

Looking Past the Pictures: Evaluating X-Ray Crystal Structure Papers

James M. Berger^{1,*}

¹Department of Molecular and Cell Biology, 374D Stanley Hall, #3220, University of California, Berkeley, Berkeley, CA 94720, USA

*Correspondence: jmberger@berkeley.edu

Introduction

There is an ever-increasing reliance on X-ray crystallography for understanding biomolecular function. This method produces atomic-resolution images of protein and nucleic acids that typically capture one or more physiologically relevant states of the molecule under investigation. Such structural information can be used to understand a wide variety of key biophysical processes, from the chemistry of enzyme catalysis and small-molecule inhibition to processes of conformational change and macromolecular assembly.

By January 2007, over 40,000 structures had been solved and deposited in the Research Collaboratory for Structural Bioinformatics (RCSB) Protein Data Bank (Berman et al., 2000). These structures have been accompanied by an explosion of papers that describe only a fraction of the tremendous variety of molecular architectures that exist in nature. Given the impact of structural data on modern molecular biological inquiry, it is important to recognize some of the limitations inherent in the method. Similarly, because “structures” are actually models that account for experimentally derived diffraction data, it is essential to understand the statistics and numbers listed in structure papers to evaluate model accuracy and veracity.

This review describes some of the important features of these papers and explains the meaning behind the numerical descriptors one is likely to encounter. For more thorough treatments of topics such as crystal growth or data acquisition and analysis, the reader is invited to see Drenth (2006), McPherson (1999), and Rhodes (2006).

Overview of the Method

To determine a crystal structure, an experimentalist first needs high-quality crystals of the protein or nucleic acid target of interest. By manipulating chemical conditions that influence solubility, many proteins and nucleic acids can be coaxed from their solution state into a crystalline array. Although crystal formation may at first seem an unnatural process that might constrain or alter the structure of a target macromolecule, numerous comparisons between crystal structures and data obtained from other spectroscopic methods (e.g., nuclear magnetic resonance; NMR) have suggested this is not the case. Indeed, crystallization is a relatively gentle process that actually captures one or more preexisting conformational states of the molecule that are already present in solution, as opposed to forcibly “wedging” the target into a rigid lattice. Moreover, protein and nucleic acid crystals are highly solvated (typically containing 40%–60% water), and possess interior macromolecule concentrations approaching those found inside cells. Finally, many enzymes retain catalytic activity in the crystalline state, permitting high-resolution imaging of their chemical reactions. Taken together, these properties mean that physiologically relevant insights into function can be derived from crystal structures.

If the molecular packing of a crystal is suitably uniform and ordered, defined diffraction patterns can be obtained from the sample upon its exposure to an intense beam of collimated X-rays (Figure 1A). X-rays are electromagnetic waves, and possess all the physical characteristics that describe a wave, including amplitude, phase, and wavelength (Figure 1B, upper). Being

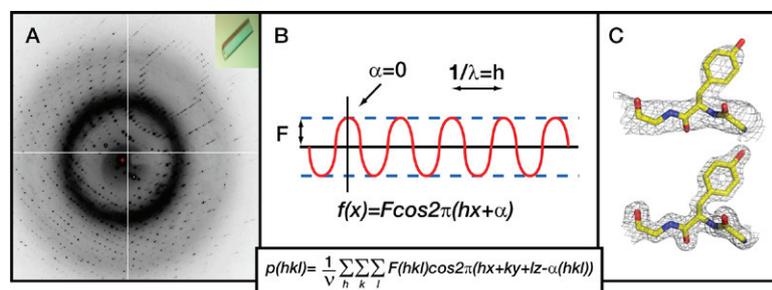


Figure 1. X-Ray Diffraction Data

(A) An example of a diffraction pattern from a protein crystal (shown in inset). Each dark “spot” on the detector corresponds to a single reflection. The large dark spot in the center marks the position of the incident X-ray beam.

(B) Waveforms and descriptors. Upper: diagram of a simple wave with amplitude (F), phase (α), and wavelength (h). Underneath is a cosine function that can be used to describe such a wave. Lower: electron-density equation. Labels are as follows: $\rho(xyz)$, electron-density value at positional coordinates x , y , and z ; $F(hkl)$, structure-factor amplitude for reflection hkl ; $\alpha(hkl)$, phase for reflection hkl ; V , volume of the unit cell.

(C) Features of electron-density maps at different resolutions. Left: segment of a fully refined structure (from Protein Data Bank ID code 1ZVT), with $2F_o - F_c$ electron density calculated to 3.0 Å resolution and contoured at 1.5 σ above the mean (Corbett et al., 2005). Right: the same segment, at 1.7 Å resolution, also contoured at 1.5 σ . As can be seen from the figure, tyrosine has a rough, “blobbish” featuredness in 3 Å resolution maps, but is defined at <2 Å resolution to the point where a hole can be seen within the aromatic ring.

high energy, X-rays typically pass through matter, but can also interact with and scatter off of atoms that happen to be in their path. As a starting point toward solving a structure, crystallographers first measure the intensities of X-rays that are scattered from a crystal; diffracted rays are termed “reflections.” The intensity of each measured reflection is proportional to the amplitude of the X-ray waveform that gave rise to it, and contains a piece of information about the spatial distribution of atoms (or, more accurately, about the “electron density” that surrounds atomic nuclei) in the crystal lattice.

Unfortunately, although phases for scattered X-rays are essential for solving a structure, they cannot be measured directly. Instead, phases must be estimated through techniques such as multiple isomorphous replacement (MIR), multiwavelength anomalous dispersion (MAD), or molecular replacement (MR). MIR relies on the incorporation of heavy-atom scatterers (e.g., mercury or platinum atoms) to subtly perturb the diffraction intensities collected from a crystal; from these deviations, and by using multiple heavy-atom “derivatives” of the crystal, phases can be calculated directly. MAD is similar to MIR in many respects, but instead of varying the type and number of different heavy atoms used, one alters the energy of the incident X-ray beam to modulate the scattering and anomalous absorptive characteristics of a heavy-atom type (typically selenium-substituted methionine) to generate differences in measured intensities. MR calculates phases from an existing model that closely matches the structure of the target (such as a homologous or mutant molecule) and that has been spatially positioned by computational means. Once the ensemble of scattered X-ray amplitudes and phasing data from one or more diffraction experiments is obtained, these data can be summed together using a Fourier transformation to reconstruct the target’s electron-density distribution in three dimensions (Figure 1B, lower). The electron-density “maps” resulting from this procedure serve as the raw material for building and interpreting a structural model (Figure 1C).

During model building, the amino acid or nucleotide sequence of the target is used as a guide to thread protein or nucleic acid chains through the electron-density map. Once a model has been built, the structure is then refined. Because the measured data are structure-factor amplitudes, a model is typically “backtransformed” into its corresponding amplitudes to allow direct comparison between observed and calculated data. The essence of refinement is to manipulate certain model parameters, such as atomic position or the magnitude of positional fluctuation about a coordinate, to achieve as good a fit as possible to the observed data. Refinement programs also reinforce appropriate stereochemical parameters, such as bond lengths and angles, until agreement between the model and data converges to an acceptable value. The end result is a “refined” structure that serves as the basis for all subsequent interpretation.

Judging the Data

Any published structure is still a *model*; it represents the structural biologist’s best interpretation of an experimentally derived distribution of electron density. To summarize the structure determination results, crystallographers usually provide a number of key parameters in a table that allows readers to judge data and model quality (Table 1). Tables are often broken into two parts: (1) data collection and/or phasing, and (2) refinement.

Data Collection

For data collection, one or more “data sets” may be listed, each of which represents a separate X-ray diffraction experiment. This particular example is for a protein structure determined using both MAD and isomorphous replacement data. The MAD data set can be seen from the two columns labeled “remote” and “peak”; these labels denote two different data-collection runs from the same crystal but using different wavelengths of incident X-ray radiation. The “peak” data set is collected at a wavelength that provides the maximal amount of anomalous scattering signal, whereas the “remote” set is often collected at a point far from this energy so as to maximize the differences in observed amplitudes between the two collection runs. The third column, labeled “HgCl₂ derivative,” contains data from a crystal that was “derivatized” by mercury atoms for phasing by MIR (technically “SIR,” in this instance, as only a single derivative is used).

Under each of the data set columns is the heading “space group.” Although the space group is not an experimental parameter one often has control over, the term is important because it refers to the ordered and symmetric arrangement of molecules within a particular crystal lattice. There are 14 types of lattice geometries (cubic, hexagonal prisms, etc.) that together encompass the 65 different kinds of space group symmetries available to proteins (Hahn and International Union of Crystallography, 1993). Space groups constrain the types of unit cell geometries; for the example shown here, P4₃2₁2 (a tetragonal lattice), two cell edges must be equal to each other and all angles must be equal to 90°. Depending on the space group and unit cell dimensions, reported values will vary greatly.

The next two parameters in Table 1 are the wavelengths used during data collection and data resolution. Wavelength (or energy) helps one to assess whether an experiment was performed to acquire anomalous scattering information (e.g., with MAD), and provides information that allows the study to be replicated. The concept of resolution is less straightforward. The resolution range is bounded by the experimenter’s ability to measure reflections over a variety of scattering angles; this term qualitatively defines the “featuredness” of an electron-density map, which in turn influences the accuracy of the final refined model (Figure 1C). However, the *precision* with which a protein/nucleic acid structure is determined is typically much better than the resolution, with positional errors in the core regions of the structure typically ranging from ~0.1–0.2 Å at 1.5 Å to ~0.5

Table 1. Data Collection and Refinement Statistics

Data Collection			
Data set	Remote	Peak	HgCl ₂ derivative
Space group	P4 ₃ 2 ₁ 2	P4 ₃ 2 ₁ 2	P4 ₃ 2 ₁ 2
Unit cell			
a, b, c (Å)	52.8, 52.8, 160.1	52.8, 52.8, 160.1	52.1, 52.1, 162.9
α, β, γ (°)	90, 90, 90	90, 90, 90	90, 90, 90
Wavelength (Å)	1.1000	0.9791	1.0093
Resolution range (Å)	45–2.4	45–2.4	43–3.1
Total reflections	206,536	72,214	41,453
Unique reflections	10,986	10,171	7,537
Redundancy	18.8 (15.5) ^a	7.1 (6.8)	5.5 (3.8)
Completeness	99.9 (100)	99.5 (100)	98.6 (99.2)
I/σ	43.3 (7.3)	41.7 (9.6)	34.9 (4.2)
R _{sym} (%) ^b	5.5 (30.8)	5.0 (20.0)	5.0 (37.4)
Phasing			
Resolution range (Å)	45–2.5	45–2.5	43–3.1
Number of sites		12	2
Phasing power ^c		2.5	1.6
Figure of merit ^d		0.48	0.46
Overall figure of merit		0.59	
Refinement Parameters			
Resolution	45–2.3		
Number of nonhydrogen atoms	29,845	Rmsd	
Number of waters	243	Bond lengths (Å)	0.013
Number of ions	3	Bond angles (°)	1.4
B factors			
Overall	30.1	Ramachandran	
Protein	29.2	Favored	88.2
Ligand/ion	39.5	Allowed	11.3
Water	34.6	Generous	0.5
R _{work} /R _{free} ^e	19.4/22.9	Disallowed	0.0

^aValues in parentheses are for highest-resolution shells.

^b $R_{sym} = \sum_j |I(h)_j - \langle I(h) \rangle| / \sum_j I(h)_j$, where $I(h)_j$ is the scaled observed intensity of the j th observation of reflection h , and $\langle I(h) \rangle$ is the mean value of corresponding symmetry-related reflections.

^c $P = \langle |F_{Hcalc}| \rangle / \langle (F_{PHobs} - F_{PHcalc}) \rangle$, where F_{PHobs} = the observed structure-factor amplitude of the derivative, F_{PHcalc} = the calculated structure-factor amplitude of the derivative, and F_{Hcalc} = the calculated structure-factor amplitude from the heavy-atom model.

^d $FOM = \int_0^{2\pi} P(\alpha)e^{i\alpha} d\alpha / \int_0^{2\pi} P(\alpha) d\alpha$, where $P(\alpha)$ is the probability that the phase angle α is correct.

^e $R_{work} = \sum ||F_{obs}| - |F_{calc}|| / \sum |F_{obs}|$, where F_{obs} and F_{calc} are observed and model structure factors, respectively. R_{free} was calculated by using a randomly selected set (5%) of reflections.

Å at 3.0 Å. This improved precision arises from the use of stereochemical constraints during refinement (discussed below). Table 1 shows that the HgCl₂ data set is of significantly lower resolution than the other two. This is not uncommon for crystals that have been treated with heavy-atom compounds, as heavy-atom binding can distort the crystal lattice. For large macromolecular complexes, resolution values are typically moderate to low (ca. 2.7–4.5 Å), whereas for smaller targets or particularly well ordered crystals, resolution can improve to better than 2 Å. In special instances, crystals can diffract to as high as 0.9–1.2 Å, or “atomic,” resolution.

The next two rows, “total reflections” and “unique reflections,” refer to the number of diffraction intensities measured for each data set. The number of total reflections simply denotes all reflections that were recorded for any particular experiment, including those that might have been measured more than once or those that are actually equivalent by virtue of crystallographic symmetry. “Unique reflections” refers to the total number of distinct reflections collected during the experiment, which accounts for the fact that crystallographically symmetric reflections can be merged into a single average measurement. The unique reflection number is defined by the resolution of the diffraction data and the size of the unit cell; when combined with the number of amino acids and/or nucleotides that occupy the asymmetric unit, these values set the “observations-to-parameters” ratio for refinement of the model.

The total number of reflections, divided by the number of unique measurements, defines the redundancy of the data. This metric lists how many times (on average) each unique reflection was measured, thus providing an estimate of the accuracy one should expect from these measurements. In Table 1, the redundancy is extremely high for one data set and moderate for the other two. Typical redundancies range from <3 to >20, depending on crystal symmetry, how many data were recorded, and/or whether one or more crystals might have been used to create a composite data set. Redundancy is often correlated with completeness, which measures how many of the total possible number of unique reflections were indeed measured. For the experiment in Table 1, the completeness lies within the upper 90th percentile, indicating that the vast majority of reflections possible for this crystal form have been measured. Most experiments should have completeness values ≥90%–95%, unless there was a specific reason why this could not be attained (e.g., extremely radiation-sensitive crystals or nonisomorphism that prevented merging data from different crystals). Likewise, highly redundant data are always desirable, though not always achievable for a variety of technical reasons.

The last two data-collection parameters are referred to as I/σ and R_{sym} . “ I ” refers to measured intensity values for reflections, while “ σ ” is the estimated standard deviation in the measurement of intensity values. Thus, I/σ , or signal to noise, refers to the average degree to which mea-

sured reflection intensities stand out over background. For example, the remote data set has an average I/σ of 43, meaning that, on average, unique reflections were approximately 43 times greater than the background noise around that reflection. Structures with high I/σ values (e.g., ≥15–20) indicate that the data are strong and imply that the quality of the data is high. R_{sym} is a measurement of how well the multiple recordings for a given unique reflection agree with one another. The formal definition of R_{sym} is listed below the table; put in simple terms, this value is a summation of the degree to which each reflection deviates from the average of all of its symmetry-related (or multiply measured) counterparts. Thus, if all data were in perfect agreement with each other (which never occurs in practice), R_{sym} should equal zero, whereas deviations from this ideal will increase R_{sym} . The R_{sym} for strong data usually ranges from ~2%–3% to 10%–15%, whereas data for weak (low I/σ) reflections can show R_{sym} values as high as 40%–50%, a phenomenon often seen in the highest-resolution shells. High R_{sym} values indicate that measured data are not in good agreement with each other, and should be taken with a degree of caution.

Phasing

The next step in the structure determination process, “phasing,” provides insights into the quality of the phase estimates and the degree of difficulty encountered solving the structure at hand. Here again, a resolution range can be associated with the experiment. This range is typically limited by the lowest resolution of the derivatives or wavelengths that are being used to determine the crystal structure.

The “number of sites” heading refers to the number of heavy-atom or anomalous scattering elements that one finds in the crystal. Such scattering centers are responsible for producing the modest intensity differences that allow crystallographers to reconstruct missing phase information using MIR or MAD. Even a single heavy-atom site can help with solving a structure, although some cases rely on dozens, or even hundreds, of sites.

The next set of values listed in Table 1 allows the reader to estimate how robustly phasing proceeded. Phasing power is a measurement of the extent to which a heavy-atom derivative contributes to phase determination: it is essentially the signal-to-noise ratio of the phasing process. Excellent phasing powers typically range from 2 to 4 or better; moderate phasing powers are usually around 1–2; and derivatives (or wavelengths) with very weak information are typically 1 or less. The figure of merit is a measurement of the probability that all of the phase angle estimates are actually correct; numerically, it is the cosine of the expected phase error. Figures of merit for MAD or MIR experiments can vary dramatically, but typically range in the order of 0.4–0.8 or so for well-estimated phases and lower for less reliable phases. In Table 1, there are actually two sets of figures of merit: one refers to the figure of merit between a particular derivative and a reference or “native” set (labeled

“remote”), whereas the other results from the combined input of the MAD/MIR data. Here, the overall figure of merit is higher than individual values, indicating that the phase information from the derivative and from the MAD experiment have reinforced one another, providing more accurate phase estimates.

Refinement

The last set of data deals with refinement. The range given for resolution generally derives from the highest-resolution reference data set. The number of nonhydrogen atoms that have been built into the model are then listed, as well as the total number of modeled waters, ions, ligands, and so forth (hydrogen atoms make a negligible contribution to X-ray scattering, are ignored in nearly all but the highest-resolution structures, and are generally not included in a final model). This information provides a gauge of the complexity of the structural problem. For example, if the number of atoms is in the hundreds of thousands, one is typically dealing with a large protein or complex, or with multiple molecules in the asymmetric unit. Because biological macromolecules are crystallized in hydrated environments, water can be included with the final model, provided it is evident in the electron-density maps. The number of water molecules added to a model varies as a function of resolution: a general rule is to add approximately one water molecule per amino acid at a resolution of 2 Å. As resolution decreases, the number of water molecules included with a structural model should decrease, and similarly should increase as resolution improves beyond 2 Å. Indeed, at >3 Å resolution, convincing density for water is typically absent, in part because the electron-density maps lack featuredness or have excess noise. At low resolutions, one should be cautious of models that have a high number of associated water molecules or ions, as there is usually insufficient information to accurately make these assignments.

A second aspect of refinement is B factor analysis. B (or temperature) factors describe the surface area of a sphere whose center corresponds to the x,y,z coordinate for each atom. These can be referred to as “ADPs,” or atomic displacement parameters, as they have actually very little to do with temperature; rather, they describe the effect of both static and dynamic disorder in the crystal. B factors thus provide an estimate of the probability that a given atom is “tightly” or “loosely” coupled to its assigned position. More precisely, B factors correspond to the root-mean-square fluctuation in position around each atom’s center.

Table 1 lists a number of different B factors, both overall B factors for the model as a whole as well as subcategories of B factors for protein, ligand, and water components. B factors are typically low for high-resolution structures, indicating that there is a high degree of certainty about each atomic position, and grow larger with medium or lower resolutions as the positional uncertainty increases. For moderate- to high-resolution structures, it is not uncommon to see B factors on the order of ~20–40

Å²; these values typically decrease as one moves to higher resolutions. By contrast, many low- or moderate-resolution structures, such as those solved at around 3.3 Å or worse, can display average B factors of >100–120 Å². B factors for protein regions typically are lower than those for waters or ligands or ions, as proteins have a well-packed hydrophobic core and a conformation that is stabilized by the crystal lattice, whereas noncovalently associated molecules may be freely exchanging with the protein’s surface. Nonetheless, it is not uncommon to see a few well-ordered waters or ligands, provided that these moieties have a suitable number of coordinating groups.

The parameters R_{work} and R_{free} are among the most important evaluators for the accuracy of the refined model. The general concept of R_{work} and R_{free} is similar to R_{sym} , except that instead of comparing the agreement between related reflections within a data set, one is now comparing the agreement between the observed structure-factor amplitudes and those calculated from the model. If the model were in perfect agreement with the data, then R_{work} would be zero; this situation never occurs in practice due to errors with the model (and to some extent the data as well). With high-quality data (~2 Å resolution), modern refinement programs typically produce a model that agrees with the observed data to an R_{work} of ~16%–22% or lower. As the resolution of the data degrades, the agreement between the model and the data will worsen, and R_{work} will concomitantly increase.

In structure papers, R_{work} is nearly always paired with R_{free} . The concept of R_{free} was first implemented by Axel Brünger as an independent validator of model quality that is unbiased by the refinement process (Brünger, 1993). To estimate and use R_{free} , one first withholds a small portion of diffraction data (the “free-R” set, ~5%–10% of the number of unique reflections), selected randomly among the available unique reflections, from refinement with the model. As refinement proceeds, model parameters are adjusted to converge with data in the working set, but are *not* exposed to measurements contained in the free-R set. At various points during refinement, the crystallographer samples the similarity between model structure-factor amplitudes and their counterparts in the free-R set, and their agreement (or disagreement) is used to monitor the process. If refinement is proceeding well and model parameters are being correctly altered, then the model structure factors should match closely with experimentally observed structure factors in the free-R set. Conversely, if the model has serious problems, the model may still agree well with the working set (and produce a low R_{work} value), but it will not agree with the free-R set.

R_{free} values are typically higher than that of the working set. The spread between R_{work} and R_{free} varies, but is generally between 2% and 6% for well-refined structures. This difference can increase for a number of reasons; however, when it does, this can serve as a warning that something is amiss with the model or the refinement.

Free-R values in the upper 30th percentile should be treated with a degree of caution, and suggest that some element of the structure may not be correctly modeled.

Perhaps the greatest caveat regarding refinement is that phases contain much more information about the electron-density distribution in the crystal than do the reflection intensities that were actually measured. Thus, refined phases become a function of the atomic coordinates and B factors. This means, particularly for low-resolution studies (say less than 2.8 Å), that a crystallographer can place an atom in a random coordinate (x,y,z) position and significant electron density will build up here, even if the position was entirely incorrect. This effect, termed “model bias,” means that it is very important, especially in low-resolution cases, that refinement decisions be guided by maps based on *measured* phases, such as those obtained from MIR or MAD experiments, or from “simulated annealing omit” approaches (Brunger et al., 1997), which help overcome this bias. It is also important that electron-density maps for critical parts of the structure be shown as figures, to be able to judge the reliability of the structural conclusions.

Two final measurements of model quality report on stereochemistry, and describe the degree to which model bond lengths, bond angles, and molecular geometries conform to accepted stereochemical standards. For good models, one typically sees root-mean-square deviations (rmsd) of bond lengths less than 0.02 Å, and rmsd bond angles less than 2°. As these rmsd values represent averages over the entire model, there may be local regions that significantly violate such limits. Ramachandran analyses are based on the empirical principle that folded macromolecular structures do not generally impose steric strain on the component residues, whereas strained or distorted local conformation is usually indicative of error in the model. Typically, the number or percentage of amino acids that fit within four different Ramachandran categories is reported. A vast majority of amino acids for a good model should fall within “favored” and “allowed” regions of Ramachandran space, although many models will at times have a few amino acids occupy the “generous” region. One typically expects that good models will not have any amino acids that fall into “disallowed” regions of Ramachandran space, although sometimes because of the size of the model or the resolution of the data, one or two amino acids may be left as outliers. When this does happen, it usually implies that the observed electron density was not sufficiently ordered to allow a particular amino acid to be properly modeled. A high number of amino acids within the disallowed region (or potentially even in the generous region) can be an indication of problems with the model.

Conclusions

This review has discussed the parameters that help a reader assess X-ray diffraction data and evaluate model quality. Because referees for peer-reviewed publications keep these values in mind as they evaluate a paper, the vast majority of structures that are published will conform to these norms. Nonetheless, it is still important to remember that all X-ray crystal structures are still models, which can vary significantly in terms of accuracy, and which often represent only one or a few structural intermediates accessed by the protein during normal function. By knowing these limitations, readers of structure papers can appreciate both the beauty and powerful insight the structures afford, while recognizing that these structures represent only an approximation of the wondrous complexity inherent in nature.

ACKNOWLEDGMENTS

The author would like to acknowledge Jacob Corn, Steve Gamblin, and Susan Marqusee for critical reading of the manuscript, and NIH grants GM071747 and CA077373 for financial support.

REFERENCES

- Berman, H.M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T.N., Weissig, H., Shindyalov, I.N., and Bourne, P.E. (2000). The Protein Data Bank. *Nucleic Acids Res.* 28, 235–242.
- Brunger, A.T. (1993). Assessment of phase accuracy by cross validation: the free R value. *Methods and applications. Acta Crystallogr. D Biol. Crystallogr.* 49, 24–36.
- Brunger, A.T., Adams, P.D., and Rice, L.M. (1997). New applications of simulated annealing in X-ray crystallography and solution NMR. *Structure* 5, 325–336.
- Corbett, K.D., Schoeffler, A.J., Thomsen, N.D., and Berger, J.M. (2005). The structural basis for substrate specificity in DNA topoisomerase IV. *J. Mol. Biol.* 351, 545–561.
- Drenth, J. (2006). *Principles of Protein X-Ray Crystallography* (New York: Springer).
- Hahn, T., and International Union of Crystallography (1993). *International tables for crystallography. In Space-Group Symmetry, Third Edition, Brief teaching edition of Volume A*, T. Hahn, ed. (Dordrecht and Boston: Kluwer Academic Publishers).
- McPherson, A. (1999). *Crystallization of Biological Macromolecules* (Cold Spring Harbor, NY: Cold Spring Harbor Laboratory Press).
- Rhodes, G. (2006). *Crystallography Made Crystal Clear: A Guide for Users of Macromolecular Models, Third Edition* (Amsterdam and Boston: Elsevier/Academic Press).
- Please cite this article as:**
Berger, J.M. (2007). Looking Past the Pictures: Evaluating X-Ray Crystal Structure Papers. In *Evaluating Techniques in Biochemical Research*, D. Zuk, ed. (Cambridge, MA: Cell Press), <http://www.cellpress.com/misc/page?page=ETBR>.