has been mutated to Ile. This particular myoglobin is part of ongoing studies on the effect of site-directed mutations of residues in the distal heme pocket on the binding of ligands to the protein. See Quillin, Arduini, Olson & Phillips (1993) for details on synthesis, crystallization, data collection and initial refinement. The protein crystallized in space group *P*6 with one molecule in the asymmetric unit, and cell constants $a = b = 91.20$, $c = 45.87$ Å, $\alpha = \beta = 90$, $\gamma = 120°$. Crystals were in the ferric met oxidation state, which has a water molecule bound to the iron.

X-ray diffraction data were collected on a Siemens X-1000 multiwire area detector with a rotating-anode source, and processed with the package *XDS* (Kabsch, 1988). The total R_{sym} based on intensities was 5.6%, monotonically increasing from 3.2% at 6 Å resolution to 23.0% at 1.8 Å. (The approximate R_{sym} statistics based on amplitudes would then be smaller by a factor of two.)

Using phases from a model of the wild-type protein, the structure of the Ile68 mutant was determined by first modeling the mutated side chain from initial difference maps, then refining the entire structure using the package *PROLSQ* (Hendrickson & Konnert, 1980). Refinement of positions, occupancy and B values for 1495 protein atoms and 185 water molecules resulted in a final R value of 16.4%, for data between 5.0–1.8 Å.

Theory of time-averaged refinement

Before discussing the refinement used in this study, we present a brief overview of standard crystallographic refinement to put this work in context.

Since X-ray experiments can only measure the diffracted amplitudes and not the phases, it is impossible to directly invert the data to obtain the average electron density of the protein molecule. Consequently, standard crystallographic structure determination (such as that used to solve the structure of the Ile68 myoglobin) is handled as an optimization problem. A target function expressing the agreement between the model and data is optimized with respect to a set of parameters describing the protein structure. The target function used most often to minimize the difference between observed structure

factors and those calculated from the model is

$$\Phi = \sum_{\mathbf{h}} [|F_o(\mathbf{h})| - K|F_c(\mathbf{h})|]^2, \quad (2)$$

where a constant, K, is introduced to place the two sets of structure factors on the same scale. It is customary to write the calculated structure factor in this expression as

$$F_c(\mathbf{h}) = \sum_m Q_m f_m(\mathbf{h}) \exp(-B_m|\mathbf{h}|^2/4) \exp(i2\pi\mathbf{h}\cdot\mathbf{r}_m) \quad (3)$$

where the sum extends over all atoms in the unit cell. In this expression, $f(\mathbf{h})$ is the atomic form factor, Q is the occupancy (usually 1.0), $\mathbf{r}$ the average position, and B is a temperature factor. The temperature factor is an attempt to account for the fact that the atoms do not always sit at their average positions $\mathbf{r}$ but are smeared isotropically with a Gaussian distribution. In this scheme, therefore, each atom is characterized by four or five parameters. The derivative, or gradient, of the residual can be calculated with respect to these parameters so that minimization algorithms can be used to find the best values of Q, B, $\mathbf{r}$ for each atom. To ensure a solution which is well determined, extra 'observations' are incorporated into the optimization by introducing stereochemical restraint terms into the residual (*e.g. PROLSQ*, Hendrickson & Konnert, 1980). Hence the final protein structure ends up fitting not only the observed X-ray data but also known restrictions on the geometry of amino acids.

Because there are so many degrees of freedom in the protein structure problem, gradient-following optimizations can become trapped in local minima. To overcome this problem simulated-annealing methods have been applied to the problem (Brünger, Kuriyan & Karplus, 1987). The system is transformed into a dynamical one by constructing a pseudo energy function related to the crystallographic residual. In this way, molecular dynamics at artificially high temperatures can be used to sample the phase space better and locate a solution closer to the global minimum. The simulated-annealing energy function used is a hybrid of the form

$$\Phi = \Phi_{\text{phys}} + \Phi_{\text{X-ray}}. \quad (4)$$

$\Phi_{\text{X-ray}}$ is set proportional to the standard crystallographic target function; that is

$$\Phi_{\text{X-ray}} = \tfrac{1}{2} W_x \sum_{\mathbf{h}} [|F_o(\mathbf{h})| - K|F_c^t(\mathbf{h})|]^2 \quad (5)$$

$$= \tfrac{1}{2} W_x \sum_{\mathbf{h}} E^2(\mathbf{h}), \quad (6)$$

where F_c^t is the instantaneous structure factor calculated from the atomic coordinates and B values at the present time ($t = t'$) in the simulation. Thus, this energy term acts like a spring [in the coordinate $E(\mathbf{h}) = |F_o(\mathbf{h})| - K|F_c^t(\mathbf{h})|$] to force calculated structure factors towards the observed equilibrium values. W_x is the weight given to the X-ray term relative to the other term, Φ_{phys}, the physical energy of atomic interaction as a result of covalent bonds, dihedral and torsion angles, and Coulombic and van der Waals forces. This physical term enforces good geometry in the structure.

Time-averaged refinement (Gros, van Gunsteren & Hol, 1990; Gros, 1990) resembles simulated-annealing refinement in that the optimization problem is cast as one in dynamics. However, there is no minimization stage to single out a unique final structure. Rather, the evolving dynamics of the system is used to generate an entire set of conformations consistent with the experimental data. This is accomplished by defining the X-ray energy, not in terms of structure factors calculated from the present conformation of the molecule, but in terms of structure factors averaged over the past history of conformations. Explicitly,

$$\Phi_{\text{X-ray}} = \tfrac{1}{2} W_x \sum_{\mathbf{h}} [|F_o(\mathbf{h})| - K|\langle F_c(\mathbf{h})\rangle_{t'}|]^2 \quad (7)$$

$$= \tfrac{1}{2} W_x \sum_{\mathbf{h}} E^2(\mathbf{h}). \quad (8)$$

The average structure factor is treated as a weighted average, where contributions from past conformations receive exponentially smaller weight. Otherwise, with an unweighted average, as the number of steps in the dynamics increased the contribution of the present conformation to the running average would decrease until the simulation effectively became unrestrained.* At time $t = t'$ this weighted average is recursively defined as

$$\langle F_c(\mathbf{h})\rangle_{t'} = \exp(-\Delta t/\tau_x)\langle F_c(\mathbf{h})\rangle_{t'-\Delta t} + [(1 - \exp(-\Delta t/\tau_x)]F_c^t(\mathbf{h}) \quad (9)$$

where Δt is the step size in the simulation, and the relaxation time τ_x defines the window size over which individual instantaneous structure factors, F_c^t, are averaged. The $[1 - \exp(-\Delta t/\tau_x)]$ factor simply normalizes the average.

The time-averaged crystallographic force on atomic coordinate x_m is

$$G(x_m) = -\partial\Phi_{\text{X-ray}}/\partial x_m \quad (10)$$

$$= -\sum_{\mathbf{h}} [\partial\Phi_{\text{X-ray}}/\partial\langle F_c(\mathbf{h})\rangle_{t'}][\partial\langle F_c(\mathbf{h})\rangle_{t'}/\partial x_m]. \quad (11)$$

* The political analogy is that as the number of people voting in an election increases, the less any one vote contributes to the outcome.

The first term in the chain-rule expansion of the force is not rigorously defined since the energy is not an analytic function of the complex variables $\langle F_c(\mathbf{h})\rangle_{t'}$. Nevertheless, the following analysis will yield equivalent results to a more formal treatment (following Lunin & Urzhumtsev, 1985) wherein crystallographic expressions are treated as functions of two real variables (F_c^{real}, F_c^{imag}) instead of a single complex one. Differentiating the energy with respect to the time-averaged structure factor gives

$$\frac{\partial \Phi_{\text{X-ray}}}{\partial \langle F_c(\mathbf{h})\rangle_{t'}} = W_x K E(\mathbf{h}) \frac{\partial |\langle F_c(\mathbf{h})\rangle_{t'}|}{\partial \langle F_c(\mathbf{h})\rangle_{t'}}. \quad (12)$$

Using the identity $\partial|F|^2/\partial F = 2|F|(\partial|F|/\partial F)$ allows us to write

$$\frac{\partial \Phi_{\text{X-ray}}}{\partial \langle F_c(\mathbf{h})\rangle_{t'}} = W_x K E(\mathbf{h}) \frac{\langle F_c(\mathbf{h})\rangle_{t'}}{|\langle F_c(\mathbf{h})\rangle_{t'}|} \quad (13)$$

$$= W_x K E(\mathbf{h}) \exp[i\varphi(\mathbf{h})] \quad (14)$$

where $\varphi(\mathbf{h})$ is the phase of the time-average structure factor.

The second half of the force $G(x_m)$ is straightforward since only the instantaneous part of $\langle F_c(\mathbf{h})\rangle_{t'}$ is a function of current atomic coordinate x_m. Indeed,

$$\frac{\partial \langle F_c(\mathbf{h})\rangle_{t'}}{\partial x_m} = [1 - \exp(-\Delta t/\tau_x)] \frac{\partial F_c^{t'}(\mathbf{h})}{\partial x_m} \quad (15)$$

$$= [1 - \exp(-\Delta t/\tau_x)](i2\pi h_x) f_m(\mathbf{h}) \times \exp(-B_m|\mathbf{h}|^2/4) \exp(i2\pi \mathbf{h}\cdot\mathbf{r}_m). \quad (16)$$

In this expression, h_x is the Miller index reciprocal to atomic coordinate x_m.

The total force is then,

$$G(x_m) = W_x K[1 - \exp(-\Delta t/\tau_x)] \sum_{\mathbf{h}} (i2\pi h_x) f_m(\mathbf{h}) \times \exp(-B_m|\mathbf{h}|^2/4) \exp(i2\pi \mathbf{h}\cdot\mathbf{r}_m) E(\mathbf{h}) \times \exp[i\varphi(\mathbf{h})]. \quad (17)$$

Except for the constant $[1 - \exp(-\Delta t/\tau_x)]$ (which can be absorbed by rescaling K) this formula is identical in form to the gradient of the standard crystallographic residual (Agarwal, 1980). The difference consists in using $\langle F_c(\mathbf{h})\rangle_{t'}$ instead of $F_c(\mathbf{h})$; that is, $E(\mathbf{h})$ must contain the difference between observed and time-averaged structure factors, and $\varphi(\mathbf{h})$ must be taken as the phase of $\langle F_c(\mathbf{h})\rangle_{t'}$. With this modification, then, existing algorithms can be used to compute the force as a result of a time-averaged structure-factor restraint.

We chose to work from the program *X-PLOR* (Brünger, 1987). This program calculates the crystallographic force by writing (17) as the Fourier transform of a product of two terms,

$$G(x_m) = W_x K \sum_{\mathbf{h}} \exp(i2\pi \mathbf{h}\cdot\mathbf{r}_m) \times \{f_m(\mathbf{h})(i2\pi h_x) \exp(-B_m|\mathbf{h}|^2/4)\} \times \{E(\mathbf{h}) \exp[i\varphi(\mathbf{h})]\}. \quad (18)$$

Application of the convolution theorem and the identity $\partial_{x_m} \exp(i2\pi \mathbf{h}\cdot\mathbf{r}_m) = (i2\pi h_x) \exp(i2\pi \mathbf{h}\cdot\mathbf{r}_m)$ yields

$$G(x_m) = W_x K \partial_{x_m} \rho(\mathbf{r}_m) * \text{FT}\{E(\mathbf{h}) \exp[i\varphi(\mathbf{h})]\}, \quad (19)$$

or the convolution of the differentiated electron density $\rho(\mathbf{r})$ with the Fourier transform (FT) of a difference-density map. This is analytically equivalent to another representation, $G(x_m) = \rho(\mathbf{r}_m) * \text{FT}\{i2\pi h_x E(\mathbf{h}) \exp[i\varphi(\mathbf{h})]\}$, obtained by grouping the $i2\pi h_x$ term with the second expression in the convolution (*e.g.* see Agarwal, 1980).

Modifications where made to the source code in *X-PLOR* so that $\langle F_c(\mathbf{h})\rangle_{t'}$ was used [instead of $F_c(\mathbf{h})$] in the terms $E(\mathbf{h})$ and $\exp[i\varphi(\mathbf{h})]$. Though computationally less efficient, it is also possible to execute this change solely within the *X-PLOR* command language, manipulating the two structure-factor arrays FCALC and FPART. With either route, however, the program's partial structure-factor array FPART is used to store the most recent time-averaged structure factors, so this array is no longer available for other refinement applications.

Application of time-averaged refinement

Two versions of time-averaged crystallographic refinement where studied. In the first, B values were not retained on the atoms. The second strategy kept B values from the standard refinement. Calculations were carried out on the single myoglobin molecule comprising the asymmetric unit of the $P6$ crystal. 175 crystallographically located waters were also included. Crystal symmetry and lattice contacts were made explicit by using closest symmetry-related neighbors in the atomic force calculations. In each of the two schemes, a series of refinements were executed using several different values for the weight W_x of the restraint. All data collected between 10.0–1.8 Å were utilized.

In scheme 1, our simulations followed essentially the same protocol introduced by Gros, van Gunsteren & Hol (1990) in their original time-averaged refinement of bovine pancreatic phospholipase A_2, where they were able to obtain the low R value of 9.8% Starting with the standard crystallographically refined coordinates and B values, molecular dynamics was run under the restraining force derived above. During the simulation, the B value of each atom was gradually reduced while the averaging window τ_x was gradually increased, until after 16 ps the B values vanished and $\tau_x = 4$ ps. For the next

10 ps, τ_x was increased to 16 ps where it was left for the remainder of the refinement. The adjustments in the B values and τ_x were handled within the *X-PLOR* command language.

Scheme 2 was identical to the protocol followed for scheme 1, except individual atomic B values were left on all atoms for the duration of the refinement.

In the results reported here, the TOPH19X and PARAM19X topology and parameter set from *X-PLOR*2.1 was used. This is essentially the original CHARMm TOPH19 and PARAM19 set (Brooks *et al.*, 1983), but with a few modifications to preserve geometry at the high temperatures used in X-ray crystallographic simulated annealing. Following this topology and parameter set, polar H atoms were treated explicitly, and all others implicitly by modifying the parameters of the hydrogen antecedent. No explicit hydrogen-bonded potential was modeled; rather, hydrogen-bond energy was incorporated into electrostatic and van der Waals terms. As in crystallographic refinement, we turned off the charges on the charged residues Lys, Arg, Glu and Asp. A cut-off of 7.5 Å was used for the non-bonded-pair list generation. The van der Waals interaction used a switching function (defined in *X-PLOR* 2.1) with cut-on and cut-off distances of 6.0 and 6.5 Å, respectively. Electrostatic forces, with dielectric constant ($\epsilon = 1$), were truncated using a shift function, with cut-off distance 6.5 Å. The TIPS3P model was used for water molecules. It is worth mentioning that simulations were also tried using *PROLSQ* parameters, which are set up to maintain stereochemical restraints and define no electrostatic interaction. The convergence properties of these simulations were virtually indistinguishable from those which were more physically realistic based on CHARMm, arguing that unlike unrestrained molecular dynamics, structure-factor restrained simulations are less sensitive to the specific details of the chemical energy.

For all simulations reported here, the total run times were between 40 and 80 ps, with an integration time step of 0.001 ps. The time-averaged structure factors were updated every 0.01 ps, at which points they and the atomic coordinates were saved on disk; thus for our purposes each picosecond of simulation corresponds to 100 structures. The non-conservative potential used in both schemes did lead to heating of the system; so the system was coupled to a 200 K heat bath using the Langevin dynamics scheme of Berendsen, Postma, van Gunsteren, DiNola & Haak (1984), with a friction coefficient of 84.0 ps^{-1}. Since the dynamics is restrained by the non-physical structure-factor 'force', it is probably preferable to think of the simulations not in terms of 'time' or 'picoseconds', but rather 'number of structures'. In such terms the simulations ran for between 4000 and 8000 structures.

Free R value

With enough parameters one can fit an elephant. So said Fermi. And so runs an uneasiness with using time-averaged crystallographic refinement. For in this strategy the data are modeled with many conformations, as opposed to one. Such ensemble modeling runs the real risk of overfitting the fixed and limited number of structure-factor observations. Without some independent consistency measure there is no way to know whether the time-averaged refinement strategy gives low R values because it succeeds in modeling the protein molecule or because it succeeds in mimicking the variation in intensities of a particular X-ray data set.

Brünger recently introduced the free R value (Brünger, 1992, 1993) as a statistical measure of the accuracy of crystal structures. To apply this test, a small subset, T, of the structure-factor data is set aside. The remaining set of data, A, is used in the refinement. Throughout the refinement, the model is compared with the test set of reflections by computing the free R value

$$R_T^{\text{free}} = \frac{\sum_{\mathbf{h}\epsilon T}||F_o(\mathbf{h})| - K|F_c(\mathbf{h})||}{\sum_{\mathbf{h}\epsilon T}|F_o(\mathbf{h})|}. \qquad (20)$$

Since the model has not been biased towards the test set of reflections T, this quantity contains an objective measure of the information content in the model. Indeed, Brunger showed that the free R value exhibits a high correlation with the accuracy of the phases in the atomic model for the protein, which makes it an ideal test since complete phase information is rarely available.

During a refinement, if drops in the standard R value are accompanied by rises in R_T^{free} then the parameters in the model at this stage of the refinement have spuriously lowered the R value by overfitting those diffraction data in the working data set A. For example, by adding 'nonsense' water molecules to an electron-density map or by relaxing stereochemical restraints the R value can be lowered in a refinement, but both these practices lead to an increase in the free R value.

Thus, to assess the meaningfulness of the two time-averaged refinement strategies, 10% of the structure-factor data were left aside as a test set, used to monitor the free R value during the course of the refinements. Only the remaining 90% of the data were included in the calculation of the time-averaged crystallographic restraint.

One important subtlety arises in the free R value analysis. Naively, the initial conditions for the time-averaged molecular dynamics refinement would be

taken to be the result of the standard crystallographic refinement of the protein. However, since this refinement was originally carried out using all of the data, the refined coordinates and B values of protein and solvent atoms are biased towards the test set of reflections. The value of R_T^{free} will consequently be artificially low. We therefore decided to re-refine the myoglobin and solvent structure using only the 90% working set in the calculations. Of course care also had to be taken to insure that the starting conditions for this revised single-site model were unbiased. To this end, the myoglobin molecule was 'scrambled' by running 0.1 ps of unrestrained molecular dynamics. Then all B values were set to a constant and all previous water sites were discarded. Starting with this state, repeated cycles of Powell coordinate minimization, occupancy and B value refinement, and water placement were executed using *X-PLOR*. Refinement was carried out against unweighted structure-factor amplitudes, using the PARAM19X parameter set from *X-PLOR*2.1. Simulated annealing was not used. This re-refinement was terminated when R_T^{free}, not R, stopped dropping. These new unbiased coordinates and B values were then used as the starting point for the time-averaged studies.

Standard refinements

The unbiased single-site model described above was also used as an objective benchmark against which to judge the results of the time-averaged refinements. To judge further the results of the two time-averaged schemes, refinements were also carried out with the crystallographic packages *PROFFT* (Hendrickson & Konnert, 1980) and *SHELXL*92 (Sheldrick, 1990, 1993). As well as isotropic refinement, these two packages allow differently constrained anisotropic temperature-factor refinements. With each program both isotropic and anisotropic refinements of the myoglobin molecule were executed, starting with the unbiased starting model and using R_T^{free} as the objective criterion for convergence.

Refinements using *SHELXL* were carried out with the gamma test version of *SHELXL*92 (Sheldrick, 1993), refining against σ-weighted intensities. George Sheldrick provided the program as well as the initial script templates for refinement of myoglobin reported in this paper. Restraints were used on bond length and angle values as provided with *SHELXL*92 which were derived from studies by Engh & Huber (1991). Additional restraints were used to maintain chirality and heme planarity, and 'anti-bumping' restraints were used to keep solvent molecules reasonable distances from the protein. The program default values for the standard deviations of these parameters were used in the isotropic case. The conjugate-gradient matrix-inversion approximation option was used in all cases reported in this study and refinement was continued to convergence, usually 20 cycles. In the case of the anisotropic refinements, the DELU option was not used, and the SIMU parameter that sets the standard deviation of the variation in the six U_{ij} temperature factors between neighboring atoms was increased to 0.1 from the default value of 0.05.

Refinements performed with *PROFFT* used the standard restraints to insure correct stereochemistry. The refinement scheme employed consisted of positional and alternating B-factor and solvent-occupancy refinement for the first 15 cycles, then positional and simultaneous B-factor and solvent-occupancy refinement for the remaining 15 cycles. After every cycle, the overall scale factor was adjusted and the values for AFSIG and BFSIG were updated to maintain an r.m.s. delta bond length of 0.015 E. For isotropic B-factor refinement, SIGB's of 1.5, 2.0, 2.0, 3.0 were applied to the main-chain bond main-chain angle, side-chain bond and side-chain angle distances. In the case of three-parameter anisotropic refinement, SIGU's were obtained from the SIGB's by dividing these values by $8\pi^2$ and taking the square root. The ellipsoidal principal axes for each atom were fixed, as defined by the default thermal ellipsoid specifiers used by *PROTIN*.

Results

Results of standard refinements

Table 1 tallies, R, R_T^{free}, r.m.s. bond-length deviations and number of water molecules for each of the standard refinements carried out using the packages *X-PLOR*, *SHELXL* and *PROFFT*. In general, all models could be refined to even lower standard R values than reported here by refining more solvent sites or by overweighting the X-ray term in the residual; however, R_T^{free} was only found to increase. The stage at which R_T^{free} attained a minimum was taken as the point of convergence.

To summarize the tabulated results, it is possible to isotropically refine the protein crystal model to free R values as low as about 20–22%, with working R values in the range 14–18%, and bond-length deviations not much exceeding 0.020 Å. Anisotropic temperature-factor refinement, though capable of reducing the working R value, did not lead to any significantly better value of R_T^{free} than an isotropic model.

R values for the time-averaged ensemble

As with ordinary crystallographic refinement, we found it possible to make R arbitrarily small in both

Table 1. *Results from the two time-averaged crystallographic schemes analyzed*

For comparison are results from single-site isotropic and anisotropic refinements using various packages. All R values are for data between 5.0–1.8 Å. In *X-PLOR*, Isotropic (2) an alternative means of de-biasing the structure was tried, somewhat less aggressive than the one described in the text. In this case, a biased refined structure (from original *PROLSQ* refinement) was subjected to repeated cycles of Powell minimization with and without the X-ray 'energy' restraint comprised of the 90% working set. In this case the free R value, which initially matched the working R value, grew asymptotically larger. When it stopped rising relative to the working R value, the structure was presumed to be unbiased. This value of R_T^{free} thus represents a reasonable lower bound for the free R value.

Method	R value	Free R value	R.m.s. bond	Waters
Time-Ave				
No B values	0.108	0.280	0.020	175
B values	0.144	0.205	0.011	175
X-PLOR				
Isotropic	0.181	0.234	0.021	175
Isotropic (2)	0.159	0.197	0.023	184
SHELX				
Isotropic	0.161	0.198	0.012	179
Anisotropic	0.103	0.191	0.012	179
PROFFT				
Isotropic	0.170	0.218	0.014	184
Anisotropic	0.173	0.224	0.015	184

schemes simply by making the weight W_x of the X-ray energy $E_{\text{X-ray}}$ large enough, that is, by effectively relaxing the geometric restrains on the molecule, which in this refinement are embodied in E_{phys}. The real issue is whether R_T^{free} is well behaved for some value of the weights. What we observed for both schemes was the existence of an optimal value for the X-ray weight; values larger and smaller than the optimal weight gave poorer (larger) values of R_T^{free}.

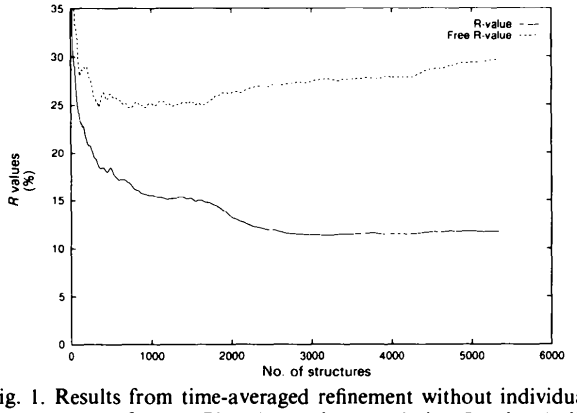

Fig. 1. Results from time-averaged refinement without individual temperature factors. Plot shows the cumulative R value (solid curve) and free R value (dotted curve) as a function of number of structures contributing to the ensemble averaged structure factors. The 5400 structures span the 'time' interval 0–54 ps. With this refinement scheme (explained in text) B values are gradually reduced until they vanish at 1600 structures. Note that R_T^{free} therefore diverges in the range of the refinement with no B values, indicating overfitting of the experimental data.

For the ensemble refinement scheme without B values (optimal weight of $W_x = 96\,000$ in *X-PLOR* units) the standard R value dropped monotonically towards 10% as more structures were added to the ensemble (Fig. 1). The minimum in R_T^{free} was 25%. However, this minimum free R value occurred within the first 16 ps, a stage in the protocol before indivi-

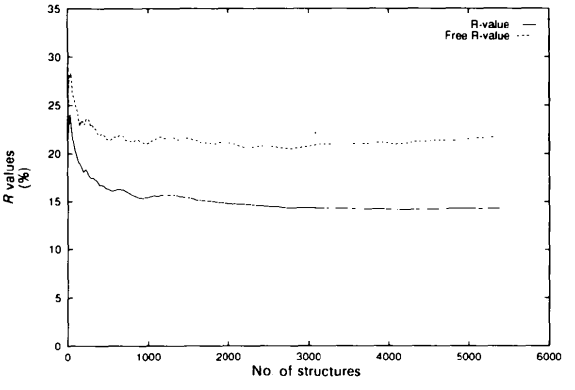

Fig. 2. Same as Fig. 1, except for a time-averaged refinement with individual temperature factors throughout the entire refinement. R_T^{free} indeed drops in the course of the refinement, until 2900 structures where it achieves its minimum. Thus the model consisting of 2900 structures, with cumulative R value 14.4%, is statistically meaningful.

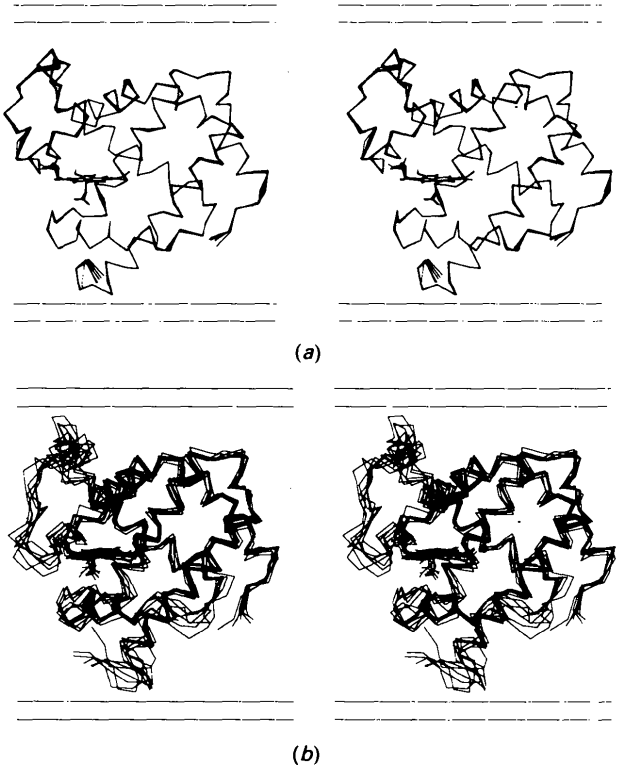

Fig. 3. Cα representation of seven uniformly spaced structures from the last 700 structures (7 ps) of the refinement. (*a*) With B values. (*b*) Without B values.

dual B values had been completely removed. Beyond the initial 16 ps, when the structures were geniunely free of B value smearing, R_T^{free} diverged quickly into the 30% range. The r.m.s. deviation on bond lengths, averaged over the last picosecond of simulation (100 structures), was 0.020 Å.

In the refinement scheme with individual B values (optimal weight of $W_x = 200\,000$ in *X-PLOR* units) the minimum value of R_T^{free} was 20.5%, with a corresponding standard $R = 14.4\%$ (Fig. 2). This best value of R_T^{free} occurred at 29 ps in the refinement, the result of 2900 structures. The r.m.s. bond-length deviation was 0.011 Å.

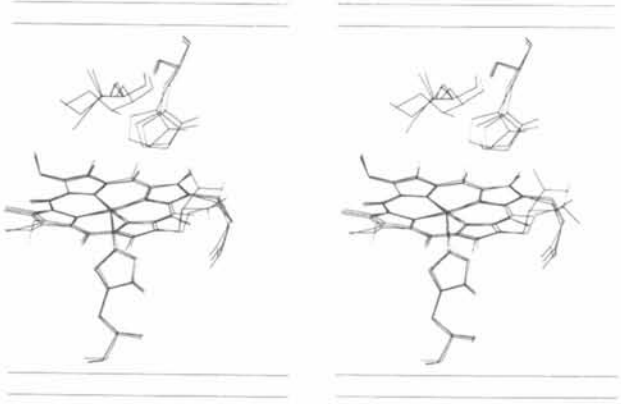

Fig. 4. Three structures from the B-value ensemble, chosen to demonstrate both anisotropy and anharmonicity in the disorder. Shown are the heme ring, the proximal histidine (below heme), the distal histidine (above and right of heme), and the mutated isoleucine (above and left of heme). Anisotropic variations are evident in the distal histidine residue. The isoleucine samples three distinct conformations. (See also Fig. 6.)

Thus we arrive at the important result: time-averaged crystallographically restrained dynamics can generate a statistically meaningful ensemble of conformations for a protein molecule, the R value and geometry of which are quite acceptable by macromolecular crystallographic standards – but only when individual B values are retained on each structure in the ensemble. Ensembles computed without B values, though having very low standard R values and acceptable geometries, have free R values much higher than a standard single-site refinement.

Furthermore, the statistical quantity (R_T^{free}) used to validate the refinement scheme can be used to obtain the ideal weight between physical and X-ray energies. Above the optimal value of the X-ray weight, R could be made even lower, but in this case R_T^{free} was always larger and would diverge sooner in the refinement. Smaller values for the weight led to larger values for both R and R_T^{free}.

Although the absolute values of the measurement errors are much lower than the model errors, R and R_T^{free} for the B-value ensemble vary with scattering angle in the same general way as R_{sym} varies. That is, over the resolution range used in the restraint (10.0–1.8 Å), R values of the model monotonically increased with resolution, parallel to the curve representing errors in the data.

Since time-averaged refinement only seems to be valid with B values included, the remainder of the paper will focus mainly on the results of this strategy.

The time-averaged ensemble and its electron density

Fig. 3(*a*) displays seven conformations, spaced 1 ps apart, from the tail end of the B-value ensemble.

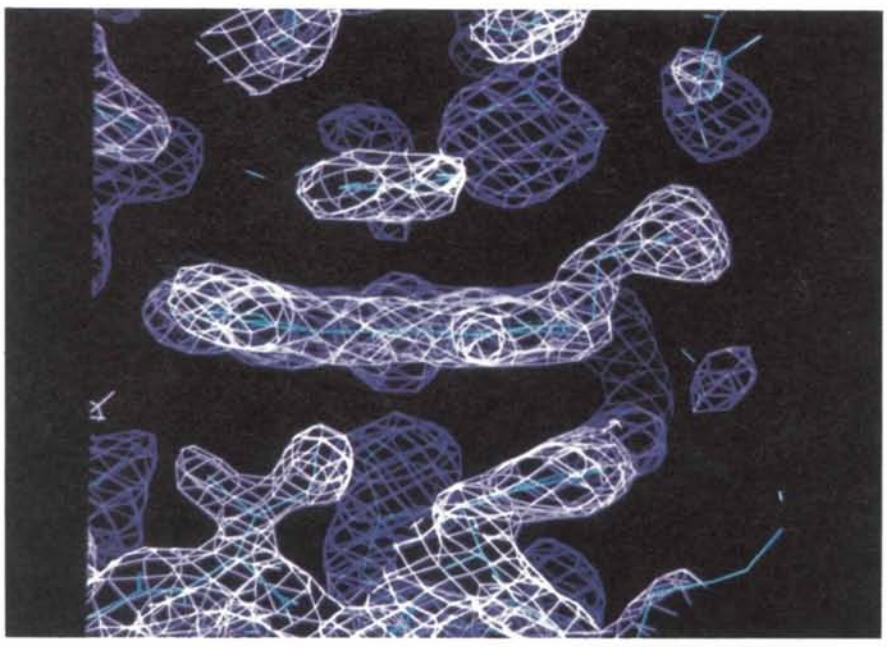

Fig. 5. $2F_o - \langle F_c \rangle_t$ map using time-averaged structure factors from the B-value ensemble. The single stick structure underlying the electron density is from the standard crystallographic refinement of the protein.

These Cα traces show the greatest mobility in two regions of the molecule: the loop bridging the *C* and *D* helices of myoglobin and the region bridging helices *E* and *F*. These regions are known, from *B* values derived from standard refinements (Quillin, Arduini, Olson & Phillips, 1993), to be the most disordered in the *P*6 crystallized protein. [For comparison, Fig. 3(*b*) shows a Cα trace for the ensemble without *B* values; obviously this set shows much more variation as it displays aspects of the random motions about the mean that are implicitly included in the *B*-factor ensemble refinements shown in 3(*a*).]

Fig. 4 depicts the atoms in the vicinity of the heme pocket, for three distinct structures. Several residues surrounding the heme exhibit considerable conformational variation. The mutated residue Ile68, for instance, assumes at least three distinct conformations.

Fig. 5 shows a $2F_o - \langle F_c \rangle_t$ map in the neighborhood of the heme pocket, where the ensemble-averaged structure factors are used for the calculated amplitudes and phases. The density map is 'clean' and is typical of maps generated by traditional methods. It is interesting to note that as well defined as the average density looks in some regions, this is not inconsistent with the fact that the conformations contributing to this average exhibit large variations about the 'ideal' atomic positions which crystallographers are accustomed to thinking in terms of when interpreting electron-density maps from protein molecules.

Anisotropy and anharmonicity in the time-averaged ensemble

The time-averaged refinement fits the observed structure-factor data while overcoming the presumptions of the standard crystallographic model, namely: isotropy and harmonicity. If time-averaged refinement overcame only the isotropy assumption, however, it would merely stand as a rather laborious, yet effective, method for refining anisotropic thermal parameters. Inspection of the structures indeed demonstrates the existence of anharmonic atomic displacements. For instance, Fig. 4 shows three structures from the ensemble, chosen to highlight the variations in two particular residues within the ligand-binding pocket of the protein. The distribution of atomic positions within the distal His64 is anisotropic; the mutated Ile68 samples three distinct conformations and is, therefore, anharmonically disordered. The departures from isotropy and harmonicity exemplified by these two residues can be illustrated further by projecting the coordinates of a particular atom onto a plane for each structure in the ensemble (Fig. 6). The set of displacements for the Cδ atom from the imidazole ring of the distal histidine traces out a highly non-isotropic, though plausibly Gaussian, distribution. Displacements of the Cγ atom in the isoleucine at position 68 fall into a tri-modal and highly non-Gaussian arrangement. The ability to distinguish between the two types of behavior is the strength of this technique.

It is straightforward to compute effective isotropic *B* values from the atomic displacements in the time-averaged ensemble ($B_{\rm eff} = 8\pi^2\langle\delta_x^2 + \delta_y^2 + \delta_z^2\rangle/3$ where

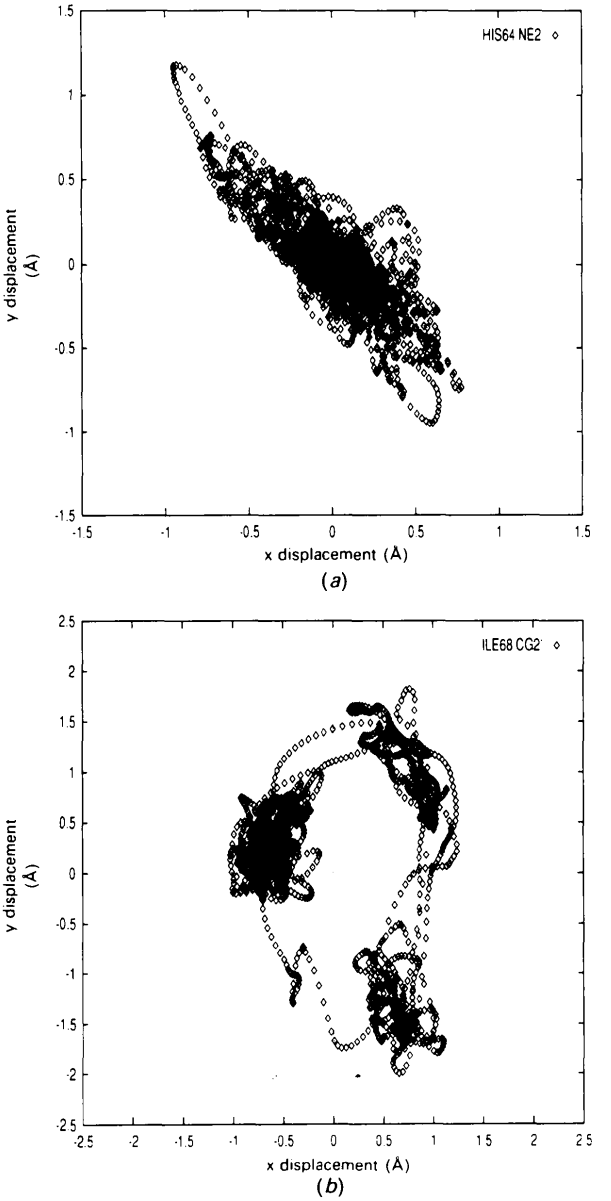

Fig. 6. Projection along the *xy* plane of the atomic displacements for two atoms in the ligand binding pocket over the course of the time-averaged refinement with *B* values. (*a*) The NE2 atom in His 64, which is anisotropically disordered. (*b*) The CG2 atom from the mutated residue Ile68, showing three conformational substates.

δ_r is the displacement of an atom from its mean coordinate $\langle r \rangle$, all averages $\langle ... \rangle$ computed in relation to the ensemble of structures). In our time-averaged refinement with B values, these effective B values represent corrections to the standard crystallographic model. We find that those parts of the protein with large effective B values have significant anisotropic and/or multi-modal distributions. These regions also correspond to those parts of the molecule which were found, through standard refinement, to have large ($B > 20$ Å^2) individual B values. Thus, the largest B values from a standard refinement still underestimate the magnitude of atomic displacement, displacements which are anisotropic and anharmonic.

This conclusion is consistent with the work of Kuriyan, Petsko, Levy & Karplus (1986) who performed a crystallographic refinement on structure-factor 'data' generated from a molecular dynamics simulation. Kuriyan and colleagues concluded that the largest atomic displacements in a protein molecule are anharmonic and that the harmonic fits imposed by standard X-ray refinement strategies greatly underestimate the true mobility except in the very well ordered regions of the molecule.

Distilling information in the time-averaged ensemble

One of the benefits of traditional refinement is that it gives an eminently interpretable, if somewhat misleading, result: a single protein structure. Since the result of time-averaged refinement is an ensemble consisting of thousands of conformations for the protein molecule, visualization, analysis and interpretation of the results is not nearly as straightforward. If we were to deposit the entire myoglobin ensemble with Brookhaven, for instance, the Protein Data Bank would more than double its present size. It is therefore desirable to develop methods for distilling the ensemble down to some smaller, representative set of structures.

One approach is to compute the quasi-harmonic modes of the distribution of conformations (Levitt, Sander & Stern, 1985; Garica, 1992). The starting point for this analysis is the atomic displacement covariance matrix,

$$\sigma(\alpha,\beta) = \langle \delta_\alpha \delta_\beta \rangle, \qquad (21)$$

where δ_α is the displacement in the atomic coordinate r_α ($\alpha = 1, 3N$, N = number of atoms) relative to its average value $\langle r_\alpha \rangle$. This matrix contains the equal-time pair correlations between the x, y, z coordinate displacements of all the atoms in the molecule, averaged over the ensemble. Diagonalization of the covariance matrix leads to an orthogonal set of $3N$ eigenvectors, which are linear combinations of the original atomic coordinates. Unlike the displacements of the physical atomic coordinates, the coordinate displacements comprising each eigenstructure are completely self-correlated, while being completely uncorrelated with any other eigenstructure. The corresponding eigenvalues represent the weight each eigenvector has in determining the character of the entire ensemble's atomic displacements. Thus, given that the trace of σ is the total mean-square displacement for all atomic coordinates, if some small number of dominant eigenvalues can account for a significant fraction of the trace, then the corresponding eigen-structures provide a reduced representation for the ensemble.

Diagonalizing the displacement covariance matrix for the myoglobin ensemble with B values, we found that 80% of the trace, or total mean-squares displacement, could be accounted for by the first 20 eigenvectors; the remaining 4000 or so modes fill in the other 20% of the trace. Although encouraging, this analysis relies solely on the quadratic covariance matrix and is thus essentially restricted to the best quasi-harmonic fit to the atomic displacements, displacements which we have shown to be manifestly anharmonic. This restriction is evident in the eigen-

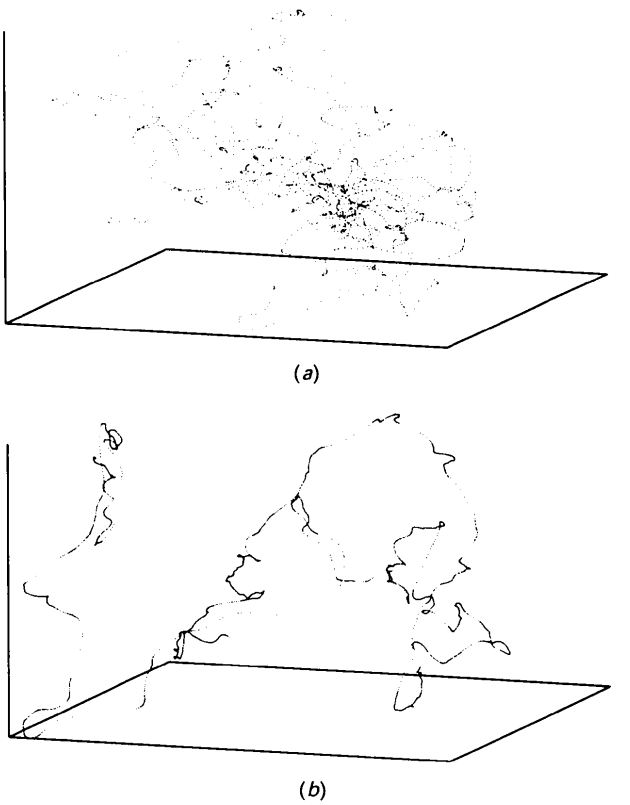

Fig. 7. (*a*) Projection of the time-averaged dynamics trajectory from $3N$-dimensional conformational phase space down onto three arbitrary orthogonal directions. (*b*) Projection of the trajectory onto the three dominant eigenvectors of the atomic displacement covariance matrix.

structures, where many regions of the molecule are simply non-physical. For example, harmonic mode fits to atoms in anharmonically disordered aromatic rings lead to stereochemically meaningless arrangements of the atoms.

One interesting use which we have found for the eigenvectors, nevertheless, is as a means of visualizing the ensemble of the protein molecule in phase space. Representing each three-dimensional conformation of the N atoms of the molecule as a single point in $3N$-dimensional phase space (each axis labeling an atomic coordinate), a molecular dynamics trajectory traces out a curve in the space. Since each conformation of the protein is now a single mathematical point, the task of making sense of an ensemble of structures is, in principle, easier – except for the fact that we are now working in thousands of dimensions. If there were a means to project everything down to two or three dimensions, though, the ensemble could be visualized. In general, projections along arbitrary orthogonal directions lead to tangled, uninterpretable curves (Fig. 7*a*). Projection of the phase space onto the three most dominant eigenvectors from the covariance matrix, though, gives a surprisingly undistorted and untangled view, or shadow, of the trajectory. Fig. 7(*b*) shows the dimensionally reduced ensemble of structures from the time-averaged refinement. Note there is no obvious clustering of the conformational states, indicating a possible difficulty in distilling the ensemble to some smaller, reprentative set based on 'overall' motions of the protein.

Discussion

The free R value was used to validate statistically one of the two schemes for time-averaged molecular dynamics refinements: modeling the structure-factor data from *P*6 myoglobin crystals with an ensemble of structures, with additional Gaussian distributions (B values) on each atom. Cross validation against 10% of the diffraction data demonstrated that up until 2900 structures, the ensemble faithfully reproduced the electron density; but, beyond this point, any new structures added to the ensemble only overfit the experimental data. The set of 2900 viable structures yields cumulative errors of $R = 14.5\%$ and $R_T^{\text{free}} = 20.5\%$. This result compares favorably against standard single-site refinements with isotropic temperature factors. Thus, as a useful alternative to a single ideal structure, this method yields an ensemble of structures consistent with the experimental Bragg scattering.

Of the several refinements included in this study, the worst model studied, from the standpoint of the more objective R_T^{free}, appears to be a time-averaged refinement scheme without individual B values. Also, models with anisotropic temperature factors, though reducing the standard R value relative to an isotropic model, actually give no better fit to the data as judged by R_T^{free}.

The time-averaged restraint was incorporated within the framework of the *X-PLOR* package (Brünger, 1987). Compared with a standard refinement of the mutant myoglobin executed within *X-PLOR* (using Powell minimization on the native structure, not full simulated annealing) the B-value ensemble refinement faired about three percentage points better in terms of R_T^{free}. *X-PLOR*, however, was able to match the accuracy of the ensemble refinements when the starting structure had been somewhat less aggressively de-biased towards the test set of reflections (see Table 1). The *SHELXL* refinements reproduced the X-ray data with the same sort accuracy as the B-value ensembles, even with aggressive de-biasing of the starting structure. Thus, the conjugate-gradient algorithm alone appears not quite as powerful as the least-squares matrix solver used by *SHELXL*. It is not inconceivable, then, that if the time-averaged restraint could be implemented within a least-squares framework such as *SHELXL* (along with the recent parameter set of Engh & Huber, 1991), the B-value ensemble refinements would yield even better results.

Though having R values not much better than a standard refinement, the time-averaged ensemble of states contains valuable structural information simply not present in a single best structure with isotropic B values. For instance, the myoglobin ensemble shows that residues involved in the biochemistry of the ligand-binding pocket exhibit both anisotropic and anharmonic multi-modal displacements. Other techniques have of course predicted such conformational variablility in protein molecules; however, the states revealed by the restrained molecular dynamics of the time-averaged refinement are now also seen to be consistent with X-ray crystallographic data.

Whereas time-averaged refinement owes its genesis to a scheme (Gros, van Gunsteren & Hol, 1990) without B values, it appears necessary to retain the notion of individual temperature factors to account faithfully for the wide range of disorder within a protein molecule. This is somewhat unsatisfying, intellectually. Perhaps further research will elucidate another means of sampling the movements which can, by itself, account for all of the atomic displacements of a macromolecule.

The isotropic Gaussian temperature factors on the atoms in the B-value ensemble account for an harmonic component in the protein's fluctuations. Atomic displacements generated from time-averaged molecular dynamics with B values, therefore, represent departures from isotropy and harmonicity in

the atomic disorder. The importance of the time-averaged refinement lies in its ability to overcome these two central assumptions about the nature of atomic displacements, which are inherent in standard crystallographic refinement. The evolving ensemble of structures is only restrained by molecular dynamics potential functions and the experimental structure-factor data, leaving the molecule free to explore a rich variety of states consistent with these two sets of information. Standard crystallographic refinements, on the other hand, by imposing Gaussian statistics, tend to locate and fit only the most harmonic component of what is often a more complex, multi-modal distribution in an atom's atomic positions; thus, the magnitudes of movement, as measured by the crystallographic B value, are underestimated for regions with high mobility, and the resulting electron density is not accurately modeled. For small molecules, or for less mobile regions of large molecules, temperature factors do adequately model the atomic disorder. But the larger the macromolecule, the more possible it becomes for atoms to make anharmonic excursions, and therefore the greater the standard crystallographic R value can tend to become - an error relation indeed born out in the history of macromolecular refinement. Hence, time-averaged refinements should prove especially valuable in investigating the structure and dynamics of large molecules, and will most likely give better R values than standard refinements.

It is clear that the time-averaged refinement can yield an ensemble which accurately fits the experimental diffraction data, at least in a scheme which includes B values. What is less obvious is why such a scheme does not suffer from catastrophic overfitting of the data. A model with thousands of conformations would seem to imply millions of floating parameters; yet cross validation with R_T^{free} demonstrates the model's meaningfullness in the sense that the ensemble can also fit data left out of the refinement process. Perhaps the conformations are not independent. After all, every conformation in the ensemble is algorithmically related to the initial one by Newton's equations of motion. Thus, perhaps there are still roughly the same number of independent parameters as in standard refinement, only in the time-averaged refinement the outside knowledge of protein dynamics is incorporated in a powerful enough way to achieve a truly accurate model for the protein molecule.

Our simulations used all solvent molecules which could be recognized as such in difference Fourier maps from the standard crystallographic refinement of myoglobin. Any other, non-ordered, solvent molecules would not contribute significantly to the structure factors in the resolution range used in this study. They would, therefore, evolve essentially unrestrained during time-averaged dynamics refinement. The argument for including bulk solvent in the ensemble refinement would be the hope that this would indirectly lead to more realistic dynamics for the ordered protein molecule itself. Bulk solvent does affect the dynamics of the protein atoms (*e.g.* by damping out atomic vibrations). Whether bulk solvent would improve the fit of the resulting ensemble to the high-resolution data is a question the answer to which lies beyond the scope of this study.

Time-averaged crystallographic refinements are computationally demanding since they involve lengthy restrained molecular dynamics simulations. Nevertheless, they yield useful, detailed structural information about macromolecules. High performance computing will be instrumental in applying these refinements to other large molecules, as well as further researching the method itself. For example, we used an averaging window of $\tau_x = 16$ps, after Gros, van Gunsteren & Hol (1990). But R_T^{free} could be used to optimize this parameter, which must be physically determined by the statistics of building up the correct electron density from a set of structures with B values, as well as by the size of the coherently scattering blocks of proteins in the crystal lattice. Also, to match the data down to the inherent noise level better, including the low-order data, all solvent in the unit cell will have to be included, a demanding computation. We have ported the structure-factor portions (electron density and FFT) of the code to a CM-2 parallel computer, achieving factor of 50 speed-ups over the SUN4 front end (Clarage & Phillips, 1992). We anticipate this refinement strategy will become more prevalent in the crystallograhic community, especially for larger macromolecules, as it is better understood and as high performance computers become more readily available.

Finally, it is still unknown whether time-averaged ensembles are unique. In other words, is the resulting ensemble the only one which could account for the experimentally measured Bragg reflections? (Or, in terms of the reduced phase space plots of the trajectory, are there other paths which could also give the same low R value?) This question is more than academic, for if the results are not unique, then there is no guarantee that the refined conformations are actually those which occur in a real protein crystal. Since the Bragg data only contain information about the average electron density in the crystal, it is not inconceivable that there could be many possible ways of generating the average density by varying the atomic positions of the polypetide chain, all the while preserving geometric constraints. The only means of unequivocally assessing the physical reality of the ensemble is to test whether the ensemble not only can predict the average electron density, but also the correct variations in the electron density.

Variations in the density, determined by the nature of the correlations in atomic displacements, produce diffuse scattering which, unlike Bragg scattering, is continuously spread throughout reciprocal space (Boylan & Phillips, 1986; Caspar, Clarage, Salunke & Clarage, 1988; Clarage, Clarage & Caspar, 1992; Chacko & Phillips, 1992). We are presently comparing computed diffuse scattering from the protein ensembles with experimentally observed diffuse scattering to resolve this crucial issue.

We acknowledge Mike Quillin and Axel Brünger for scientific discussions, and Todd Romo and Dan Sorensen for help diagonalizing a 3800×3800 matrix. We also are indebted to George Sheldrick for his help in carrying out the anisotropic refinements. Work supported by NIH AR40252, the Robert A. Welch Foundation, and the W. M. Keck Center for Computational Biology. JBC is the recipient of NRSA Postdoctoral Fellowship GM-13945 from NIH.

References

AGARWAL, R. C. (1980). *Computing in Crystallography*, edited by R. DIAMOND, S. RAMASESHAN & K. VENKATESAN, pp. 18.01–18.14. Bangalore: Indian Academy of Sciences.

BADGER, J. & CASPAR, D. L. D. (1991). *Proc. Natl Acad. Sci. USA*, **88**, 622–626.

BERENDSEN, H. J. C., POSTMA, J. P. M., VAN GUNSTEREN, N. F., DINOLA, A. & HAAK, J. R. (1984). *J. Chem. Phys.* **81**, 3684–3690.

BLAKE, C. C. F., PULFORD, W. C. A. & ARTYMIUK, P. J. (1983). *J. Mol. Biol.* **167**, 693–723.

BOYLAN, D. & PHILLIPS, G. N. (1986). *Biophys. J.* **49**, 76–78.

BROOKS, B., BRUCCOLERI, R., OLAFSON, B., STATES, D., SWAMINATHAN, S. & KARPLUS, M. (1983). *J. Comput. Chem.* **4**, 187–217.

BRÜNGER, A. T. (1987). *X-PLOR Manual.* Version 2.1. Yale Univ., New Haven, USA.

BRÜNGER, A. T. (1992). *Nature (London)*, **355**, 472–475.

BRÜNGER, A. T. (1993). *Acta Cryst.* D**49**, 24–36.

BRÜNGER, A. T., KURIYAN, J. & KARPLUS, M. (1987). *Science*, **235**, 458–460.

CASPAR, D. L. D., CLARAGE, J., SALUNKE, D. M. & CLARAGE, M. (1988). *Nature (London)*, **332**, 659–662.

CHACKO, S. & PHILLIPS, G. N. JR (1992). *Biophys. J.* **61**, 1256–1266.

CLARAGE, J., CLARAGE, M. & CASPAR, D. L. D. (1992). *Proteins*, **12**, 145–157.

CLARAGE, J. & PHILLIPS, G. N. JR (1992). *Biophys. J.* **61**, A448.

DEBYE, P. (1914). *Ann. Phys.* **43**, 49–95.

ENGH, R. A. & HUBER, R. (1991). *Acta Cryst.* A**47**, 392–400.

GARCIA, A. E. (1992). *Phys. Rev. Lett.* **68**, 2696–2699.

GROS, P. (1990). PhD thesis, Univ. of Groningen, The Netherlands.

GROS, P., VAN GUNSTEREN, W. F. & HOL, W. G. J. (1990). *Science*, **249**, 1149–1152.

HENDRICKSON, W. A. & KONNERT, J. H. (1980). *Computing in Crystallography*, edited by R. DIAMOND, S. RAMASESHAN & K. VENKATESAN, pp. 13.01–13.06. Bangalore: Indian Academy of Sciences.

HONG, M. K., BRAUNSTEIN, D., COWEN, B. R., FRAUENFELDER, H., IBEN, I. E. T., MOURANT, J. R., ORMOS, P., SCHOLL, R., SCHULTE, A., STEINBACH, P. J., XIE, A. H. & YOUNG, R. D. (1990). *Biophys. J.* **58**, 429–436.

KABSCH, W. (1988). *J. Appl. Cryst.* **21**, 916–924.

KARPLUS, M. & MCCAMMON, J. A. (1983). *Annu. Rev. Biochem.* **53**, 263–300.

KURIYAN, J., OSAPAY, K., BURLEY, S. K., BRÜNGER, A. T., HENDRICKSON, W. A. & KARPLUS, M. (1991). *Proteins*, **10**, 340–358.

KURIYAN, J., PETSKO, G. A., LEVY, M. & KARPLUS, M. (1986). *J. Mol. Biol.* **190**, 227–254.

LEVITT, M., SANDER, C. & STERN, P. (1985). *J. Mol. Biol.* **181**, 423–447.

LUNIN, V. YU. & URZHUMTSEV, A. G. (1985). *Acta Cryst.* A**41**, 327–333.

QUILLIN, M., LI, T., OLSON, J. S. & PHILLIPS, G. N. JR (1993). Submitted.

SCHOENBORN, B. P. (1988). *J. Mol. Biol.* **201**, 741–749.

SHELDRICK, G. M. (1990). *Acta Cryst.* A**46**, 467–473.

SHELDRICK, G. M. (1993). *J. Appl. Cryst.* In preparation.