









qFit 3: Protein and ligand multiconformer modeling for X-ray crystallographic and single-particle cryo-EM density maps

Blake T. Riley¹  | Stephanie A. Wankowicz^{2,3}  | Saulo H. P. de Oliveira⁴  |
Gydo C. P. van Zundert⁵  | Daniel W. Hogan²  | James S. Fraser²  |
Daniel A. Keedy^{1,6,7}  | Henry van den Bedem^{2,8} 

¹Structural Biology Initiative, CUNY Advanced Science Research Center, New York, New York

²Department of Bioengineering and Therapeutic Sciences, University of California San Francisco, San Francisco, California

³Biophysics Graduate Program, University of California San Francisco, San Francisco, California

⁴Frontier Medicines Corporation, South San Francisco, California

⁵Schrödinger, New York, New York

⁶Department of Chemistry and Biochemistry, City College of New York, New York, New York

⁷Ph.D. Programs in Biochemistry, Biology, and Chemistry, The Graduate Center, City University of New York, New York

⁸Atomwise, Inc., San Francisco, California

Correspondence

Daniel A. Keedy, Structural Biology Initiative, CUNY Advanced Science Research Center, New York, NY 10031.
Email: dkeedy@gc.cuny.edu
and

Henry van den Bedem, Atomwise, Inc., San Francisco, California, USA.
Email: vdbedem@atomwise.com

Funding information

National Institute of General Medical Sciences, Grant/Award Numbers: GM123159, GM124149, GM133769; National Science Foundation, Grant/Award Number: 16501133, STC-1231306

Abstract

New X-ray crystallography and cryo-electron microscopy (cryo-EM) approaches yield vast amounts of structural data from dynamic proteins and their complexes. Modeling the full conformational ensemble can provide important biological insights, but identifying and modeling an internally consistent set of alternate conformations remains a formidable challenge. qFit efficiently automates this process by generating a parsimonious multiconformer model. We refactored qFit from a distributed application into software that runs efficiently on a small server, desktop, or laptop. We describe the new qFit 3 software and provide some examples. qFit 3 is open-source under the MIT license, and is available at <https://github.com/ExcitedStates/qfit-3.0>.

KEYWORDS

conformational ensemble, cryo-electron microscopy, multiconformer models, protein dynamics, X-ray crystallography

1 | INTRODUCTION

Conformational dynamics play an essential role in many aspects of protein function, including ligand binding, allostery, and enzyme turnover.^{1,2} In each of these processes, the protein does not adopt a single conformation, but rather a conformational ensemble including a number of low-energy states. This ensemble can then be redistributed or reshaped by small-molecule binding, post-translational modifications, or other perturbations, thereby controlling biological function. To fully understand the fundamental interplay between protein conformational heterogeneity and function, it is necessary to develop experimental and computational techniques to reveal alternative protein conformations in atomic detail.

X-ray crystallography is a powerful tool for addressing this need. Because individual protein molecules in the crystal lattice sample different conformations, there is a growing appreciation that crystallographic electron density maps contain a wealth of information about sparsely populated, alternative protein conformations.³ Moreover, crystallography is undergoing an experimental renaissance: new tools are emerging with the potential to bias conformational distributions in crystals and gain new mechanistic insights into the links between protein dynamics and function.

For example, crystallographic data sets collected across multiple temperatures—as opposed to at a single cryogenic temperature—often reveal ensembles with more conformational diversity,^{4–8} including at dynamic enzyme active sites.⁹ High-throughput crystallographic protein:ligand screening can identify otherwise undetectable low-occupancy ligand-bound protein states.^{7,10,11} And time-resolved diffraction experiments, triggered by a variety of stimuli,^{12–15} can offer detailed windows into how protein conformational ensembles evolve in real time. Time-resolved experiments are becoming more accessible as serial microcrystallography experiments can take place not only at X-ray free-electron lasers with up to fs time resolution, but also at third-generation synchrotrons with microfocus beamlines with up to ms and perhaps even μ s time resolution.^{16,17} Serial microcrystallography can also help reveal alternative protein states by dissecting distinct crystal polymorphs within the microcrystal population.¹⁸ These advances, coupled with an ever-growing level of automation and faster X-ray detectors,¹⁹ are yielding larger amounts of data that highlight the need for automated (rather than manual) computational methods for modeling alternative conformations and their correlations in electron density maps.

In parallel to the renaissance for X-ray crystallography, cryo-electron microscopy (cryo-EM) is in the midst

of a “resolution revolution.”²⁰ Recently, cryo-EM structures of apoferritin at “atomic resolution” (1.2–1.25 Å)^{21,22} demonstrated how far this method has come in recent years. Similar to electron density maps from X-ray crystallography, Coulomb potential maps from cryo-EM reveal evidence for alternative protein states, which in this case are sampled by individual protein molecules on the microscopy grid. Unfortunately, so far no methods exist for unbiased and automatic modeling of correlated alternative conformations in cryo-EM maps. Additionally, many cryo-EM structures feature large protein complexes with thousands of amino acids, posing a significant challenge to traditional model building approaches. Efficient, automated algorithms²³ could meet this challenge for cryo-EM.

There is thus a clear need for computational model-building methods that better explain X-ray and cryo-EM data by incorporating alternative conformations. Protein conformational heterogeneity can be represented using various approaches, including B-factors, multi-copy ensembles, or multiconformer models.¹ First, B-factors are present for every atom in the Protein Data Bank (PDB)²⁴ file format. Theoretically, B-factors represent the harmonic, thermal displacement of each atom about its mean position, either isotropically with one parameter or anisotropically with six parameters.²⁵ However, in practice, B-factors often absorb uncertainty in a more general sense about each atom's position, and are insufficient representations of anharmonic motions such as transitions between side-chain rotamers.²⁶ Second, multi-copy ensemble models consist of some number (>1) of full copies of the protein with distinct xyz coordinates and B-factors that collectively explain the experimental data.²⁷ Ensemble models can successfully describe discrete conformational heterogeneity such as rotamer transitions—but they unnecessarily inflate the number of model parameters for those regions of the protein with essentially a single, unique conformation.²⁸ Finally, multiconformer models, such as those generated by qFit, lie somewhere in the middle in terms of model complexity. A multiconformer model represents local, anharmonic features in the data with a small number (2–5) of discrete conformations, but represents regions of the protein that show little to no evidence of flexibility with a single conformation. These conformations are assigned labels (“alternative locations” or “altlocs”), such as A, B, and so forth, with corresponding occupancies in the PDB format on a per-atom basis. Groups of atoms whose alternate positions are correlated (side chains, stretches of contiguous backbone, collective exchange across an active site, etc.) are assigned the same label and occupancy. When constructed in a parsimonious manner, multiconformer models can limit a model's complexity while maximizing

its explanatory power. Algorithms such as Ringer³ can identify amino acids with putative alternative conformations, but do not perform model building.

To efficiently generate parsimonious multiconformer models for protein X-ray crystal structures, we previously introduced the software package qFit.²⁹ Besides providing mechanistic insights, for example, by revealing hidden protein contact signaling networks³⁰ and allosteric pathways,⁷ multiconformer qFit models have also established that the conformational ensemble at room temperature is not dominated by radiation damage,³¹ and that the effect of crystal dehydration on the conformational ensemble is similar to that of cryocooling.³² We recently introduced multiconformer treatment of ligands in complex with proteins in a standalone version, *qFit-ligand*.³³ However, previous versions of qFit were computationally demanding (requiring a high-performance computing cluster), and were restricted to density maps from X-ray crystallography only, among other limitations.

Here we report a new, refactored version of qFit, which we call qFit 3, with several key improvements. qFit 3 operates on maps from either X-ray crystallography or cryo-EM. It combines multiconformer modeling of proteins and of ligands complexed with proteins (from *qFit-ligand*) in a single software package written in Python. The software distribution includes a script to refine the multiconformer model generated by qFit with Phenix.³⁴ Importantly, we reduced the runtime by two orders of magnitude compared with qFit 2. The gains in efficiency are largely due to removing redundancies, for example by pre-computing maps, parallelizing the code wherever possible, and reducing the number of Phenix refinement applications prior to qFit modeling. The R_{free} values are largely unaffected by these changes (Figure S4), exemplified by the qFit 3 model of peptidyl-prolyl cis-trans isomerase CypA in the results section below which has an R_{free} of 0.1515 compared with 0.1520 with qFit 2.⁶ qFit 3 typically runs for a ~300 residue protein in several hours on a laptop, making it significantly more accessible to users.

Overall, qFit 3 reveals hidden alternative conformations in protein structures in a rapid, automated, and unbiased manner. This new software will allow a broader array of users to explore conformational heterogeneity in their systems of interest. It will also smooth the path toward integrating new and exciting types of structural biology data, including series of data sets related by temperature, ligands, or time, as well as biologically important and/or large protein systems from X-ray free electron lasers (XFELs) or cryo-EM. qFit 3 will thus empower novel studies of the relationship between protein dynamics and biological function.

2 | RESULTS

qFit was completely refactored in the Python programming language and released as open-source software; see Methods and the GitHub repository (<https://github.com/ExcitedStates/qfit-3.0>) for more details. A typical qFit 3 workflow is illustrated in Figure 1. qFit 3 takes as minimal input a starting model and either a real-space map in the MRC/CCP4 format or map coefficients in the MTZ format. For X-ray crystallography, the preferred map is a composite omit map to minimize model bias, which can be readily generated with Phenix. For cryo-EM, the input is a real-space map together with the resolution of the data and a flag to use electron scattering factors for generating synthetic densities. qFit 3 relies on a sample-and-select procedure based on constrained optimization to identify alternative conformations of proteins and their ligands. To ensure optimal model selection and prevent overfitting, qFit 3 evaluates increasing model complexities, selecting the model with the lowest Bayesian Information Criterion (BIC). qFit 3 now also provides all functionality to model ligand alternate conformations, previously available separately in *qFit-ligand*. A distinctly important new feature is qFit 3's capability to model alternate conformations into cryo-EM maps. Numerous additional options and details are described in the Methods section and can be found in the qFit 3 GitHub repository. Descriptions and default values for all available options can be seen in the—help text of each program. Here, we demonstrate typical use cases of qFit for protein systems and their ligands. All analyses in this section used default parameters, unless otherwise stated.

2.1 | X-ray diffraction data

We first carried out qFit 3 modeling on a previously deposited cryogenic X-ray structure of a protein tyrosine phosphatase, PTPN18 (PDB ID: 2oc3).³⁵ While the deposited model includes 10 residues with alternate conformers, a difference density map shows unmodeled positive density over 3σ around Phe30 and Gln34 (Figure 2a, left panel). qFit 3 models suggest that an alternate conformer for Phe30 and an ensemble of three side-chain conformers for Gln34 better fit the density, and reduce nearby difference density peaks (Figure 2a, right panel). Running on a quad-core processor, qFit sampled and selected alternative conformations for this 290-residue protein in 12.75 hr. While runtime depends on $n_{\text{processors}}$, as well as the size, resolution, and tertiary structure of the model, around 10 CPU-minutes per residue can be expected for a 1.5 Å structure of this size.

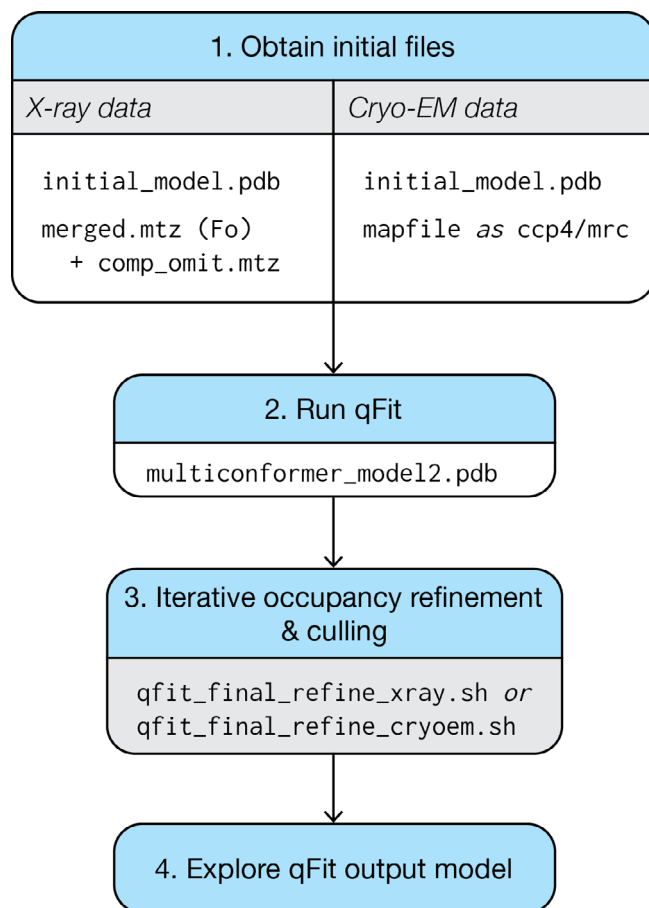


FIGURE 1 Usage flowchart for qFit 3 for either protein or ligand inputs and for either X-ray or cryo-EM data. (1) qFit requires an initial model and map information. In the case of X-ray diffraction data, qFit will require both the structure factors and a high-quality, unbiased map, such as a composite omit map. (2) With these files, qFit will generate a parsimonious model (`multiconformer_model2.pdb`) containing the fewest number of sampled conformers that explain the experimental data. (3) This intermediate/preliminary model should proceed through an iterative procedure to refine the occupancies of conformers in the model, and cull those conformers that have <9% occupancy. (4) The resulting model can then be used to explore conformational diversity. See Figures S1 and S2 for more detail on usage for X-ray versus cryo-EM data

The default algorithm of qFit 3 changed slightly compared with earlier versions. Previously, each amino acid in turn was truncated at the C β atom and refined anisotropically. This had two advantages: (a) it generally positioned the C β atom at the peak *average* density of potential alternate conformations, and (b) the anisotropy of the atomic displacement parameter provided guidance for backbone motions. Although this earlier version often better captured subtle backbone movements, it led to significant increased computational expense and complexity.³⁶ Nonetheless, the present version of qFit can be made to mimic the behavior of the earlier algorithm on a

single residue by providing an alternative input. A previously studied room-temperature structure of the peptidyl-prolyl cis-trans isomerase CypA (PDB ID: 3k0n) displays multiple conformers for Phe113.⁹ Starting from a single conformer (Figure 2b, left panel), we truncated Phe113 at C β , refined the structure anisotropically, calculated a composite omit map, and used this as input to qFit 3 with default parameters. This preprocessing enabled qFit to recapitulate the alternative conformations observed in the published model (Figure 2b, right panel). With Phe113 in place, qFit 3 ran in 460 min over the other 161 residues. This computationally expensive preprocessing procedure is provided as an option (`-phenix-per-residue-aniso`), and improved backbone modeling will be a focus of future development (Discussion).

2.2 | Cryo-EM data

qFit 3, for the first time, also accepts cryo-EM density maps as input. We have adopted the simplified scattering factor calculation of averaging the contributions of all atoms to calculate synthetic maps, as is used in real-space refinement in Phenix.³⁷ As an example application of this new functionality, we ran qFit 3 on two ultra-high-resolution cryo-EM structures: β 3 GABA receptor²² (1.2 Å resolution) and apoferritin²¹ (1.7 Å resolution). qFit 3 was run on both chain A and the entire structure for both examples. Chain A of apoferritin (176 residues) had a runtime of 112 minutes using four cores.

For these examples, qFit 3 captured both previously modeled and newly modeled alternative conformations (Figure 3). Within chain A, there were originally 19 residues with modeled alternative conformers. qFit 3 successfully identified alternate conformations for 16 (84.2%) of these residues and suggested 66 additional residues with alternative conformations. In Figure 3a, we demonstrate the ability of qFit 3 to recapitulate alternative conformers in Ser124. In Figure 3b, we demonstrate the ability of qFit 3 to detect a new alternative rotamer for Gln14 (**mm-40** and **pt20**,²⁶ RMSF 1.16 Å).

2.3 | Alternative conformations of ligands

Additionally, qFit 3 can determine alternative conformations of ligands.³³ Distinct ligand conformations can play an important role in determining binding affinities, activity, and disassociation from the protein. Visualizing ligand alternate conformations can help determine the role of entropy in binding affinity, or help guide lead optimization in drug discovery.^{38–40} `qFit-ligand` takes a model, map, and information about the position of the




FIGURE 2 qFit 3 recapitulates deposited alternate conformations in X-ray crystallography density maps, and suggests additional conformations to explain unmodeled density. (a) *Left*: PTPN18 (PDB ID: 2oc3) displays regions of unmodeled density near Phe30 and Gln34 in the deposited *mFo-DFc* difference density map at $+3\sigma$ (green cloud). These are visible in a *2mFo-DFc* composite omit density map contoured at 1σ (blue mesh), which is clarified by a low-density 0.5σ contour (blue cloud). *Right*: qFit 3 adds extra conformers to model these residues. Gln34 is modeled by three conformers (corresponding to the rotamers **mm110**, **mt0**, **mt0²⁶**); Phe30 is also modeled by two conformers (both in the “favored” **t80** rotamer space). The distance between Phe30 and Gln34 does not lead to steric hindrance between any of the conformers of either residue. Note that qFit 3 sets the minimum number of conformers in Ile33 to three (because of Gln34) to ensure backbone consistency; Phe30 is part of another backbone segment. (b) *Left*: Following the methodology in qFit 2,³⁶ Phe113 was truncated at C β and refined. Both the composite omit map and the difference map indicated the presence of at least two conformers for this residue. *Right*: qFit 3 sampled and selected two conformers of Phe113 (matching the two known ones) to explain the density in the composite omit map

ligand of interest (chain and residue number). The output is a set of conformations of the ligand. In Figure 4, we show two examples of ligands taking on multiple conformations to two different proteins, CDK2⁴¹ and Human Leukotriene A4 Hydrolase.⁴²

3 | DISCUSSION

qFit 3 is a significantly faster implementation of the qFit algorithm that can now run on commodity computer

hardware like a laptop. It is open-source and freely available, with simple installation instructions. qFit 3's speed enables application of the qFit approach to series of multiple data sets generated by new high-throughput methods in crystallography; to large, increasingly high-resolution cryo-EM structures with many thousands of amino acids; and to many more structural bioinformatics studies that focus on conformational heterogeneity.

qFit 3 can help users generate novel hypotheses about connections between conformational heterogeneity and function in protein or protein:ligand systems of interest.

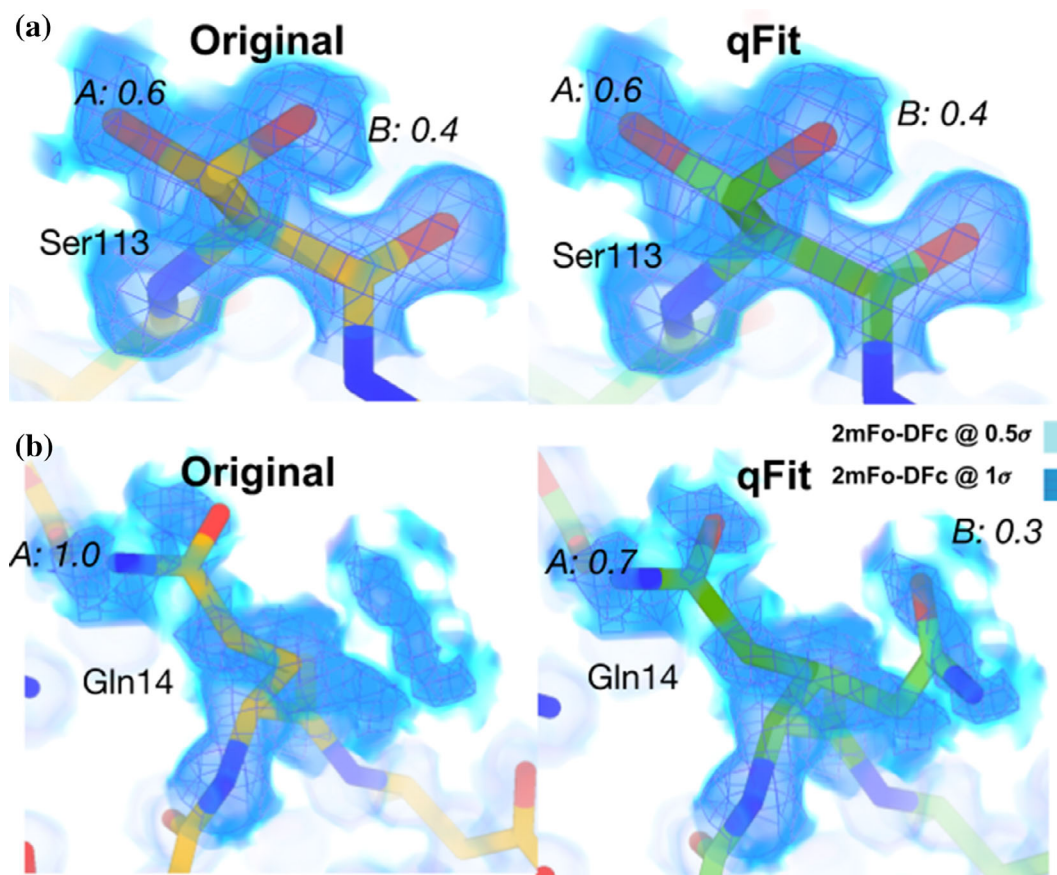


FIGURE 3 qFit 3 recapitulates deposited alternate conformations in cryo-EM density maps, and suggests alternate conformations to explain noisy data. (a) *Left*: Deposited alternative conformations for Ser113 in a high-resolution published cryo-EM structure of apoferritin (PDB ID: 6v21). These are visible in a $2mFo-DFc$ composite omit density map contoured at 1σ (dark blue cloud) and at 0.5σ (light blue cloud and blue mesh). *Right*: qFit 3 and subsequent refinement successfully modeled identical alternative conformations. Occupancies are indicated in italics. (b) *Left*: Deposited single conformation for Gln14 in the same structure of apoferritin. *Right*: qFit 3 and subsequent refinement identifies the original conformer, plus an alternative conformer (**mm-40** and **pt20** rotamers²⁶)

It readily identifies alternative protein conformations, including transitions between side-chain rotamers or even within rotameric energy wells.²⁶ Because it combines side-chain and backbone sampling, qFit often models side-chain transitions that are coupled to local backbone motions, including implicitly captured backrub-like motions.⁴³ In many cases, alternative conformations scattered throughout the protein can interact with one another, forming coupled networks⁴⁴ that may have the capacity for allosteric transmission.^{6,7} Due to the simulated-annealing labeling protocol in qFit, alternative conformation labels (A, B, etc.) are consistent for flexible regions that are adjacent in the tertiary structure, and thus may be used to infer possible energetic coupling between nearby conformations; however, the relationships between these labels for flexible regions that are separated from one another by nonflexible regions are essentially arbitrary. For protein:ligand systems, qFit can be used to contrast protein conformational heterogeneity

in the apo versus liganded states to gain mechanistic insights into pocket formation and binding. Further, qFit can identify undetected, sparsely populated secondary ligand binding poses within the same pocket, opening doors to rational structure-based ligand design to stabilize poses and to exploring the role of conformational entropy in binding thermodynamics. The alternative conformations typically visualized by qFit are often relatively local in nature, as opposed to larger-scale motions involving for example, rotations between domains, in large part because the protein is constrained by the crystal lattice. However, functional protein motions are observed within crystals,^{9,45} and thus the shifts identified by qFit can be quite relevant to functional interpretations.

Although qFit 3 can be run in an automated fashion on large (numbers of) structures, the user should apply caution in interpreting its multiconformer models. False positives can occur when qFit 3 selects spurious alternative protein conformations based on density that




FIGURE 4 qFit 3 generates occupancy-weighted multiconformer models for bound ligands. (a) *Left*: Deposited alternative conformations of thiazolylpyrimidine, an inhibitor of CDK2, in a co-crystal structure (PDB ID: 5hq5). The $2mF_o-DF_c$ composite omit density map is contoured at 1σ (dark blue cloud) and at 0.5σ (light blue cloud and grey mesh). Occupancies of alternative conformations are labeled in italics. *Right*: *qFit-ligand* successfully identifies both deposited alternative thiazolylpyrimidine conformations, as well as an additional, similar conformer. (b) *Left*: Deposited conformation of 4-(4-benzylphenyl)thiazol-2-amine, an epoxide hydrolase selective inhibitor, co-crystallized with human Leukotriene A4 Hydrolase (PDB ID: 4l2l).⁴² *Right*: *qFit-ligand* models both the deposited 4-(4-benzylphenyl)thiazol-2-amine conformation and suggests two additional conformations that, unlike the deposited conformation, fit entirely within the 1σ density contour

corresponds to other atoms such as water molecules. False negatives can occur when qFit 3 fails to sample backbone conformational space sufficiently. Development of qFit is ongoing and the user community is invited to contribute to the open-source project at <https://github.com/ExcitedStates/qfit-3.0>.

To improve qFit further, we envision several new developments. For example, qFit's backbone sampling methodology has ample room for improvement. Currently in qFit, each amino acid's backbone is translated along the principal axes of the anisotropic ellipsoid of the C β atom (or O for Gly), while closure of the backbone is maintained by torsion-based nullspace inverse kinematics (Methods), thus positioning it to accommodate suitable alternative side-chain rotamers (Methods). Although this current backbone sampling is powerful for capturing small-scale motions, it is limited in its ability to capture larger ones (Figure 2b). A suite of backbone sampling methods in qFit, ranging from explicit backrubs⁴³ and helix "shear"^{46,47} to inverse-kinematics-based loop modeling,⁴⁸ would be able to overcome this limitation. These new methods will allow qFit to model alternative conformations that are related to each other by larger, biologically relevant motions that occur in crystals, as with loops in protein tyrosine phosphatase 1B (PTP1B)⁷ and helices in isocyanide hydratase (ICH).¹⁵ A related challenge is that hierarchical alternative

conformations—such as alternative loop or helix backbone positions that each have alternative side-chain rotamers—are not supported in the existing PDB format. It may be possible to use additional restraints to bypass this limitation. For example, PanDDA helps users create a (potentially multiconformer) model of the bound state of a protein:ligand crystal system, which can then be combined with a (potentially multiconformer) model of the unbound state for crystallographic refinement using appropriate positional and chemical restraints between conformations within these states.⁴⁹ Alternatively, the new PDBx/mmCIF format that was recently adopted by the PDB could be used to explicitly define hierarchical relationships between alternative conformations.

Another important direction is improving ligand models, and correlating protein alternate conformations with alternate ligand binding modes. Currently, qFit lacks chemical knowledge of ligand atoms such as hybridization and protonation. Incorporating this knowledge, for example with the help of sophisticated force fields that work in tandem with crystallography maps,⁵⁰ will greatly improve ligand model quality and help determine the precise interactions between protein and ligand.

Finally, the problem of compositional heterogeneity must be addressed. Some of the alternative conformations in the protein may be in response to the ordering of

other components in the unit cell (heteroatoms such as ligands, crystallographic additives, and solvent). While multi-data set approaches, such as PanDDA,¹¹ may increase confidence in modeling partially occupied ligands and crystal additives, addressing the problem of partially occupied solvent may be bootstrapped by using stereotypical interactions in a solvated rotamer library.⁵¹ Solving this problem will also help to better define the border between proteins or ligands and bulk solvent,⁵² which is likely to be key to reducing the “R-factor gap in crystallography.”⁵³

4 | CONCLUSION

X-ray crystallography and cryo-electron microscopy remain the dominant experimental techniques to obtain structural information for proteins and their complexes with other macromolecules or with ligands, like therapeutic chemical compounds. New, emerging experimental techniques in X-ray crystallography and ever-increasing resolution limits in cryo-EM can reveal an ensemble of protein and ligand conformations that can provide insights into molecular mechanisms and function. qFit 3 automates interpreting an ensemble from X-ray or cryo-EM density maps, and generates an

unbiased, internally consistent, parsimonious model of conformational heterogeneity. We refactored qFit with a specific focus on efficiency and ease-of-use, so that it effortlessly installs and runs on a standard laptop to facilitate advanced interpretation of experimental structural biology data.

5 | METHODS

5.1 | qFit algorithm

5.1.1 | Overview

qFit samples numerous conformers and uses a deterministic approach to select a small ensemble of these conformers that parsimoniously explains local density. The method starts from an initial single-conformer model and generates candidate conformers for each residue/ligand in the initial structure. It evaluates all possible combinations of these conformers to determine an optimal ensemble. A final relabeling step ensures that conformers of different residues/ligands have consistent altloc labels. For all analyses in this manuscript, default parameters were used unless otherwise stated. Figure 5 provides a graphical overview of both the *qFit-protein* and *qFit-*




FIGURE 5 A flowchart of the sample-and-select protocols for (a) *qFit-protein*, and (b) *qFit-ligand*. QP = quadratic program; MIQP = mixed-integer quadratic program. See Figure S3 for a flowchart of the subsequent refinement stage

ligand algorithms, the two main command-line utilities of the qFit 3 package for automatic multiconformer modeling of proteins or ligands.

5.1.2 | Input

The qFit 3 protocol accepts input density maps or map coefficients in several commonly accepted crystallographic or cryo-EM file formats (MTZ, CCP4). For best performance, we recommend the use of a composite omit map for crystallographic densities.⁵⁴ All runs of qFit 3 on crystal structures described in this manuscript used an input composite omit map generated with the *phenix.composite_omit_map* command from the Phenix software suite.³⁴ Refinement was carried out on each partial model (*omit-type = refine*) and default parameters were used for this calculation. qFit 3 also expects a PDB file containing the structure of interest as input. Hydrogens are automatically removed to provide uniform treatment of input models. Note that during the final refinement stage, hydrogens will be (re-)added (see Final refinement script). For analyses described in this manuscript, we removed all alternate conformers (except for altloc A) using the *phenix.pdbtools* executable and used the resulting single-conformer input structure as input for all subsequent modeling.

5.1.3 | Map treatment

qFit 3 converts the input maps to absolute scale following the protocol described in Reference 55. The software creates a lookup table corresponding to the theoretical spatial density value distribution for each atomic element for radial shells spaced at 0.01 Å. The mask radius for this calculation is resolution-dependent (default radius = 0.5 Å + resolution/3). qFit 3 indirectly avoids clashes during sampling by means of a real-space density subtraction. It uses all atoms whose conformations are not being sampled to calculate a density map to perform this real-space subtraction. This prevents undesirable modeling into density from neighboring residues/side chains. The mask radius and an option to use excluded volume for clash detection instead, as detailed in Reference 33, can be determined via the command line (*-scale-rmask*, *-external-clash*). Different sets of scattering factors are used for electron density maps from X-ray crystallography vs. Coulomb potential maps from cryo-EM. For convenience, we refer to both types of maps as “density maps” in this paper.

5.1.4 | Conformational sampling for residues

qFit 3 exhaustively samples residue conformations in three stages: backbone sampling, C α -C β -C γ bond angle sampling (for certain residues), and side-chain sampling (Figure 5a). These are all enabled by default, but can be individually disabled via command-line options (*-no-backbone*, *-no-sample-angle*, *-no-sample-rotamers*).

5.1.5 | Backbone sampling

qFit 3 samples backbone conformations by means of a nullspace inverse kinematics algorithm.^{29,36,48} Backbone sampling for each residue extends to neighboring residues, two on each side. Backbone sampling is not performed if a residue lacks two neighbors on both sides (e.g., close to terminal residues).

The C β atom of the residue of interest (or O atom for Gly) is moved in the direction of the major and minor axes of its thermal ellipsoid. That motion is accommodated using its surrounding five-residue fragment to project adjustments to its dihedral angle degrees of freedom onto the subspace of motions that keeps the fragment closed (i.e., nullspace inverse kinematics). By default, three amplitudes along the ellipsoid axes are used for this sampling (0.1 + σ , 0.2 + σ , 0.3 + σ), where σ is randomly selected in the interval [-0.125, 0.125]. Altogether, three amplitudes times six directions = 18 positions for the C β (O in case of glycine) are tested. The input conformation is also added to the ensemble, leading to 19 backbone conformations after backbone sampling. Command-line options allow the user to tune amplitudes and the maximum value of σ . Expanding the number of sampled amplitudes can improve the ability to capture backbone motions, at the expense of additional computational complexity. Peptide flips³⁶ are not yet implemented in qFit 3.

5.1.6 | C α -C β -C γ bond angle sampling

For amino acids with large planar aromatic groups (Phe, Tyr, Trp, His), qFit samples around the C α -C β -C γ bond angle of the 19 backbone conformations resulting from the previous sampling step. For each conformation, we sample the C α -C β -C γ bond angles as follows: [$\theta - 7.5^\circ$, $\theta - 3.75^\circ$, θ , $\theta + 3.75^\circ$, $\theta + 7.5^\circ$]. Both the range and the step of the bond angle sampling can be adjusted via the command line. This step expands the number of sampled conformations to 95 for the large planar aromatic residues.

5.1.7 | Side-chain sampling

Side-chain sampling in qFit 3 is performed by iteratively rotating around the χ angles of ideal rotamers. The protocol begins by rotating around χ_1 . For each of the (19 or 95) backbone conformations, we rotate around each of the rotamers for the target residue in the penultimate rotamer library.²⁶ For each rotamer, we explore a sampling window using a rotamer neighborhood of $[-60^\circ, +60^\circ]$ at 10° intervals. Both the sampling window and the step size can be defined via command-line options (-rotamer-neighborhood, -dihedral-stepsizes). For the default parameters, at most $19 \times 5 \times (8+1) \times 13 = 11,115$ conformations are generated (with either Phe, Tyr, Trp, or His), which provides a balance between performance and accuracy. From this set, we remove conformations that lack support from the subtracted density map (voxel with minimum density intensity $< 0.3 \text{ e}^{-1} \text{ \AA}^{-3}$), conformations that contain self-collisions (based on hard spheres), and conformations that are redundant (using an all-atom RMSD threshold of 0.01 \AA). These exclusion strategies can be adjusted via command-line options. For protein and ligand atoms, B-factor sampling is also a non-default option.

Once the backbone and χ_1 sampling is complete, the protocol initiates a selection step based on our optimization strategy (see Optimization protocol for more details). We select all atoms starting from the backbone up to the atoms involved in the χ angle being sampled (χ_1 in this first iteration). The remaining atoms are rendered inactive, and their density contribution is not taken into account during optimization. Up to five conformers can be selected at each iteration, which then serve as the basis for sampling of subsequent χ angles.

From the second iteration onwards, we sample up to two χ angles simultaneously (also defined via the command line, -dofs-per-iteration). After sampling χ_i we exclude unsupported, clashing, and redundant conformers (as outlined above) and use this filtered conformer ensemble to sample around χ_{i+1} . In the worst-case scenario (Arg), χ_i leads to $5 \times (34+1) \times 13 = 2,275$ conformers and up to $2,275 \times (34+1) \times 13 = 1,035,125$ conformations are produced for χ_{i+1} . In practice, this number of conformations is never produced owing to redundancy. We limit the number of conformations that can be used during optimization to 15,000 for computational efficiency and memory (RAM) constraints. If sampling two χ angles in a single iteration leads to more than 15,000 conformers, we reduce sampling to a single χ for that iteration. Side-chain sampling concludes when all χ angles have been sampled.

5.1.8 | Conformational sampling for ligands

Ligand sampling in qFit 3 is performed in two steps: a local rigid body search followed by an iterative step which samples the degrees of freedom about the flexible areas of the ligand³³ (Figure 5b). For the local search, we identify all possible roots, that is, rigid fragments of atoms. Rigid fragments are defined as a set of connected atoms that do not contain a rotatable bond. We sample conformations starting from each possible root. Around the center of each ligand root, we test 100 possible rotations, by sampling rotation space in intervals of $(0^\circ, 10^\circ)$. For each rotation, we enumerate possible translations for x, y, and z coordinates in the interval $(-0.2 \text{ \AA}, 0.2 \text{ \AA})$ at 0.1 \AA increments. The local search leads to $100(\text{rotations}) \times 125$ (translations) = 12,500 conformers. We then exclude conformers that do not have support from the density (voxel with minimum density intensity $< 0.3 \text{ e}^{-1} \text{ \AA}^{-3}$) and conformations that are redundant, using an all-atom RMSD threshold cutoff of 0.01 \AA . Additionally, conformers with internal (ligand) or external clashes (receptor) are removed using a spatial hashing algorithm, which efficiently converts the 3D coordinates to a 1D hash table to determine if the sampled portion of the ligand occupies the same spatial coordinates as any other part of the ligand and/or receptor. After this exclusion step, remaining conformations are used as input for the optimization routine (see below), which selects up to five conformers of each root to best represent the local density.

Still treating each root independently, we take the root fragments selected by the local rigid body search and “expand” each fragment to the full ligand, by iteratively sampling around rotatable bonds. The protocol follows a rotatable bond hierarchy from the root to the extremities of the molecule. For each rotatable bond, we sample all angles in a $[0^\circ, 360^\circ]$ interval at 10° increments. Two rotatable bonds are sampled at a time, leading to $5 \times 36 \times 36 = 6,480$ conformations per iteration. At each iteration, we exclude conformers that do not have support from the density (voxel with minimum density intensity $< 0.3 \text{ e}^{-1} \text{ \AA}^{-3}$), those with an all-atom RMSD threshold cutoff of 0.01 \AA , or that contain internal or external clashes. After exclusion, qFit uses the optimization routine to select up to five conformers to be used for the next iteration. After all rotatable bonds have been sampled, up to five conformers can be output for each root. One final optimization step is used to select up to five consensus conformers from the pool of conformers produced across all roots.

5.1.9 | Optimization protocol

We frame the problem of selecting a subset of conformers that best represents local density as an optimization problem. Each conformer has an occupancy ω_i associated with it. The vector of all occupancies ω^T represents the regression coefficients for the optimization, with the extra constraints that the coefficients ω_i are non-negative ($\omega_i \geq 0$) and their sum lies in the unit interval ($0 \leq \sum_i \omega_i \leq 1$). We use a regression model to optimize real-space residuals, calculated from the observed density (ρ^{obs}) against the occupancy-weighted sum of the calculated densities (ρ_i^{calc}) for all conformers. We can formulate this problem as constrained quadratic optimization:

$$\begin{aligned} \min_{\omega} \quad & \|\rho^{obs} - \sum_i \omega_i \rho_i^{calc}\|_2 \\ \text{subject to} \quad & 0 \leq \sum_i \omega_i \leq 1 \\ & \omega_i \geq 0 \quad \text{for } i = 1, \dots, N. \end{aligned}$$

Residuals are calculated over all voxels within ($0.5 + \text{resolution} / 3$) Å from any active atoms across all input conformers. To prevent overfitting conformers with arbitrarily small occupancies, we require a threshold $\omega_i \geq t_{d_{min}}$ for $i = 1, \dots, N$ constraint on the occupancies, turning the problem into a mixed-integer quadratic program (MIQP):

$$\begin{aligned} \min_{\omega} \quad & \|\rho^{obs} - \sum_i \omega_i \rho_i^{calc}\|_2 \\ \text{subject to} \quad & z_i t_{d_{min}} \leq \omega_i \leq z_i, \quad z_i \in \{0, 1\}^N \\ & 0 \leq \sum_i \omega_i \leq 1, \quad \text{for } i = 1, \dots, N. \end{aligned}$$

Note that this ensures that the number of conformers selected is at most $1/t_{d_{min}}$. The optimal threshold parameter is determined using a penalized-likelihood criteria (see below). An MIQP is NP-hard (i.e., it is unknown if the problem can be solved or even verified in polynomial time), thus applying an MIQP solver directly to the conformers output from our sampling step is computationally inefficient.^{29,36} Applying a QP solver to the thousands of conformers output from our sampling routine, and then selecting the QP-fitted conformers with non-zero occupancy as input for MIQP, allows for near-optimal solutions to be calculated within a tractable time. Our protocol uses cvxopt (<https://cvxopt.org/>) and a proprietary, freely available implementation of the IBM ILOG CPLEX Optimization Studio (Python API, version 12.10) to solve QP and MIQP programs.

5.1.10 | Achieving parsimony by means of the BIC

To prevent overfitting and to ensure optimal model selection, we use the BIC to decide on model complexity. For every optimization call in qFit, we iteratively test increasing values of the threshold parameter $t_{d_{min}}$ and determine if the gain of information justifies the use of a more complex model. We fit iteratively, allowing the maximum number of conformers to vary from 1 up to 5 conformers ranked according to real-space correlation. For each iteration, we use our combined QP/MIQP routine to optimize the real-space residual sum of squares (RSS). We calculate the BIC for each level of complexity according to the following formula:

$$\text{BIC} = n \ln(\text{RSS}/n) + k \ln(n),$$

where n is the number of voxels in our resolution-dependent mask (see previous section for details) and $k = 4n_{\text{active atoms}}/t_{d_{min}}$ is the number of parameters in the model. Each active atom during sampling has four parameters: x, y, z, and B-factor. Note that the occupancies are the regression coefficients and not parameters. The factor $1/t_{d_{min}}$ is a proxy for model complexity and imposes a limit on the maximum number of conformations. We select the number of conformers that minimizes the BIC.

The -threshold option allows the user to tune the limits on model complexity relative to the default of 0.2 (i.e., max 5 conformations).

5.1.11 | Parallelization

qFit 3 can be run individually for a single residue or ligand of interest, or in parallel across a whole protein using Python's *multiprocessing* module to spawn embarrassingly parallel subprocesses that run qFit across all residues in a target protein. Up to this point, qFit samples and scores residues individually and independently: there are no interdependencies between these processes.

Once a residue has finished sampling & selecting, the selected conformers (QP/MIQP) are saved. These checkpoints are used to resume operation in the case of an abrupt exit.

5.1.12 | Validation metrics

For each residue/ligand modeled by qFit 3, we output several validation metrics, which include the BIC and the

related Akaike information criterion $AIC = 2k + n \ln(\text{RSS})$ with n and k as above. qFit 3 also reports a confidence interval for the real-space cross-correlation of the proposed conformers. The confidence interval is calculated from the Fisher z -score of the real-space cross-correlation r ⁵⁶:

$$z = 0.5 \ln((1 + r)/(1 - r))$$

Note that the z -score is approximately normally distributed with a standard deviation of $\sigma = (n - 3)^{-1/2}$, where n is the number of voxels in our resolution-dependent mask around the set of conformers being assessed. qFit 3 reports the 95% confidence interval $z \pm 1.96\sigma$ for the cross-correlation. Overlapping intervals suggest that the gain in cross-correlation is statistically not significant; we cannot reject the null hypothesis that the cross-correlations are the same at 95% confidence.

These auxiliary validation metrics are not used to filter results, but provide a guideline for balancing gain of information vs. model complexity.

5.1.13 | Building an internally consistent structural model

In the procedure above, residues are modeled independently, that is, without taking into account multi-conformer models for neighboring residues. This leads to two modeling inconsistencies. First, consecutive residues may have different occupancies for each altloc, or even a different number of alternate conformations. Second, alternate conformers of (not necessarily consecutive) side chains in a spatial neighborhood can clash owing to inconsistent assignment of altloc identifiers. To resolve these two inconsistencies, we execute two routines: qFit-segment, which addresses the problem of inconsistency along the backbone, and qFit-relabel, which resolves clashing alternate conformers between neighboring residues by reassigning altloc labels.

The qFit-segment routine starts by identifying all segments along the backbone for which all residues have at least two backbone conformers. To mark the start and end points of such backbone segments, we identify residues for which either (a) a single conformer was output, or (b) where the backbone C α and O atoms of that residue's conformers do not deviate by more than 0.05 Å. A segment is then delimited by these single-backbone-conformer residues. To create consistent segments, we proceed iteratively. We break the segments in fragments of up to 4 residues (adjustable via the command line). We enumerate all possible combinations of conformers for the fragment, which at worst case leads to $5^4 = 625$

possible conformers. We use our optimization strategy (QP/MIQP iteratively, using the BIC) to select up to five conformers per fragment based on optimal fit to the experimental map (not based on covalent geometry). To ensure consistency with the PDB file format and compatibility with refinement software, we duplicate conformers for some residues within a fragment as needed to ensure that all consecutive residues have the same number of backbone conformers. Once all 4-residue fragments have been modeled in this fashion, we proceed to enumerate all possible combinations of such length-4 fragments. This leads to fragments of at most length 16, and, again, at worst case $5^4 = 625$ possible conformers. We continue to iterate in this fashion, enumerating all possible combinations and solving/modeling, until the segment is completed. The output of the qFit-segment routine is segments, each with up to five conformers, for which the backbone is consistent, that is, for which all atoms for each conformer have the same label and occupancy.

Next, qFit-relabel relies on simulated annealing (SA) optimization of a Lennard-Jones potential to reassign altloc labels. We calculate the pairwise Lennard-Jones potential across every atom of all conformers output by qFit. Parameters for the Lennard-Jones calculation were taken from the Amber ff99SB forcefield.⁵⁷ The procedure selects five segments at random (a segment can include a single residue in this case) and randomly shuffles their labels. The user can use the -random-seed flag to ensure consistency between runs. We then assess the change in the Lennard-Jones potential and either accept or reject this move. The probability of accepting an unfavorable move is defined as:

$$P = \exp(-\Delta LJ / \text{Temperature})$$

The temperature begins at 273 K, and decreases by 10% every 10,000 perturbations. By default, 100,000 perturbations are sampled during relabeling. Benchmarking suggests that this value is sufficient for the scoring function to converge (data not shown). The output of the relabeling routine is a multiconformer model with up to five conformers per residue, in which backbones are consistent and in which alternate conformers for side chains are not clashing.

5.1.14 | Final refinement script

We performed iterative refinement on the qFit multiconformer models using version 1.18 of the Phenix software suite³⁴ to normalize the initially distorted covalent geometry, to ensure that the output models are properly fit into density (Figure S3), and to remove any unnecessary conformers.

For X-ray crystallography structures, this iterative refinement protocol uses the *phenix.refine* executable (script name: *qfit_final_refine_xray.sh*). If the resolution is >1.55 Å, ADP refinement is disabled. The initial round of refinement is done without hydrogens and uses the `strategy = *individual_sites`. We then (re-)add hydrogens to the model.⁵⁸ The next rounds of refinement use the following parameters: `"strategy = *individual_sites *individual_adp *occupancies"`, `"number_of_macro_cycles = 5"`. At each iteration, we remove all conformers for which the occupancy fell below a cutoff of 0.09. This iterative cycle continues for as long as atoms are being removed due to this occupancy cutoff criterion. We then perform one last refinement round.

For cryo-EM structures, we use a similar refinement protocol as described above, but using *phenix.real_space_refine*³⁷ (script name: *qfit_final_refine_cryoem.sh*). All rounds of real-space refinement use the default parameters.

5.2 | High-performance and cloud computing

qFit is capable of scaling from single laptops to large high-performance computing clusters. The following instructions enable qFit on Amazon's AWS, and should readily generalize to other cloud providers and RPM-based Linux distributions. We describe configurations at two different scales: a single instance and an autoscaling cluster with a free master instance.

5.2.1 | Single instance

Launch an instance that will be used to execute qFit. AWS's c5.9xlarge instance has an appropriate number of cores and amount of memory for most proteins.

The following Bash script, reproduced from `docs/aws_deploy.sh` in the qFit repository, installs qFit and its dependencies within a conda environment:

```
#!/usr/bin/env bash.
# Tested on Amazon Linux 2, but should work on most
RPM-based Linux distros
# install Anaconda RPM GPG keys
sudo rpm -import https://repo.anaconda.com/pkgs/misc/
gpgkeys/anaconda.asc
# add Anaconda repository
cat <<EOF | sudo tee /etc/yum.repos.d/conda.repo
[conda]
name = Conda
```

```
baseurl=https://repo.anaconda.com/pkgs/misc/rpmrepo/
conda
enabled = 1
gpgcheck = 1
gpgkey=https://repo.anaconda.com/pkgs/misc/gpgkeys/
anaconda.asc
EOF
```

```
sudo yum -y install conda
sudo yum -y install git gcc
```

```
source /opt/conda/etc/profile.d/conda.sh
conda create -y-name qfit
conda activate qfit
```

```
conda install -y -c anaconda mkl
conda install -y -c anaconda -c ibmdecisionoptimization
cvxopt cplex
```

```
git clone https://github.com/ExcitedStates/qfit-3.0.git
cd qfit-3.0/
# Optionally, uncomment the following line to set a specific
version of qFit
#git checkout v3.2.0
pip install.
```

Consider creating an image of the instance at this point to avoid executing the above script each time an instance is launched from a base instance.

After installation, it is necessary to execute `source /opt/conda/etc/profile.d/conda.sh` to set up conda within your Bash shell then activate the conda environment by executing `conda activate qfit`.

Using the example described in qFit's README.md, alternative conformers for all residues in 3K0N can be calculated by executing `qfit_protein 3K0N.mtz -l 2FOFCWT, PH2FOFCWT 3K0N.pdb -p 36` for 3K0N.mtz and 3K0N.pdb in the current working directory, utilizing up to 36 cores.

5.2.2 | Autoscaling cluster

Additionally, ParallelCluster can be used to create an autoscaling cluster to maximize efficiency of cloud resources. See Supplementary Methods for details.

ACKNOWLEDGMENTS

SAW is supported by NSF GRFP 16501133. JSF is supported by NIH GM123159, NIH GM124149, and NSF STC-1231306. DAK is supported by NIH GM133769.

AUTHOR CONTRIBUTIONS

Blake Riley: Investigation; methodology; software; validation; visualization; writing-original draft; writing-review and editing. **Stephanie Wankowicz:** Investigation; software; validation; writing-original draft; writing-review and editing. **Saulo de Oliveira:** Formal analysis; methodology; software. **Gydo van Zundert:** Formal analysis; methodology; software. **Daniel Hogan:** Methodology; software; writing-original draft. **James Fraser:** Conceptualization; funding acquisition; writing-original draft; writing-review and editing. **Daniel Keedy:** Conceptualization; methodology; supervision; writing-original draft; writing-review and editing. **Henry van den Bedem:** Conceptualization; methodology; supervision; writing-original draft; writing-review and editing.

ORCID

Blake T. Riley  <https://orcid.org/0000-0003-2176-0503>
 Stephanie A. Wankowicz  <https://orcid.org/0000-0002-4225-7459>
 Saulo H. P. de Oliveira  <https://orcid.org/0000-0002-9230-7610>
 Gydo C. P. van Zundert  <https://orcid.org/0000-0001-9924-3317>
 Daniel W. Hogan  <https://orcid.org/0000-0003-3375-408X>
 James S. Fraser  <https://orcid.org/0000-0002-5080-2859>
 Daniel A. Keedy  <https://orcid.org/0000-0002-9184-7586>
 Henry van den Bedem  <https://orcid.org/0000-0003-2358-841X>

REFERENCES

- van den Bedem H, Fraser JS. Integrative, dynamic structural biology at atomic resolution—it's about time. *Nat Methods*. 2015;12:307–318.
- Aviram HY, Pirchi M, Mazal H, Barak Y, Riven I, Haran G. Direct observation of ultrafast large-scale dynamics of an enzyme under turnover conditions. *Proc Natl Acad Sci U S A*. 2018;15:3243–3248.
- Lang PT, Ng H-L, Fraser JS, et al. Automated electron-density sampling reveals widespread conformational polymorphism in proteins. *Protein Sci*. 2010;19:1420–1431.
- Fraser JS, van den Bedem H, Samelson AJ, et al. Accessing protein conformational ensembles using room-temperature X-ray crystallography. *Proc Natl Acad Sci U S A*. 2011;108:16247–16252.
- Keedy DA, van den Bedem H, Sivak DA, et al. Crystal cryocooling distorts conformational heterogeneity in a model Michaelis complex of DHFR. *Structure*. 2014;22:899–910.
- Keedy DA, Kenner LR, Warkentin M, et al. Mapping the conformational landscape of a dynamic enzyme by multi-temperature and XFEL crystallography. *eLife*. 2015;4:e07574.
- Keedy DA, Hill ZB, Biel JT, Kang E, Rettenmaier TJ. An expanded allosteric network in PTP1B by multitemperature crystallography, fragment screening, and covalent tethering. *eLife*. 2018;7:e36307.
- Doukov T, Herschlag D, Yabukarski F. A robust method for collecting X-ray diffraction data from protein crystals across physiological temperatures. *bioRxiv*. 2020. 995852.
- Fraser JS, Clarkson MW, Degnan SC, Erion R, Kern D, Alber T. Hidden alternative structures of proline isomerase essential for catalysis. *Nature*. 2009;462:669–673.
- Pearce NM, Bradley AR, Krojer T, Marsden BD, Deane CM, von Delft F. Partial-occupancy binders identified by the pan-dataset density analysis method offer new chemical opportunities and reveal cryptic binding sites. *Struct Dyn*. 2017;4:032104.
- Pearce NM, Krojer T, Bradley AR, et al. A multi-crystal method for extracting obscured crystallographic states from conventionally uninterpretable electron density. *Nat Commun*. 2017;8:15123.
- Tenboer J, Basu S, Zatsepin N, et al. Time-resolved serial crystallography captures high-resolution intermediates of photoactive yellow protein. *Science*. 2014;346:1242–1246.
- Hekstra DR, White KI, Socolich MA, Henning RW, Šrajcar V, Ranganathan R. Electric-field-stimulated protein mechanics. *Nature*. 2016;540:400–405.
- Thompson MC, Barad BA, Wolff AM, et al. Temperature-jump solution X-ray scattering reveals distinct motions in a dynamic enzyme. *Nat Chem*. 2019;11:1058–1066.
- Dasgupta M, Budday D, de Oliveira SHP, et al. Mix-and-inject XFEL crystallography reveals gated conformational dynamics during enzyme catalysis. *Proc Natl Acad Sci U S A*. 2019;116:25634–25640.
- Ebrahim A, Moreno-Chicano T, Appleby MV, et al. Dose-resolved serial synchrotron and XFEL structures of radiation-sensitive metalloproteins. *IUCrJ*. 2019;6:543–551.
- Pearson AR, Mehrabi P. Serial synchrotron crystallography for time-resolved structural biology. *Curr Opin Struct Biol*. 2020;65:168–174.
- Ebrahim A, Appleby MV, Axford D, et al. Resolving polymorphs and radiation-driven effects in microcrystals using fixed-target serial synchrotron crystallography. *Acta Cryst D*. 2019;75:151–159.
- Casanas A, Warshamanage R, Finke AD, et al. EIGER detector: Application in macromolecular crystallography. *Acta Cryst D*. 2016;72:1036–1048.
- Callaway E. Revolutionary cryo-EM is taking over structural biology. *Nature*. 2020;578:201.
- Yip KM, Fischer N, Paknia E, Chari A, Stark H. Breaking the next Cryo-EM resolution barrier – Atomic resolution determination of proteins! *bioRxiv*. 2020;106740.
- Chirgadzhe DY, Murshudov G, Aricescu AR, Scheres S. Single-particle cryo-EM at atomic resolution. *bioRxiv* 2020. Available from: <https://www.biorxiv.org/content/10.1101/2020.05.22.110189v1.abstract>
- Li P-N, de Oliveira SHP, Wakatsuki S, van den Bedem H (2020) Sequence-guided protein structure determination using graph convolutional and recurrent networks. *arXiv*. 2020. Available from: <http://arxiv.org/abs/2007.06847>.
- Berman HM, Westbrook J, Feng Z, et al. The Protein Data Bank. *Nucleic Acids Res*. 2000;28:235–242.

25. Merritt EA. Expanding the model: Anisotropic displacement parameters in protein structure refinement. *Acta Cryst D*. 1999; 55:1109–1117.
26. Lovell SC, Word JM, Richardson JS, Richardson DC. The penultimate rotamer library. *Proteins*. 2000;40:389–408.
27. Burnley BT, Afonine PV, Adams PD, Gros P. Modelling dynamics in protein crystal structures by ensemble refinement. *eLife*. 2012;1:e00311.
28. Babcock NS, Keedy DA, Fraser JS, Sivak DA. Model selection for biological crystallography. *bioRxiv*. 2018. Available from: <https://www.biorxiv.org/content/10.1101/448795v1.abstract>
29. van den Bedem H, Dhanik A, Latombe JC, Deacon AM. Modeling discrete heterogeneity in X-ray diffraction data by fitting multi-conformers. *Acta Cryst D*. 2009;65:1107–1117.
30. Brock JS, Hamberg M, Balagunaseelan N, et al. A dynamic Asp-Arg interaction is essential for catalysis in microsomal prostaglandin E2 synthase. *Proc Natl Acad Sci U S A*. 2016;113:972–977.
31. Russi S, González A, Kenner LR, Keedy DA, Fraser JS, van den Bedem H. Conformational variation of proteins at room temperature is not dominated by radiation damage. *J Synchrotron Radiat*. 2017;24:73–82.
32. Atakisi H, Moreau DW, Thorne RE. Effects of protein-crystal hydration and temperature on side-chain conformational heterogeneity in monoclinic lysozyme crystals. *Acta Cryst D*. 2018; 74:264–278.
33. van Zundert GCP, Hudson BM, de Oliveira SHP, et al. qFit-ligand reveals widespread conformational heterogeneity of drug-like molecules in X-ray electron density maps. *J Med Chem*. 2018;61:11183–11198.
34. Liebschner D, Afonine PV, Baker ML, et al. Macromolecular structure determination using X-rays, neutrons and electrons: Recent developments in Phenix. *Acta Cryst D*. 2019;75: 861–877.
35. Barr AJ, Ugochukwu E, Lee WH, et al. Large-scale structural analysis of the classical human protein tyrosine phosphatome. *Cell*. 2009;136:352–363.
36. Keedy DA, Fraser JS, van den Bedem H. Exposing hidden alternative backbone conformations in X-ray crystallography using qFit. *PLoS Comput Biol*. 2015;11:e1004507.
37. Afonine PV, Poon BK, Read RJ, et al. Real-space refinement in PHENIX for cryo-EM and crystallography. *Acta Cryst D*. 2018; 74:531–544.
38. Su H, Zou Y, Chen G, et al. Exploration of fragment binding poses leading to efficient discovery of highly potent and orally effective inhibitors of FABP4 for anti-inflammation. *J Med Chem*. 2020;63:4090–4106.
39. Mobley DL, Dill KA. Binding of small-molecule ligands to proteins: “What you see” is not always “what you get.”. *Structure*. 2009;17:489–498.
40. Cappel D, Jerome S, Hessler G, Matter H. Impact of different automated binding pose generation approaches on relative binding free energy simulations. *J Chem Inf Model*. 2020;60: 1432–1444.
41. Kang YN, Stuckey JA. Crystal structure of the cdk2 in complex with thiazolylpyrimidine inhibitor. <https://doi.org/10.2210/pdb4EK8/pdb>.
42. Stsiapanava A, Olsson U, Wan M, et al. Binding of pro-Gly-pro at the active site of leukotriene A4 hydrolase/aminopeptidase and development of an epoxide hydrolase selective inhibitor. *Proc Natl Acad Sci U S A*. 2014;111:4227–4232.
43. Davis IW, Arendall WB 3rd, Richardson DC, Richardson JS. The backrub motion: How protein backbone shrugs when a sidechain dances. *Structure*. 2006;14:265–274.
44. van den Bedem H, Bhabha G, Yang K, Wright PE, Fraser JS. Automated identification of functional dynamic contact networks from X-ray crystallography. *Nat Methods*. 2013;10: 896–902.
45. Kiefer JR, Mao C, Braman JC, Beese LS. Visualizing DNA replication in a catalytically active *Bacillus* DNA polymerase crystal. *Nature*. 1998;391:304–307.
46. Smith CA, Kortemme T. Backrub-like backbone simulation recapitulates natural protein conformational variability and improves mutant side-chain prediction. *J Mol Biol*. 2008;380: 742–756.
47. Hallen MA, Keedy DA. Dead-end elimination with perturbations (DEEPer): A provable protein design algorithm with continuous sidechain and backbone flexibility. *Proteins*. 2013;81: 18–39.
48. van den Bedem H, Lotan I, Latombe JC. Real-space protein-model completion: An inverse-kinematics approach. *Acta Cryst D*. 2005;61:2–13.
49. Pearce NM, Krojer T, von Delft F. Proper modelling of ligand binding requires an ensemble of bound and unbound states. *Acta Cryst D*. 2017;73:256–266.
50. van Zundert GCP, Moriarty NW, Sobolev OV, Adams PD, Borrelli KW. Macromolecular refinement of X-ray and cryo-electron microscopy structures with Phenix/OPLS3e for improved structure and ligand quality. *bioRxiv*. 2020:198093.
51. Jiang L, Kuhlman B, Kortemme T, Baker D. A “solvated rotamer” approach to modeling water-mediated hydrogen bonds at protein–protein interfaces. *Proteins*. 2005;58(4): 893–904.
52. Liebschner D, Afonine PV, Moriarty NW, et al. Polder maps: Improving OMIT maps by excluding bulk solvent. *Acta Cryst D*. 2017;73:148–157.
53. Holton JM, Classen S, Frankel KA, Tainer JA. The R-factor gap in macromolecular crystallography: An untapped potential for insights on accurate structures. *FEBS J*. 2014;281: 4046–4060.
54. Terwilliger TC, Grosse-Kunstleve RW, Afonine PV, et al. Iterative-build OMIT maps: Map improvement by iterative model building and refinement without model bias. *Acta Cryst D*. 2008;64:515–524.
55. Lang PT, Holton JM, Fraser JS, Alber T. Protein structural ensembles are revealed by redefining X-ray electron density noise. *Proc Natl Acad Sci U S A*. 2014;111:237–242.
56. Volkman N. Confidence intervals for fitting of atomic models into low-resolution densities. *Acta Cryst D*. 2009;65:679–689.
57. Hornak V, Abel R, Okur A, Strockbine B, Roitberg A, Simmerling C. Comparison of multiple AMBER force fields and development of improved protein backbone parameters. *Proteins*. 2006;65:712–725.
58. Word JM, Lovell SC, Richardson JS, Richardson DC. Asparagine and glutamine: using hydrogen atom contacts in the

choice of side-chain amide orientation. *Journal of Molecular Biology*. 1999;285:1735–1747. <http://dx.doi.org/10.1006/jmbi.1998.2401>.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Riley BT, Wankowicz SA, de Oliveira SHP, et al. qFit 3: Protein and ligand multiconformer modeling for X-ray crystallographic and single-particle cryo-EM density maps. *Protein Science*. 2021;30:270–285. <https://doi.org/10.1002/pro.4001>