Received 12 October 2016

Accepted 12 December 2016

Edited by P. Langan, Oak Ridge National Laboratory, USA

Keywords: quantum refinement; structural biology; X-ray diffraction; neutron diffraction; cryo-EM; *cctbx*; *Q|R*.

Q|R: quantum-based refinement

Min Zheng,^{a,b} Jeffrey R. Reimers,^{a,c} Mark P. Waller^{a,b,*} and Pavel V. Afonine^{a,d,*}

^aDepartment of Physics and International Centre for Quantum and Molecular Structures, Shanghai University, Shanghai, 200444, People's Republic of China, ^bTheoretische Organische Chemie, Organisch-Chemisches Institut und Center for Multiscale Theory and Computation, Westfälische Wilhelms-Universität Münster, Corrensstrasse 40, 48149 Münster, Germany, ^cSchool of Mathematical and Physical Sciences, University of Technology Sydney, Sydney, 2007, Australia, and ^dMolecular Biophysics and Integrated Bioimaging Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA. *Correspondence e-mail: waller@shu.edu.cn, pafonine@lbl.gov

Quantum-based refinement utilizes chemical restraints derived from quantum-chemical methods instead of the standard parameterized library-based restraints used in refinement packages. The motivation is twofold: firstly, the restraints have the potential to be more accurate, and secondly, the restraints can be more easily applied to new molecules such as drugs or novel cofactors. Here, a new project called *Q|R* aimed at developing quantum-based refinement of biomacromolecules is under active development by researchers at Shanghai University together with *PHENIX* developers. The central focus of this long-term project is to develop software that is built on top of open-source components. A development version of *Q|R* was used to compare quantum-based refinements with standard refinement using a small model system.

1. Introduction

Crystallography accounts for about 90% of all structures in the Protein Data Bank (PDB; Bernstein *et al.*, 1977, 2000), and is therefore the leading tool for obtaining three-dimensional structures of biomacromolecules. Cryo-electron microscopy (cryo-EM) is rapidly becoming its major competitor (Bai *et al.*, 2015; Cheng, 2015). These two methods are rather different from technological and conceptual perspectives (Frank, 2006; Rupp, 2010); however, they both yield a map into which an initial atomic model is built. Model refinement against experimental data is the next common step in the process for both of these structure-solution techniques. For cryo-EM, the experimental data are used to construct a map, and this map normally does not change during the refinement procedure. For crystallography, the experimental data are the measured intensities of reflections and, since the phases are lost in the diffraction experiment, the map is typically calculated using model phases. This implies that the map is constantly changing, since it depends on the model, which changes during refinement. It turns out that despite these technical and methodological nuances, the computational refinement tools are very similar, if not identical, for both techniques. Therefore, we now refer to crystallographic or cryo-EM experimental data as 'experimental data' or simply 'data'.

A general refinement protocol is shown schematically in Fig. 1. Given an atomic model and experimental data, the refinement engine calculates a refinement target and its derivatives with respect to atomic parameters, which are then sent to an optimizer (typically, a minimizer). The minimizer updates the model parameters and then sends them back to the refinement engine, which then calculates a new target value and set of derivatives and returns them back to the minimizer. This process is carried out iteratively until

convergence is achieved. Finer details and specific implementation depend on the particular software and experimental data (X-ray, neutron or cryo-EM, for example). Model refinement against experimental data is an optimization process of changing the parameters that describe the model to satisfy a goal (or target) function. A target function relates the model parameters to experimental data and, if needed, *a priori* knowledge (for reviews, see Tronrud, 2004; Watkin, 2008; Afonine *et al.*, 2015). In the case of biomacromolecules the data are almost always of insufficient quality to be used alone in refinement, and thus the use of *a priori* knowledge is almost always needed, with the exception being ultrahigh-resolution data, which constitute less than 0.5% of all entries in the PDB.

A priori knowledge is typically introduced as constraints or as a weighted term ($wT_{\text{restraints}}$) to the refinement target function,

$$T = T_{\text{data}} + wT_{\text{restraints}}, \quad (1)$$

and is hereafter called 'restraints'. T_{data} is referred to as the experimental or the data term, the term that scores model to data fit, and w is the relative weight that balances the contributions of experimental data and restraints.

Most popular refinement packages such as *REFMAC* (Murshudov *et al.*, 2011), *SHELXL* (Sheldrick, 2008), *CNS* (Brünger *et al.*, 1998), *BUSTER-TNT* (Bricogne *et al.*, 2016) and *phenix.refine* (Afonine *et al.*, 2012) use a sum of potentials (e.g. harmonic) to restrain specific features of the atomic model, such as bond lengths or angles, or planes of planar groups. Typically, it is a sum of six terms,

$$T_{\text{restraints}} = T_{\text{bond}} + T_{\text{angle}} + T_{\text{planarity}} + T_{\text{chirality}} + T_{\text{torsion}} + T_{\text{nonbonded_repulsion}}, \quad (2)$$

where each term is responsible for a particular feature: covalent bonds and angles, planes, chiral volumes, torsion angles and preventing nonsensical steric clashes.

This kind of restraint is sufficient most of the time at data resolutions of 2.5–3 Å or better. However, for lower resolutions, which account for about 20% of the crystallographic

data in the PDB, or for the resolutions typically found in cryo-EM, these restraints are insufficient. Indeed, at a typical macromolecular resolution (around 2 Å) there is insufficient information to determine the atomic level of detail, but it does contain information about secondary and higher order structure. However, lower resolution data may not even contain enough information to accurately describe the secondary structure. Restraints such as those in (2) are needed to compensate for this lack of information. The impact of poorly performing restraints in (2) during refinement against low-resolution data is at least twofold. Firstly, the geometry of a refined model may not be sound; for example, α -helices and β -sheets may be distorted while still fitting the map and satisfying the restraints in (2). Secondly, data overfitting may be significant because the amount of data (experimental plus restraints) may be severely outweighed by the number of model parameters.

To address these problems, additional restraints have been used to augment (2) (see, for example, Oldfield, 2001; Echols *et al.*, 2010; Headd *et al.*, 2012; Sobolev *et al.*, 2015),

$$T_{\text{restraints_plus}} = T_{\text{restraints}} + T_{\text{SS}} + T_{\text{Ramachandran}} + T_{\text{rotamer}} + T_{\text{reference}}. \quad (3)$$

Here, T_{SS} represents secondary-structure restraints, which are essentially restraints on hydrogen-bond distances and angles. $T_{\text{Ramachandran}}$ restrains the torsion angles of protein main chain against the Ramachandran plot (Ramachandran *et al.*, 1963). T_{rotamer} restrains amino-acid side chains to valid rotameric states. $T_{\text{reference}}$ can restrain a model refined against low-resolution data to a reference model that was solved against higher resolution data.

Restraints for a standard refinement target are functions of (2) or (3); while these additional restraints are clearly an improvement, they are not without problems. For example, they require manual annotation (a user needs to tell the program what the secondary structure is) and they are still simple potentials. These potentials are fitted to reproduce some average value taken from a compiled library, and do not take into account finer details such as local environment and nearby charges. Refinement that uses such parameterized restraints is hereafter referred to as standard refinement.

A fundamentally different style of refinement is known as quantum refinement, where the restraints are derived from a quantum-chemical calculation. More specifically, the restraints are set to be the total electronic energy E , which is computed using standard quantum-chemical methods such as Hartree–Fock (Szabo & Ostlund, 2000) or density functional theory (Koch & Holthausen, 2001).

Quantum-chemical methods have the potential to play a transformational role in refinement by delivering restraints in much less of an *ad hoc* way, and this can potentially lead to more chemically meaningful structures (Carlsen & Røgen, 2015). Quantum-based refinement does not include any of the parameterized restraints such as (2) or (3), which means that no ligand-specific parameters need to be created whenever new ligands are encountered. Importantly, the choice of an

Figure 1
A general model refinement workflow. See text for details.

appropriate quantum-chemical method for a given molecular system requires a trade-off between the accuracy of the geometries and the computational cost.

Performing an accurate and efficient quantum-chemical calculation for macromolecules remains a challenge in computational chemistry (Borbulevych *et al.*, 2014, 2016; Goerigk *et al.*, 2013, 2014). However, several attempts at using quantum-chemical calculations as a source of restraints for crystallographic refinement have been reported before and can be categorized as follows.

1.1. Hybrid QM/MM

A refinement procedure can be focused on an ‘active’ region of a molecule. The advantage is that one does not waste computational resources trying to better describe the (potentially uninteresting) environment region. The QM/MM (quantum mechanics/molecular mechanics)-based refinement method advocated by Ryde and coworkers (Ryde, 2003; Ryde & Nilsson, 2003a,b; Nilsson *et al.*, 2004) was pioneering in this area. The *ComQum* software package (Ryde, 1996) was developed for this task. In particular, *ChemShell* (Sherwood *et al.*, 2003), a modular software package for QM/MM simulations, was modified to perform QM/MM-based refinement of protein X-ray structures (Hsiao *et al.*, 2010). The challenge of hybrid QM/MM-based methods is that one needs to carefully select the active QM region, ensuring that a sufficiently large region is taken. It can be time-consuming and labor-intensive to carry out convergence studies and, furthermore, finding a balanced force-field and *ab initio* combination remains an open area for QM/MM modeling.

1.2. Semi-empirical

Seminal work by Merz and coworkers has managed to address many of the issues in quantum refinement (Yu *et al.*, 2005; Yu, Li *et al.*, 2006; Yu, Hayik *et al.*, 2006; Li *et al.*, 2012; Fu *et al.*, 2013; Borbulevych *et al.*, 2014). Using semi-empirical calculations more or less alleviates the issue of computational scaling. The *DivCon* software (Dixon & Merz, 1996) was used for this purpose and has been interfaced with *PHENIX* (Adams *et al.*, 2010). Employing semi-empirical methods is attractive owing to their inherent computationally more efficient scaling (Korth & Thiel, 2011). However, the accuracy and robustness (*e.g.* metalloenzymes) issues may prove to be too much of a drawback in the long run.

1.3. Linear-scaling density functional theory

The work by Canfield *et al.* (2006) employed a divide-and-conquer-based QM/MM optimization approach to study a 150 000-atom photosystem I trimer. The whole protein is divided into individually optimized regions, with each region (and its immediate environment) treated by density functional theory (DFT) and the remaining protein by molecular mechanics. This study used the forces coming from the DFT calculation to optimize the structure. This calculation found the structural feature that held the trimer together. Serious errors in the coordinates of the chlorophyll ‘special pair’ were

identified. The orientations of 35 residue side chains were optimized to make improved hydrogen-bonding networks.

Quantum methods such as semi-empirical, Hartree–Fock or DFT can be used to calculate restraints for cofactors, co-crystals, drugs bound to active sites *etc.* The primary concern with quantum-based methods is the tremendous amount of computing resources that are required; however, recent progress in developing very efficient code, accelerated by general purpose graphical processing units (GPUs), now offers an exciting glimpse into a promising future for quantum-based refinement.

To facilitate the future development of quantum-based refinement, we set out to develop a new software package. We want to apply quantum-based restraints to the whole structure during refinement. This is a key differentiating feature compared with other previous quantum-based refinement packages, which typically only considered the active site using a QM/MM-based approach. A full quantum description has a number of benefits over a hybrid QM/MM-based approach; for example, we can alleviate the laborious step of preparing the force-field parameters of ligands and we can avoid spurious QM/MM boundary effects. Another design goal of our new software package was to make generic interfaces with many different quantum-chemical packages. This means that we are not explicitly coupled to a single quantum-chemistry package. A generic interface will facilitate rapid incorporation of any newly implemented methods from any one of the many quantum-chemical software suites. Here, we have developed a quantum-refinement package for crystallographic and cryo-EM structures called *Q|R*, and the implementation details along with an illustrative example are now reported.

2. Methods

The *Q|R* source code was written as a lightweight standalone Python (for example v.2.7) program and the source code is freely available at <https://github.com/qrefine/qrefine-core>. The command-line user interface of *Q|R* is simple, and only requires the X-ray scattering data (MTZ or CIF file) or cryo-EM map and a fully atom-complete structure (PDB or mmCIF file), along with the total charge and spin multiplicity of the molecular system, which is necessary for the quantum-chemical calculation. *Q|R* uses the *cctbx* open source library (Grosse-Kunstleve & Adams, 2002; Grosse-Kunstleve *et al.*, 2002) to construct a standard refinement protocol, very much like *phenix.refine* and most other *PHENIX* tools. *cctbx* is used to compute the data term in (1) and its derivatives, perform scaling and account for bulk solvent, and drive refinement using an L-BFGS (Liu & Nocedal, 1989; Byrd *et al.*, 1995) minimizer.

Both standard and quantum refinements in *Q|R* perform minimization under the condition (1). The difference between these two types of refinements is the source of the $T_{\text{restraints}}$. *cctbx* is used to compute $T_{\text{restraints}}$ and its derivatives using the expression (2) in standard refinement. However, the restraints ($T_{\text{restraints}}$) and its derivatives in quantum refinement are the electronic energy and analytical gradients calculated by

external quantum-chemistry software. In order to achieve this, *Q|R* interfaces with *ASE* v.3.8.1 (Bahn & Jacobsen, 2002) to enable easy access to many quantum-chemical calculators. These calculators are thin wrappers around major quantum-chemical codes. We generated a custom *ASE* calculator for *TeraChem* v.1.5 (Ufimtsev & Martinez, 2009) and modified some existing *ASE* calculators, and they are all available at <https://github.com/qrefine/qr-plugin-ase>. In this study, we have chosen to investigate three different quantum methods: semi-empirical (PM7; Stewart, 2013) in *MOPAC* v.2016 (Stewart, 2016), *ab initio* (HF/6-31G-D3; Grimme *et al.*, 2010) using *TeraChem* v.1.5 and a density functional [RI-BP86/SV(P); Von Arnim & Ahlrichs, 1998; Becke, 1988; Schäfer *et al.*, 1992] from *TURBOMOLE* v.7.0.1 (Furche *et al.*, 2013). The choice of individual quantum methods was arbitrary at this point, because the goal of this present study was to validate quantum-based refinement, not to carry out a systematic survey of candidate methods. The three different quantum-chemical approaches chosen here are vastly different methodologies.

The relative weight, denoted w in (1), is initially taken as the ratio of the gradient norm of the restraint and data terms. This weight is scaled up or down using a heuristic approach based on crystallographic statistics such as R_{work} , R_{free} and $R_{\text{free}} - R_{\text{work}}$ and geometric descriptors of the atomic model (Afonine *et al.*, 2011).

To validate the approach and its implementation, the following test was carried out. A short 13-amino-acid well ordered and resolved helix was taken as a reference from the X-ray structure of aldose reductase (PDB entry 1us0; Howard *et al.*, 2004) refined at 0.66 Å resolution. A helix was extracted

from the structure; all of the side chains were then removed to form a polyglycine reference model and saturating H atoms were added to complete the helix model (see Fig. 2). The removal of the side chains leads to so-called 'dangling bonds', and these are then capped with hydrogen to give a realistic model that is suitable for a quantum-chemical calculation (Sherwood, 2000). Since this is a very high-quality structure derived from high-resolution data, the geometry of this helix is likely to be very close to representing reality. This helix was then placed into a $16 \times 18 \times 30$ Å *P1* unit-cell box, which should be sufficiently large to have only minimal boundary effects. A low-resolution and highly incomplete set of structure factors describing all reflections in the 4–6 Å range was calculated from this model. After adding 5% of random noise to the amplitudes of these structure factors, we refer to this set as the experimental data F_{obs} .

As starting coordinates for quantum refinement, five sets of structures were constructed by applying increasing amounts of perturbation ranging from 0.3 to 1.5 Å by running molecular-dynamics simulations starting from the original model. The simulation used a simplified potential (2) from *phenix.dynamics*. This potential does not include an explicit hydrogen-bonding term and therefore cannot maintain the hydrogen-bonding interactions during the simulation (see Fig. 3). This diverse set of structures obtained from the different perturbation strengths should provide insight into the behavior of the quantum-based refinement. The structures were considered to be within a typical convergence radius of refinement. To test the robustness of our implementation, each degree of perturbation was repeated ten times using different snapshots sampled from the molecular-dynamics simulations. The original model (prior to perturbation) is taken as the reference structure in all subsequent analyses. All data presented in this work, including the scripts to reproduce reported statistics, figures and plots, are available at <https://github.com/qrefine/qr-tests-1us0>.

3. Results and discussion

In order to validate our approach and exercise the implementation (for example to eliminate bugs, optimize runtime performance and investigate the convergence radius), we choose to work with the semi-artificial system described in §2. The advantage of working with such a system is twofold. Firstly, it is small and therefore allows the sampling of diverse refinement scenarios and different parameters in a manageable amount of time (minutes to hours and not days or weeks of computer time). This is extremely important during the development stage of a project as this allows a quick

Figure 2

Aldose reductase PDB structure (left) and extracted helix model (right), with hydrogen-bond distances shown in Å.

turnaround, which in turn promotes a continuous and fast development process. Secondly, since we have constructed this system (as opposed to using real experimental data) we have full control over all of its properties and, most importantly, we know what the expected answer is. This development model has been used for more than a decade during the development of *cctbx* and many *PHENIX* tools, including writing from scratch its refinement engine *phenix.refine*, and has proven to be very efficient. Here, we adopt this paradigm for the development of our *Q|R* code.

The perturbed models described in §2 were subjected to quantum-based or standard *cctbx*-based refinement in *Q|R* using the calculated scattering data. Since the starting model is known and data are calculated from it, it is trivial to score the refinement outcomes against a known answer and compare the scores between different refinement approaches, namely standard and quantum.

It is clear from Fig. 3 that perturbed models become increasingly further removed from the reference model as the perturbation strength increases. Fig. 3 shows a lengthening of the chain owing to a greater loss of hydrogen bonds with an increasing degree of perturbation. In the set of smallest perturbations (0.3 Å r.m.s.d.) the majority of hydrogen bonds (about 60%) are retained, while in the set of most heavily perturbed structures (1.5 Å r.m.s.d.) around 97% of the hydrogen bonds are destroyed. Hence, the challenge for refinement becomes greater as the perturbation strength increases. This gives a well controlled set of models that can challenge quantum-based and standard refinement methods. Refinement is expected to return the structure back to the

original reference model, but this task becomes more challenging as the perturbation becomes stronger.

The model parameters for the test refinements were both the non-H-atom coordinates and the H-atom coordinates. The atomic displacement parameters (ADPs) were not included for this test refinement. We refined all of the 50 perturbed structures with semi-empirical (PM7), *ab initio* (HF/6-31G-D3) and density functional theory [RI-BP86/SV(P)] quantum-based methods and a standard method (*cctbx*), and the results are displayed in Figs. 4 and 5.

Fig. 4 shows the crystallographic R factors R_{work} , R_{free} and the gap $R_{\text{free}} - R_{\text{work}}$. Since we introduced a 0.05 error into the calculated F_{obs} , the expected R_{work} for converged refinement is 0.05, which corresponds to the refined structure perfectly matching the structure of the known answer. It is desirable that R_{free} stays close to R_{work} , indicating less overfitting. R_{work} and $R_{\text{free}} - R_{\text{work}}$ are marginally but systematically higher for *cctbx*-based refinements using standard restraints when compared with the quantum-based refinements across all perturbation sizes (see Fig. 4). The lower R_{work} and $R_{\text{free}} - R_{\text{work}}$ the better the fit, therefore quantum-based refinements outperform standard refinement for this model system. It is remarkable that almost all of the refinements converged to an R factor of around 5% irrespective of the restraint type (R_{work} is expected to be systematically lower, and R_{free} is expected to be systematically larger than the target 5% due to unavoidable overfitting). This R factor is to be expected because we introduced a 5% error into F_{obs} , and therefore we consider all of the refinements to be converged. Refinement of the most heavily perturbed starting structures (1.5 Å) resulted in R

Figure 3

Perturbed models with r.m.s. deviations from the starting model of 0.3, 0.6, 0.9, 1.2 and 1.5 Å, overlaying ten models per perturbation. The average percentage of conserved hydrogen bonds at each perturbation level is shown in parentheses.

factors that ranged between 5 and 9%. This is not too surprising since 1.5 Å is known to be about the limit of convergence for reciprocal-space refinement.

In addition to R factors, the number of hydrogen bonds recovered from the perturbed starting points is also monitored to check the quality of the refined helix structures. The range of valid hydrogen-bond lengths was considered to be 1.7–2.2 Å between corresponding H and non-H atoms; bonds outside this range were considered to be ‘distorted’. Hydrogen-bond distances in the helix extracted from the 1us0 model range from 1.8 to 2.1 Å (see Fig. 2). We can clearly see in Fig. 5 that refinements using the quantum-based restraints recover more of the hydrogen bonds than the refinements that employed standard restraints. In the largest perturbation (1.5 Å r.m.s.d.) only 3% of the hydrogen bonds were retained; the lowest percentage of recovered hydrogen bonds by quantum refinement is 63%, while standard refinement only recovered 25% of the hydrogen bonds. It is worth noting that while the refinement outcomes seem very similar in terms of the model-to-data fit (R factors), models refined using

quantum refinement are of much better quality based on geometric similarity (percentage of hydrogen bonds recovered) to the known reference structure. As expected, the geometry of the helix cannot recover the perturbed hydrogen bonds during refinement with standard restraints, because the restraints and low-resolution data do not contain the relevant information. This can be understood as the standard refinement does not contain any explicit hydrogen-bonding term, or even electrostatic interactions, which are dominant in hydrogen-bonding interactions.

Finally, we have demonstrated that the present version of $Q|R$ has the capability of carrying out quantum-based refinement using a number of different QM methods. The refinements of the α -helical model that were carried out using each of the quantum-based methods were remarkably similar (see Figs. 4 and 5). Ultimately, finding the single best QM method for universal application would make quantum refinement a black-box tool. This would be very useful for people wanting to perform quantum refinement who have limited knowledge of quantum chemistry. However, the search for such a method

Figure 4
Average (a) R_{work} , (b) R_{free} and (c) $R_{\text{free}} - R_{\text{work}}$ as a function of perturbation strength (Å) for semi-empirical (PM7), *ab initio* (HF/6-31G-D3), density functional [RI-BP86/SV(P)] and standard (*cctbx*) refinement. The average (ten trials per perturbation) starting R_{work} values are 0.15, 0.27, 0.35, 0.44 and 0.55, respectively, for each perturbation dose from 0.3 to 1.5 Å. Random noise (5%) was added to F_{obs} ; therefore, R is expected to be around 0.05, which would correspond to the ideal structure.

may indeed turn out to be a fool's errand, because different computational tools have been tailor-made to address particular types of problems, and a universally superior method (in terms of accuracy *versus* computational expense) remains elusive. In practice, a number of carefully carried out benchmarking studies (see, for example, Goerigk *et al.*, 2013, 2014) are required to derive a knowledge base that is sufficiently broad to cover the diversity in the PDB. These benchmarking studies are first required to make a meaningful evidence-based method selection, and they will be carried out in future work.

4. Conclusions

The *Q|R* project is focused on developing software and methods for refining biomacromolecules using chemical restraints derived from quantum mechanics. We have detailed our initial development implementation built on open-source components, which we consider as a solid starting point. In addition, we have shown a validation example in which quantum-based refinement was able to recover more of the disrupted hydrogen-bonded network in a model system, providing a glimpse of what quantum refinement can provide in the future.

Previous attempts to develop software for quantum-based refinement have been made. A *PHENIX* plugin for their linear-scaling semi-empirical *DivCon* code was developed by QuantumBio (<http://www.quantumbioinc.com>). Prior to this, the *ComQum* code was developed to locally improve a crystal structure using hybrid QM/MM methods. The development of the *Q|R* code is different from these two codes for three main reasons. Firstly, we have a multi-disciplinary team of developers from biocrystallography and quantum chemistry working together. Secondly, we see *Q|R* as being a stable bridge between the well established large quantum-chemical code bases and the open-source biocrystallographic refinement tools that are available, for example in the *cctbx* library. Therefore, we are strictly adhering to best practices in

software development for long-term sustainability. Thirdly, we are focused on developing a high-quality code base using an open-source model, and are welcoming new contributors.

It is well known that QM calculations require significant computational resources, and therefore issues related to scalability will need to be addressed in future work. Further challenges also await us, such as crystallographic symmetry and static disorder, to name but a few. To overcome these scientific and technical challenges will require significant teamwork sustained over a long period of time. Quantum refinement has the potential to become a standard technique for assisting structural biologists in obtaining high-quality structures.

Acknowledgements

MZ and MPW would like to acknowledge financial support from the Deutsche Forschungsgemeinschaft (DFG) for funding from the SFB858 project. MPW would like to acknowledge support from the Shanghai Eastern Scholar Program. PVA thanks the NIH (grant GM063210) and the PHENIX Industrial Consortium. *PyMOL* (<http://www.pymol.org>) was used for molecular graphics. ICQMS is gratefully acknowledged for support.

References

- Adams, P. D. *et al.* (2010). *Acta Cryst.* **D66**, 213–221.
- Afonine, P. V., Echols, N., Grosse-Kunstleve, R. W., Moriarty, N. W. & Adams, P. D. (2011). *Comput. Crystallogr. Newsl.* **2**, 99–103. http://www.phenix-online.org/newsletter/CCN_2011_01.pdf.
- Afonine, P. V., Grosse-Kunstleve, R. W., Echols, N., Headd, J. J., Moriarty, N. W., Mustyakimov, M., Terwilliger, T. C., Urzhumtsev, A., Zwart, P. H. & Adams, P. D. (2012). *Acta Cryst.* **D68**, 352–367.
- Afonine, P. V., Urzhumtsev, A. & Adams, P. D. (2015). *Arbor*, **191**, a219. <https://doi.org/10.3989/arbor.2015.772n2005>.
- Bahn, S. R. & Jacobsen, K. W. (2002). *Comput. Sci. Eng.* **4**, 56–66.
- Bai, X.-C., McMullan, G. & Scheres, S. H. W. (2015). *Trends Biochem. Sci.* **40**, 49–57.
- Becke, A. D. (1988). *Phys. Rev. A*, **38**, 3098–3100.
- Berman, H. M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T. N., Weissig, H., Shindyalov, I. N. & Bourne, P. E. (2000). *Nucleic Acids Res.* **28**, 235–242.
- Bernstein, F. C., Koetzle, T. F., Williams, G. J., Meyer, E. F. Jr, Brice, M. D., Rodgers, J. R., Kennard, O., Shimanouchi, T. & Tasumi, M. (1977). *J. Mol. Biol.* **112**, 535–542.
- Borbulevych, O., Martin, R. I., Tickle, I. J. & Westerhoff, L. M. (2016). *Acta Cryst.* **D72**, 586–598.
- Borbulevych, O. Y., Plumley, J. A., Martin, R. I., Merz, K. M. Jr & Westerhoff, L. M. (2014). *Acta Cryst.* **D70**, 1233–1247.
- Bricogne, G., Blanc, E., Brandl, M., Flensburg, C., Keller, P., Paciorek, W., Roversi, P., Sharff, A., Smart, O. S., Vonnrhein, C. & Womack, T. O. (2016). *BUSTER*. Cambridge: Global Phasing Ltd.
- Brünger, A. T., Adams, P. D., Clore, G. M., DeLano, W. L., Gros, P., Grosse-Kunstleve, R. W., Jiang, J.-S., Kuszewski, J., Nilges, M., Pannu, N. S., Read, R. J., Rice, L. M., Simonson, T. & Warren, G. L. (1998). *Acta Cryst.* **D54**, 905–921.
- Byrd, R. H., Lu, P., Nocedal, J. & Zhu, C. (1995). *SIAM J. Sci. Comput.* **16**, 1190–1208.
- Canfield, P., Dahlbom, M. G., Hush, N. S. & Reimers, J. R. (2006). *J. Chem. Phys.* **124**, 024301.
- Carlsen, M. & Røgen, P. (2015). *Proteins*, **83**, 1616–1624.
- Cheng, Y. (2015). *Cell*, **161**, 450–457.
- Dixon, S. L. & Merz, K. M. Jr (1996). *J. Chem. Phys.* **104**, 6643–6649.

Figure 5

The average percentage of recovered hydrogen bonds as a function of perturbation strength (Å) after refinement using either semi-empirical (PM7), *ab initio* (HF/6-31G-D3), density functional [RI-BP86/SV(P)] or standard (*cctbx*) refinement. The percentage of hydrogen bonds that remained in the perturbed models is also shown for comparison.

- Echols, N., Headd, J. J., Afonine, P. & Adams, P. (2010). *Comput. Crystallogr. Newsl.* **1**, 12–17. http://www.phenix-online.org/newsletter/CCN_2010_07.pdf.
- Frank, J. (2006). *Three-Dimensional Electron Microscopy of Macromolecular Assemblies: Visualization of Biological Molecules in Their Native State*. Oxford University Press.
- Fu, Z., Li, X., Miao, Y. & Merz, K. M. Jr (2013). *J. Chem. Theory Comput.* **9**, 1686–1693.
- Furche, F., Ahlrichs, R., Hättig, C., Klopper, W., Sierka, M. & Weigend, F. (2013). *Wiley Interdiscip. Rev. Comput. Mol. Sci.* **4**, 91–100.
- Goerigk, L., Collyer, C. A. & Reimers, J. R. (2014). *J. Phys. Chem. B*, **118**, 14612–14626.
- Goerigk, L. & Reimers, J. R. (2013). *J. Chem. Theory Comput.* **9**, 3240–3251.
- Grimme, S., Antony, J., Ehrlich, S. & Krieg, H. (2010). *J. Chem. Phys.* **132**, 154104.
- Grosse-Kunstleve, R. W. & Adams, P. D. (2002). *J. Appl. Cryst.* **35**, 477–480.
- Grosse-Kunstleve, R. W., Sauter, N. K., Moriarty, N. W. & Adams, P. D. (2002). *J. Appl. Cryst.* **35**, 126–136.
- Headd, J. J., Echols, N., Afonine, P. V., Grosse-Kunstleve, R. W., Chen, V. B., Moriarty, N. W., Richardson, D. C., Richardson, J. S. & Adams, P. D. (2012). *Acta Cryst. D* **68**, 381–390.
- Howard, E. I., Sanishvili, R., Cachau, R. E., Mitschler, A., Chevrier, B., Barth, P., Lamour, V., Van Zandt, M., Sibley, E., Bon, C., Moras, D., Schneider, T. R., Joachimiak, A. & Podjarny, A. (2004). *Proteins*, **55**, 792–804.
- Hsiao, Y. W., Sanchez-Garcia, E., Doerr, M. & Thiel, W. (2010). *J. Phys. Chem. B*, **114**, 15413–15423.
- Koch, W. & Holthausen, M. C. (2001). *A Chemist's Guide to Density Functional Theory*, 2nd ed. Weinheim: Wiley-VCH.
- Korth, M. & Thiel, W. (2011). *J. Chem. Theory Comput.* **7**, 2929–2936.
- Li, X., Fu, Z. & Merz, K. M. Jr (2012). *J. Comput. Chem.* **33**, 301–310.
- Liu, D. C. & Nocedal, J. (1989). *Math. Program.* **45**, 503–528.
- Murshudov, G. N., Skubák, P., Lebedev, A. A., Pannu, N. S., Steiner, R. A., Nicholls, R. A., Winn, M. D., Long, F. & Vagin, A. A. (2011). *Acta Cryst. D* **67**, 355–367.
- Nilsson, K., Hersleth, H.-P., Rod, T. H., Andersson, K. K. & Ryde, U. (2004). *Biophys. J.* **87**, 3437–3447.
- Oldfield, T. J. (2001). *Acta Cryst. D* **57**, 82–94.
- Ramachandran, G. N., Ramakrishnan, C. & Sasisekharan, V. (1963). *J. Mol. Biol.* **7**, 95–99.
- Rupp, B. (2010). *Biomolecular Crystallography: Principles, Practice, and Application to Structural Biology*. New York: Garland Science.
- Ryde, U. (1996). *J. Comput. Aided Mol. Des.* **10**, 153–164.
- Ryde, U. (2003). *Curr. Opin. Chem. Biol.* **7**, 136–142.
- Ryde, U. & Nilsson, K. (2003a). *J. Mol. Struct.* **632**, 259–275.
- Ryde, U. & Nilsson, K. (2003b). *J. Am. Chem. Soc.* **125**, 14232–14233.
- Schäfer, A., Horn, H. & Ahlrichs, R. (1992). *J. Chem. Phys.* **97**, 2571–2577.
- Sheldrick, G. M. (2008). *Acta Cryst. A* **64**, 112–122.
- Sherwood, P. (2000). *Modern Methods and Algorithms of Quantum Chemistry*, edited by J. Grotendorst, pp. 285–305. Jülich: John von Neumann Institute for Computing. <https://core.ac.uk/download/pdf/35009880.pdf>.
- Sherwood, P. et al. (2003). *J. Mol. Struct.* **632**, 1–28.
- Sobolev, O. V., Afonine, P. V., Adams, P. D. & Urzhumtsev, A. (2015). *J. Appl. Cryst.* **48**, 1130–1141.
- Stewart, J. J. P. (2013). *J. Mol. Model.* **19**, 1–32.
- Stewart, J. J. P. (2016). *MOPAC2016*. Stewart Computational Chemistry, Colorado Springs, USA. <http://openmopac.net>.
- Szabo, A. & Ostlund, N. S. (2000). *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*. New York: McGraw-Hill.
- Tronrud, D. E. (2004). *Acta Cryst. D* **60**, 2156–2168.
- Ufimtsev, I. S. & Martinez, T. J. (2009). *J. Chem. Theory Comput.* **5**, 2619–2628.
- Von Arnim, M. & Ahlrichs, R. (1998). *J. Comput. Chem.* **19**, 1746–1757.
- Watkin, D. (2008). *J. Appl. Cryst.* **41**, 491–522.
- Yu, N., Hayik, S. A., Wang, B., Liao, N., Reynolds, C. H. & Merz, K. M. Jr (2006). *J. Chem. Theory Comput.* **2**, 1057–1069.
- Yu, N., Li, X., Cui, G., Hayik, S. A. & Merz, K. M. Jr (2006). *Protein Sci.* **15**, 2773–2784.
- Yu, N., Yennawar, H. P. & Merz, K. M. Jr (2005). *Acta Cryst. D* **61**, 322–332.