Gentle and fast all-atom model refinement to cryo-EM densities via Bayes' approach

Christian Blau, Linnea Yvonnesdotter, and Erik Lindahl

(Dated: September 30, 2022)

Abstract

Better detectors and automated data collection have generated a flood of high-resolution cryo-EM maps, which in turn has renewed interest in improving methods for determining structure models corresponding to these maps. However, automatically fitting atoms to densities becomes difficult as their resolution increases and the refinement potential has a vast number of local minima. In practice, the problem becomes even more complex when one also wants to achieve a balance between a good fit of atom positions to the map, while also establishing good stereochemistry or allowing protein secondary structure to change during fitting. Here, we present a solution to this challenge using Bayes' approach by formulating the problem as identifying the structure most likely to have produced the observed density map. This allows us to derive a new type of smooth refinement potential - based on relative entropy - in combination with a novel adaptive force scaling algorithm to allow balancing of force-field and density-based potentials. In a low-noise scenario, as expected from modern cryo-EM data, the Bayesian refinement potential outperforms alternatives, and the adaptive force scaling appears to also aid existing refinement potentials. The method is available as a component in the GROMACS molecular simulation toolkit.

I. INTRODUCTION

Cryo-electron microscopy (cryo-EM) has undergone a revolution the last few years due to better detectors, measurement techniques and algorithms[1], and the technique now allows for rapid reconstruction of biomolecular "density maps" at near-atomic resolution[2, 3]. These density maps describe interactions between the sample and an electron beam in real space. They (including similar ones derived from X-ray structure factors) provide the basis for reasoning in structural biology. In particular for cryo-EM, Bayesian statistics has revolutionised the reconstruction of the density maps from micrographs. This provides a framework to soundly combine prior assumptions about the three-dimensional density map model with the likelihood function that connects this model to the measured data, and determine the density map most likely to have generated the observed data instead of directly trying to solve the underdetermined inverse problem[4].

However, to understand the structure and function of biological macromolecules, merely having an overall electron density is typically not sufficient - it is also necessary to model

coordinates of individual atoms into the maps[5]. This enables understanding of e.g. binding site properties, interactions with lipids or other subunits, structural rearrangements between alternative conformations, and in particular it makes it possible to model structural dynamics on nanosecond to microsecond time scales via molecular dynamics (MD) simulation[6]. If the interaction descriptions (force fields) used in these simulations were perfect and one had access to infinite amounts of sampling, computational methods should be able to further improve the structure just by starting refinement from a rough initial density, but in practice both force fields and sampling have shortcomings. Nevertheless, it remains an attractive idea to combine the best of both worlds by using cryo-EM data as large-scale constraints while force fields are employed to fine-tune details - in particular details such as local stereogeometry or interactions on resolution scales that go beyond what the cryo-EM data can resolve. Cryo-EM data and stereochemical constraints have been combined favourably in the past to aid structure modelling into three-dimensional cryo-EM densities either by adding force field terms enforcing desired stereochemistry to established modelling tools[7–11] or by adding a heuristic density-based biasing potential to molecular dynamics simulations[12, 13] or elastic network models[14].

In practice it is not straightforward how to best combine experimental map data with simulations and achieve both satisfactory results and rapid convergence. Density-based biasing potentials can in principle achieve arbitrarily good fits to a map, but it comes at a cost of distorting the protein structure. To address the challenges of balancing desired stereochemical properties with cryo-EM data, refinement protocols have been expanded to include secondary structure restraints[12], multiple resolution ranges [11, 15], as well as multiple force constants[11, 16]; the latter two either consecutively in individual simulations[11] or via Hamiltonian replica exchange[15, 16]. A common challenge of all these approaches is the increase in ruggedness of the applied bias potential function as the resolution of the cryo-EM density maps increases, and how to correctly balance molecular mechanics and forces from the biasing potential. This leads to an apparent modelling paradox that further improving structural models for cryo-EM densities with molecular dynamics appears to be harder the more high-resolution data is available for these models.

To attack this challenge from a fundamental standpoint, Bayes approach has been used to derive probabilities for all-atom structural models given a cryo-EM density[17] and to weigh cryo-EM data influence against other sources of data[18]. These modelling approaches offer

valuable insight into the data content in cryo-EM maps and provide promising new ways to model cryo-EM densities by treating them as generic experimental data. However, they do not reflect the underlying physics of data acquisition and density reconstruction from micrographs and have previously not yielded refinement potentials of a new quality to be applicable, e.g., in molecular dynamics simulations.

One way of circumventing the number of model assumptions that are necessary to reflect the reconstruction of three-dimensional cryo-EM densities is to employ Bayesian models that directly connect micrographs and all-atom ensembles[19]. These attempts have previously proven to be prohibitively costly as a way to derive driving forces for molecular dynamic simulation, because projections of model atom coordinates onto millions of cryo-EM particle images (i.e., images of molecules) are required for a single force evaluation.

In this work, we show how it is possible to borrow the highly successful approach to density reconstruction and use Bayesian modeling of cryo-EM density maps from structures to derive a new biasing potential that is smooth, long-ranged, and provides fewer barriers to refinement than established potentials based on cross-correlation[11, 13] or inner product (equivalent to using potentials proportional to inverted electron density[12]). This provides a number of advantages, including an ability to overcome density barriers and in particular avoid excessive biasing forces resulting from large local gradients in electron density. It also avoids the need for constraints e.g. on secondary structure, and rather allows the simulation to freely explore local conformational space, while the experimental data is used to bias sampling to experimentally favored regions of the global conformational space.

We further demonstrate how better balancing between the force field and cryo-EM density components can be achieved by adaptive force scaling derived from thermodynamic principles. This allows refinement with a single fixed parameter at low computational cost for a range of system sizes and initial model qualities. Additionally, to evaluate biasing potentials based on model to cryo-EM density comparison, we suggest a transformation of all-atom structure to model density that reflects cryo-EM specific characteristics while minimising computational effort.

We investigate advantages and drawbacks of the newly derived potential in practical applications when compared to established inner-product and cross-correlation biasing potentials in a noise-free and experimental cryo-EM data. Finally, we show how the proposed refinement methods rectifies a distorted initial model with cryo-EM data. A full open source

implementation is freely available as part of the GROMACS molecular dynamics simulation code[20].

II. RESULTS

A. Deriving refinement forces via Bayes approach

Our canonical algorithm to refine all-atom models into a cryo-EM density map ρ with molecular dynamics is based on initially roughly aligning density map and target structure, generating a model density from coordinates $\vec{x}$, and then determining a dimensionless similarity score S between the generated model density and the target cryo-EM density (Fig. 1). The similarity is used to derive fitting forces, which are then combined with the force field potential U_{ff} based on a heuristic balance between the density-derived forces and force field determined by a force constant k. The combined driving forces are determined by the total potential energy,

$$U_{\mathrm{tot}}(\vec{x}, \rho) = U_{\mathrm{ff}}(\vec{x}) - kS(\rho, \vec{x})\,. \tag{1}$$

We find that applied density forces imply a similarity score between the model structure and target density, and vice-versa. Assuming that a single configuration of atoms gives rise to the observed cryo-EM density, Bayes' approach quantifies the probability density that the given model describes the cryo-EM data as[17]

$$p(\vec{x}|\rho) \propto p(\vec{x})p(\rho|\vec{x})\,. \tag{2}$$

Boltzmann inversion at temperature T connects the left-hand sides of Eqs. (1) and (2)[21] where c is an arbitrary potential energy offset

$$\log p(\vec{x}) = -\frac{1}{k_B T} U_{\mathrm{ff}}(\vec{x}) + c \tag{3}$$

$$\log p(\rho|\vec{x}) = \frac{k}{k_B T} S(\rho, \vec{x}) + c\,. \tag{4}$$

In this formalism, the force field provides the prior $p(\vec{x})$ that would have determined the model without any additional observations, while the similarity measure provides the conditional probability that a particular given structure yields a target density $p(\rho|\vec{x})$, scaled with force-constant k.

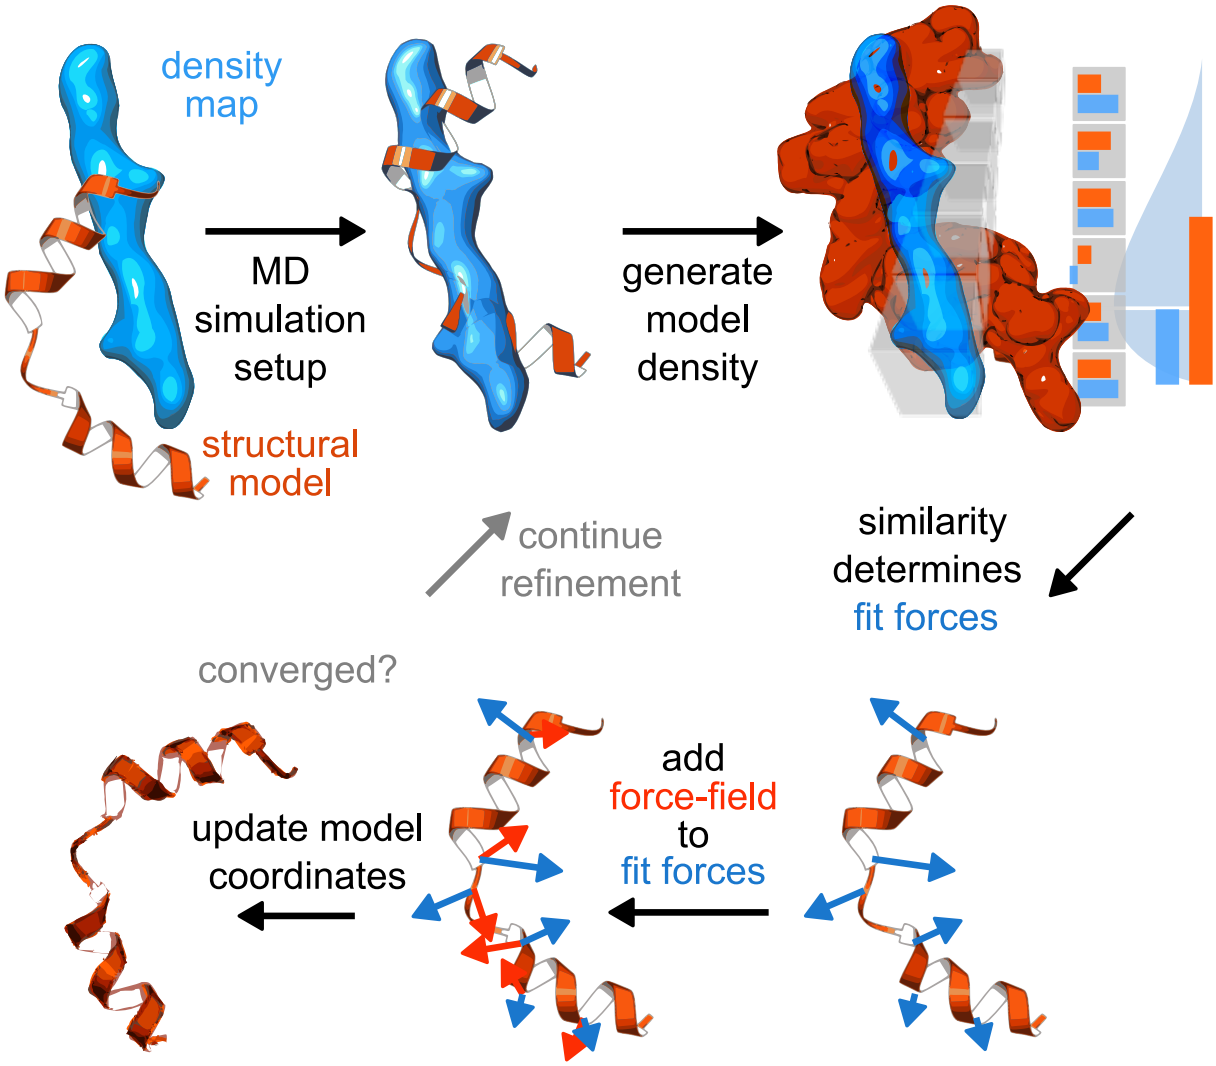

Figure 1. Atomic structure models are refined into a cryo-EM density using biasing forces that maximise similarity between model and map. A refinement/simulation is initialized with an atomic model (orange) and a density map (blue). A model density is generated in each voxel (grey boxes). Voxel-wise similarity scores between model density and cryo-EM density are akin to a noise model (light blue curve). The gradient of the similarity score determines the fitting forces (blue arrows). Together with a molecular dynamics force field (red arrows), the fitting forces enable model coordinate updates (dark orange) that make the model more similar to the density under force field constraints. New model densities are generated iteratively from the updated model in each time step of the simulation until acceptable convergence is reached.

Note that this does not assume any particular form of the similarity, but it provides a general relationship that also relates established similarity measures like cross-correlation or inner-product to their implicit assumptions about the likelihood function above, and in turn enables construction of new similarity scores that drive refinement procedures depending on the assumptions about the underlying measurement process.

B. Bayes approach yields negative relative entropy as the natural similarity score

To derive a new refinement potential from the likelihood of measuring a density given the structure, we assume that cryo-EM densities represent atom-electron scattering probabilities, where electron-atom interaction leads to a phase shift in electrons. With this assumption, two steps are necessary to calculate the density likelihood $p(\rho, \vec{x})$ from coordinates. First, an electron-scattering probability density ρ^s is created from a given structure. Second, this density is compared to a given measured density by the likelihood $p(\rho|\vec{x}) = p(\rho|\rho^s(\vec{x}))$. Using this result, two further assumptions enable the derivation of a new similarity score.

First, reported cryo-EM electron densities at each voxel are assumed to be proportional to the number of interactions of N incident electrons that each interact with probability ρ_v^s. This assumes that vitrious ice is not visible and contributing to the scattering, which is commonly achieved by shifting the offset of cryo-EM densities so that water density is represented with voxel values that fluctuate around zero. Only accounting for positive density, we describe this scattering interaction process by a Poisson distribution with parameter $\lambda = N\rho_v^s$. While it is theoretically possible to expand the model to include noise fluctuations and negative densities, we omit this for the sake of reducing model complexity. Second, we assume that measured scattering probabilities per voxel are independent of other voxel values. This does not exclude spatial correlation between density data, but states that the *scattering process* in one voxel does not influence the electron interaction in other voxels.

With these assumptions (the detailed algebraic transformations are laid out in the Supporting Information S1 Appendix), we obtain a similarity score between simulated model density and cryo-EM density proportional to the negative relative entropy, or Kullback-Leibler divergence,

$$\frac{k}{k_B T} S(\rho, \rho^s) = \sum_{v \in \text{voxels}} \rho_v \left(\log \frac{\rho_v^s}{\rho_v} \right) = -D_{\text{KL}}(\rho \| \rho_v^s) \,. \tag{5}$$

Actual cryo-EM micrographs are typically normalized to unity variance around the particle region[22]. As a consequence, cryo-EM densities are scaled by a free parameter, and they may thus be rescaled as $\hat{\rho}_v = r\rho_v$. This results in an unknown scaling in the force constant, which can seen from so-called re-gauging with an additional constant $k_B T \sum_{v \in \text{voxels}} \log(\hat{\rho}_v / r)$,

$$\frac{r \cdot k}{k_B T} S(\rho, \rho^s) = \sum_{v \in \text{voxels}} \hat{\rho}_v \left(\log \frac{\rho_v^s}{\hat{\rho}_v} \right) . \tag{6}$$

To ensure that the cryo-EM density has the properties of a probability distribution, we choose our arbitrary re-gauging such that $\sum_v \hat{\rho}_v = 1$. The force constant balancing cryo-EM data vs. force field / stereogeometry is still a free parameter with these model assumptions and is chosen adaptively with a protocol described below.

The newly derived relative-entropy-based similarity score has a domain of $[-\infty, 0]$ with perfect agreement at zero. Due to the $\log \rho_v^s$ term it differs prominently from established similarity scores like cross-correlation[11] and inner-product (formulated as force following the gradient of a smoothed inverted density which are equivalent in this approach[12]; see Supporting Information S1 Appendix). In contrast, the relative-entropy based score receives the largest contribution from voxels where cryo-EM data has no corresponding model density data.

This leads to a substantially different behaviour from established similarity scores with local minima for locally good agreement with cryo-EM data while the relative-entropy based potential will only have minima where there is good *global* agreement between structure and density. As a consequence, the relative-entropy based density potential is expected to perform better in situations where other potentials cannot escape local minima, at the cost of higher sensitivity to noise in the data.

C. The potential energy landscapes based on relative entropy are smooth

The proposed relative entropy density-to-density similarity measure has favourable properties in one-dimensional model refinement of one and two particles to a reference density (Fig. 2).

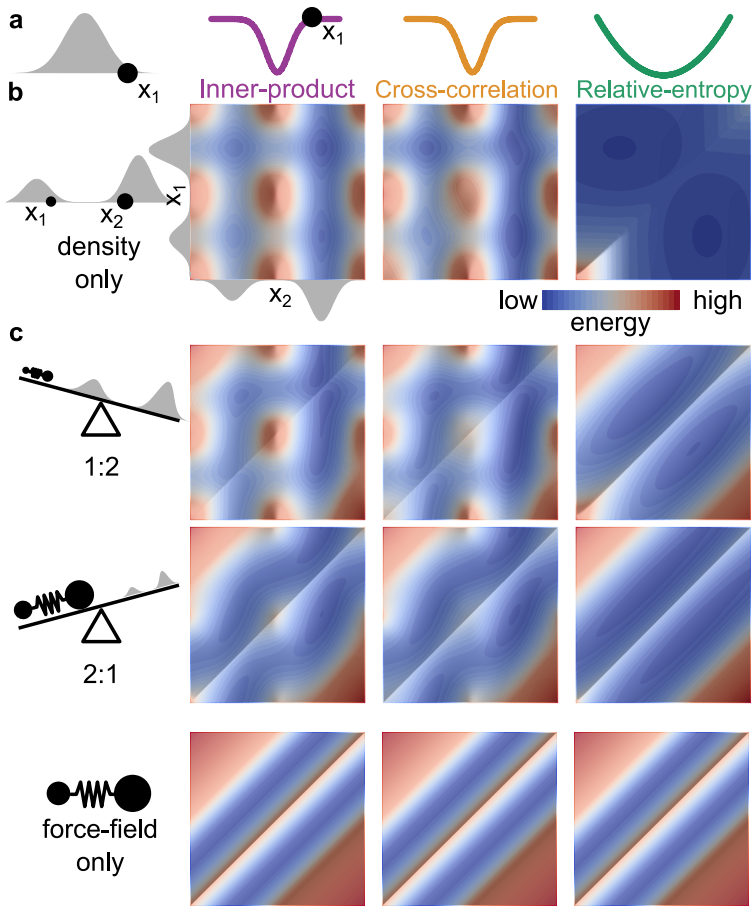

Figure 2. Similarity score determines ruggedness of the effective refinement potential energy landscape, also when balancing it with structural bias. From top to bottom: **a** One-dimensional refinement of a single particle (black circle) towards a Gaussian-shaped density (gray) with inner-product (purple), cross-correlation (ochre), and relative-entropy (green) as similarity scores. **b** Expanded model with two particles (black circles, x_1 smaller and x_2 larger) with two amplitude peaks in a one-dimensional density and target distribution (gray), and the resulting two-dimensional effective potential energy landscapes for inner-product (left), cross-correlation (middle), and relative-entropy similarity measures (right). **c** Combination of the similarity measure and force field contribution to the potential energy landscape, exemplified by a harmonic bond that keeps particles at half the distance between the Gaussian centers. For all relative weights of the contributions of the refinement potential and bond potential energy landscape (ratio 1:2 upper panel, 2:1 middle panel, as illustrated by the scale on the left), the relative entropy similarity score produces smooth landscapes single correct minima.

In contrast to cross-correlation and inner product similarity measures that have a steep and sudden onset for refinement forces in one dimension, the relative entropy similarity score has a harmonic shape with long-ranged interactions that allow for efficient minimisation. Using relative-entropy, the particle to be refined is attracted by a harmonic spring-like potential to the best fitting position; forces are large far away from the minimum, but their magnitude decreases monotonically as the minimum is approached. Inner-product and cross-correlation based fitting potentials, however, exert almost no force on the particle outside the Gaussian spread width, while exerting a suddenly increasing force when moving closer to the Gaussian center, and are only insignificant very close to the minimum.

This advantage for the relative-entropy-based potential is maintained for refinement of two particles, where the relative-entropy-based potential energy landscape is less rugged and has fewer pronounced features and minima than the corresponding landscapes for the inner-product and cross-correlation based potentials (Fig. 2). Only a single diagonal barrier is found in the relative-entropy-based potential landscape, corresponding to a "swapping" of particle positions, which alleviates the search for a global minimum. The inner-product-based free energy landscape has its minimum at a configuration where both particles are at the same position at the highest density. This issue can alleviated in practical applications through a force-field prior that would enforce a minimum distance between the atoms (e.g. through van der Waals interactions).

To model the influence of a force field, the two particles were connected with a harmonic bond with increasing influence. The balance between density-based forces and bond strongly determines the shape of the resulting energy landscape, but here too relative entropy provides a smoother landscape less sensitive to the specific relative weight of refinement and bond potentials.

D. Adaptive force scaling reduces work exerted during refinement and allows for comparison of density-based potentials

The force constant for density-guided simulations cannot be derived from the cryo-EM density alone and thus needs to be set heuristically. Established protocols where the force constant has to be determined manually require an iterative trial-and-error approach. We address this by introducing an adaptive force-scaling as depicted in Fig. 3a to automatically

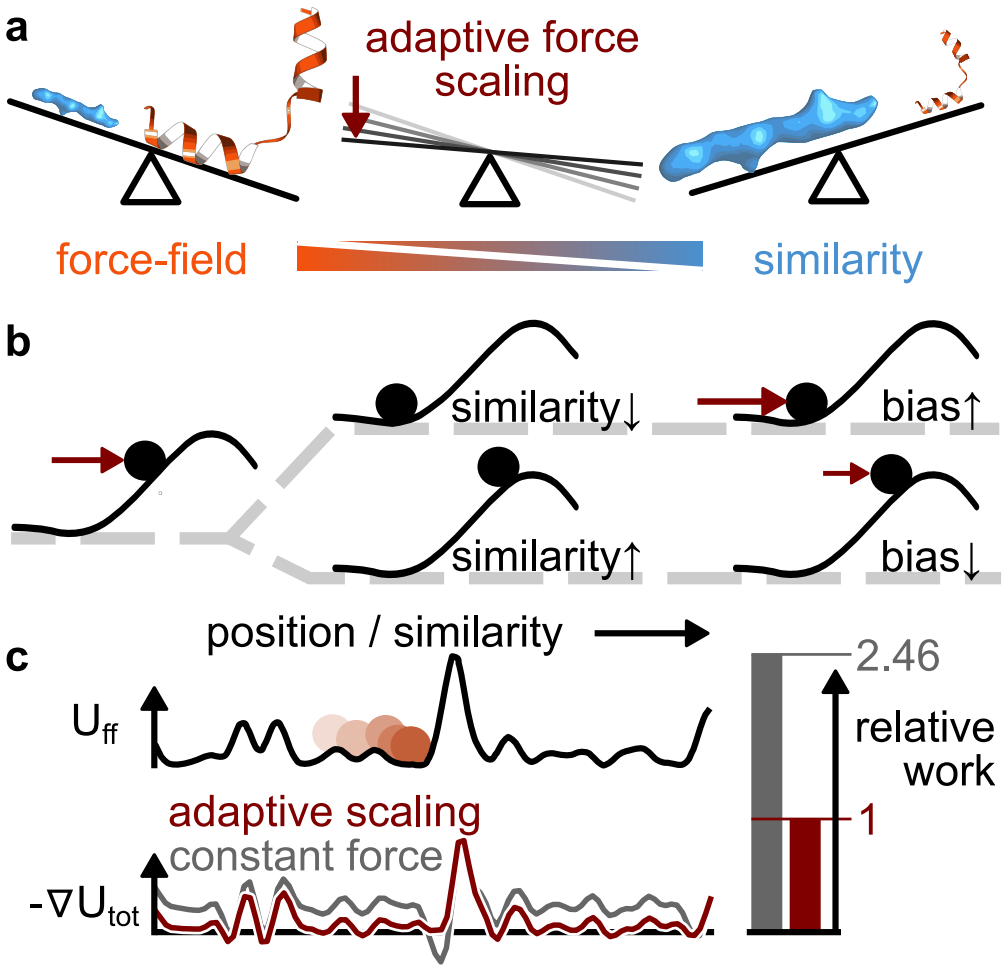

Figure 3. Adaptive scaling of contributions from force-field and cryo-EM density data overcomes potential energy barriers without excessive work input. **a** Adaptive force scaling heuristically balances force-field and density influence during refinement simulations. **b** Particle in energy landscape where density similarity increases from left to right along the black curve. For the upper leg alternative, the similarity decreases despite biasing forces (burgundy arrow), which causes the bias strength to be increased. Conversely, in a scenario where the similarity remains high (lower leg), the biasing force will gradually be reduced to allow the system to better sample the local landscape. **c** Brownian diffusion in a potential with fixed (grey) and adaptive (burgundy) biasing forces, respectively. The constant biasing force is scaled such that both force adding schemes yield the same average mean first passage time moving from left to right. The relative-entropy approach leads to significantly lower exerted work on the system (area under the grey and burgundy curves, respectively), which reduces perturbation of the dynamics of the system.

balance force-field and density-based forces during the refinement.

Cryo-EM refinement simulations are non-equilibrium simulations with the aim to drive a system from an initial model state to a final state that is as similar to the cryo-EM density as possible while avoiding structural distortions that result e.g. from unphysical paths. To avoid or at least reduce the latter during refinement, a heuristic protocol has been devised that aims to minimize work exerted on the system while still requiring as little time as possible for the refinement. To minimize the exerted work formulated as

$$W_{\text{fit}} = \int_{S=S_{\text{start}}}^{S_{\text{end}}} \partial_S k S \, \mathrm{d}S \approx \sum_{\text{frame}} k_{\text{frame}} \Delta S_{\text{frame}} \,, \tag{7}$$

the adaptive-force scaling starts from a low force constant k. This is then increased if similarity decreases, and conversely decreased if the similarity is increasing (Fig. 3b). Any feedback protocol of this type is guaranteed exert less work to reach the same similarity score than keeping the force constant fixed at the final value of the adaptive scaling protocol, given that the score is monotonically increasing.

A one-dimensional Brownian diffusion model system (Fig. 3c) is used to test the performance of the concrete scaling protocol as described in the methods section of this paper. In this model, the similarity score simply increases with increasing particle coordinate value. Biasing the system towards increasing coordinate values with adaptive force scaling in contrast to a constant force allows for the particle to reach a given coordinate value in the same average first-passage time at much lower average work input. Without any coupling of the free energy landscape to the adaptive force scaling protocol other than through the particle trajectory, the adaptive force scaling increases the force just sufficiently to allow overcoming energy barriers, but then reduces it again.

Adaptive force scaling further enables comparison of relative entropy to other established density-based potentials in simulations with cryo-EM data, because it disentangles the effect of the force constant choice from the choice of refinement potential. To carry out this comparison on cryo-EM data with our newly derived similarity score within the Bayesian framework, a model density generation protocol is required which is shown below.

E. Deriving an optimal model density generation for cryo-EM data refinement

To evaluate similarities of structural models to cryo-EM densities, a model electron scattering probability density is generated from atom positions. Two dominant effects are convoluted when modeling electron scattering probabilities: The scattering cross section of each atom and their thermal motion. Both are approximated with Gaussian functions of amplitude A and width σ. While the scattering cross sections determines A (Appendix C in Ref. [23]), the magnitude of thermal fluctuation of atoms at cryogenic temperatures determines the spread width σ.

However, these limits to the model resolution are superseeded by the finite performance of the measurement instrument and the reconstruction process where structural heterogeneity, detector pixel size, microscope lenses, and particle alignment limit the resolution. Here, structural heterogeneity is not accounted for, because it is an ensemble effect. A connection between the approach presented here and an ensemble model may be made though by employing a probability distribution $p(\vec{x})$ instead of $\vec{x}$ in Eq. (2) and leveraging ensemble simulations[24]. Other resolution-limiting effects are approximated by additional convolution of the generated maps with a Gaussian kernel. Rather than aiming to reproduce the same blur as in the experimental map, we strive to preserve as much information as possible from the physical model.

A balance between information loss due to under-sampling on the grid on the one hand and information loss due to coarse blurring is found where the Gaussian width at half maximum height equals the resolution. The maximum representable resolution on a grid corresponds to twice the Nyqvist frequency δ (corresponding to the pixel and voxel size) so that the Gaussian width σ is approximated in refinement simulations from the highest local resolution or, where that data is not applicable, from the voxel-size,

$$\sigma = \frac{\mathrm{res}_{\mathrm{max}}}{2\sqrt{\log 4}} \approx \frac{2\delta}{2\sqrt{\log 4}}. \tag{8}$$

For computational efficiency Gaussian spreading is truncated at 4σ for all simulations in this publication, accounting for more than 99.8% of the density (Fig. 1 in Supporting Information S1 Appendix). The small limitation on the maximal distance between initial model structure and the cryo-EM density through this cutoff has proven to be irrelevant for all practical purposes, as density-based forces will "pull" structures into densities as

soon as there is minimal overlap between model density and cryo-EM density, which can easily be achieved with an initial alignment. Interestingly, this approach results in a smaller Gaussian spreading width than previously applied ones that aim to reproduce a density map with the same overall resolution as the experimental cryo-EM density. As a result, it maintains as much structural information as possible in the model density while still reducing computational cost.

F. Refinement against noise-free data

To separate additional noise effects in experimental data and possible limitations in the above model, we first assess refinement with ideal data where a small straight helix model system[14] has been refined against a synthetically generated target density of the same helix in a kinked configuration. As illustrated in Fig. 4a, adaptive force scaling and relative-entropy as similarity score efficiently fit the helix into the synthetic cryo-EM density[25].

To follow the balancing of density-based forces and the force-field, the stereochemical distortion of the helix was assessed via the force-field potential energy against the unweighted FSC average[26]. This serves as an established similarity score that is not related to the biasing potentials which were used to refine the helix (Fig. 4b). With low biasing force, no structural distortion is seen while there is also no improvement of fit to the target density. With high forces from the refinement potential, the helix structure fits the density very well - but at the cost of larger local structural deviations. The less distortion necessary on average to achieve given FSC average, the better the performance of the refinement potential.

The time-averaged force-field potential energy was found to give a good indication of structural distortion. Since the reference structure will not be known for actual applications, this provides a practical way to determine how far to push refinement in practice. In this work we have chosen this threshold as the level where the influence from the density bias on the simulation excees three standard deviations of the averaged potential energy in the same simulation without additional density-based forces, but this choice is admittedly somewhat arbitrary and can no doubt be improved.

The combination of adaptive force scaling and different similarity scores consistently worked well when the helix was aligned to the density, despite small fluctuations of the results (Fig. 3c) due to the stochastic nature of molecular dynamics simulations. When the helix

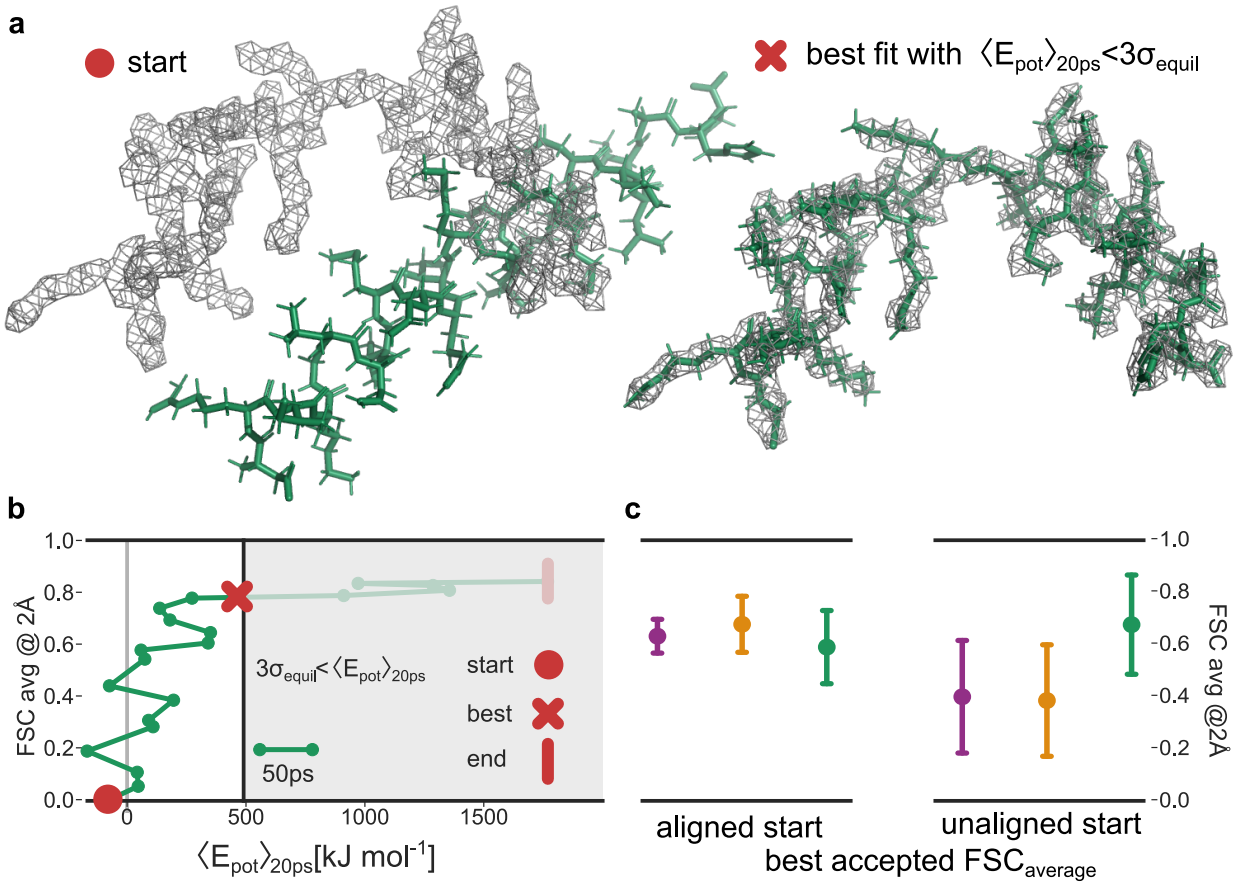

Figure 4. Refinement into noise-free data with adaptive force scaling. **a** Starting (left) and final (right) conformations (green sticks) of a helix subject to refinement simulation into a synthetically generated cryo-EM density (gray mesh) using relative-entropy similarity score and adaptive force scaling. The refined model structure is selected as the one with highest $\mathrm{FSC}_{\mathrm{avg}}$ within a 3σ excess deviation from the average potential energy of an equilibration simulation. **b** Change in force field potential energy and quality of fit as measured by the unweighted FSC average. In the course of the simulation (green line starting from circle to end at vertical line), the fit improves at the cost of larger potential energy averages, eventually exceeding (cross) 3σ of the system equilibration simulation (grey area). **c** Average and standard deviation of the highest FSC average within the 3σ potential energy criterion starting from aligned and unaligned positions deriving density fitting forces from inner-product (purple), cross-correlation (ochre) and relative-entropy (green) similarity measures (n=7).

was initially not aligned, the relative-entropy based potential shows markedly better results for the refinement. While the inner-product and cross-correlation based potentials in some instances fail to align the helix and get stuck in local minima that provide faulty fits(Fig. 4 in Supporting Information S1 Appendix), the relative-entropy measure consistently places and bends the model helix correctly. For a single helix this is a slightly artificial case, but in a large structure undergoing significant transitions it will be common for some secondary structure elements to not overlap with the target density.

G. Refinement against experimental cryo-EM data

By using adaptive-force scaling refinement of a previously published X-Ray structure of rabbit-muscle-aldolase[27] against a recently published independently determined cryo-EM structure[28], we consistently achieve accurate refinement using experimental cryo-EM data that contains noise that cannot be fully accounted for by the model assumptions. Figure 5a shows the final models of refinement against one half-map with a global deviation of less than 1Å heavy-atom root mean square deviation (RMSD) from the deposited manually build model using inner product and cross-correlation measures and slightly above 1Å for the relative-entropy based density potential (Tab. 1 in Supporting Information S1 Appendix).

The close agreement with the cryo-EM data is reflected in the FSC of the models refined against the density (Fig. 5b) being nearly indistinguishable from the manually built model. The relative-entropy-based potential, while still providing good agreement to the cryo-EM density, somewhat emphasises agreement with global features at the cost of local resolution (Fig. 5 in Supporting Information S1 Appendix).

The three-sigma equilibration energy shows that refinement reached almost saturated agreement with the density before the need for larger structural rearrangements, indicating balance between force-field and cryo-EM data influence (Fig. 5c). In this case it is possible to continue refinement which further improves the fit while only marginally affecting the model statistics (Tab. 3 in Supporting Information S1 Appendix). This could indicate our threshold is too conservative, but we prefer to err on the side of not introducing too much distortion. Convergence to this threshold was achieved in less than 3 ns for all underlying potentials as shown in Fig. 5c. The less rugged and long-range potential properties of relative-entropy based density forces are reflected in a rapid rigid-body like initial fit to global structural

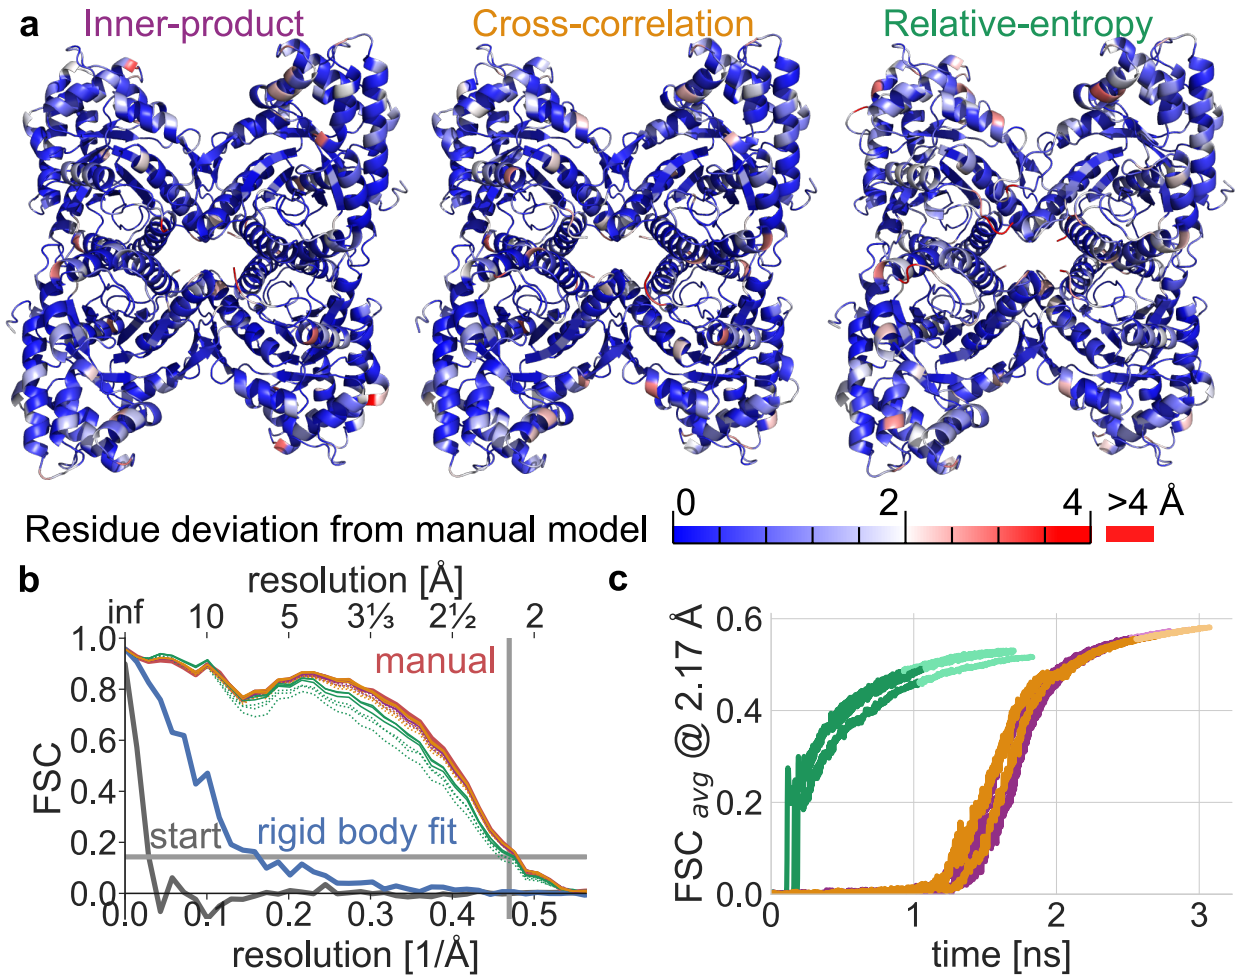

Figure 5. Refinement of an all-atom X-Ray structure (PDB id 6ALD) into experimental cryo-EM density (EMD-21023:halfmap1). **a** Final structure models from density-guided simulations using different similarity scores colored by unaligned root mean square coordinate deviation (RMSD) per residue from the manually built model (PDB id 6V20). **b** Fourier shell correlation of starting structure (gray line), rigid-body fit of the starting model to the target density (blue) as well as refinement results in the last simulation frame (solid lines) and the best fitting frame with a potential energy running average below 3σ deviations of an equilibrium simulation (dotted lines). The reported cryo-EM map resolution and 0.143 FSC are indicated with grey lines. **c** Unweighted FSC average over the course of refinement simulation. Light colors indicate crossing the 3σ running average potential energy threshold.

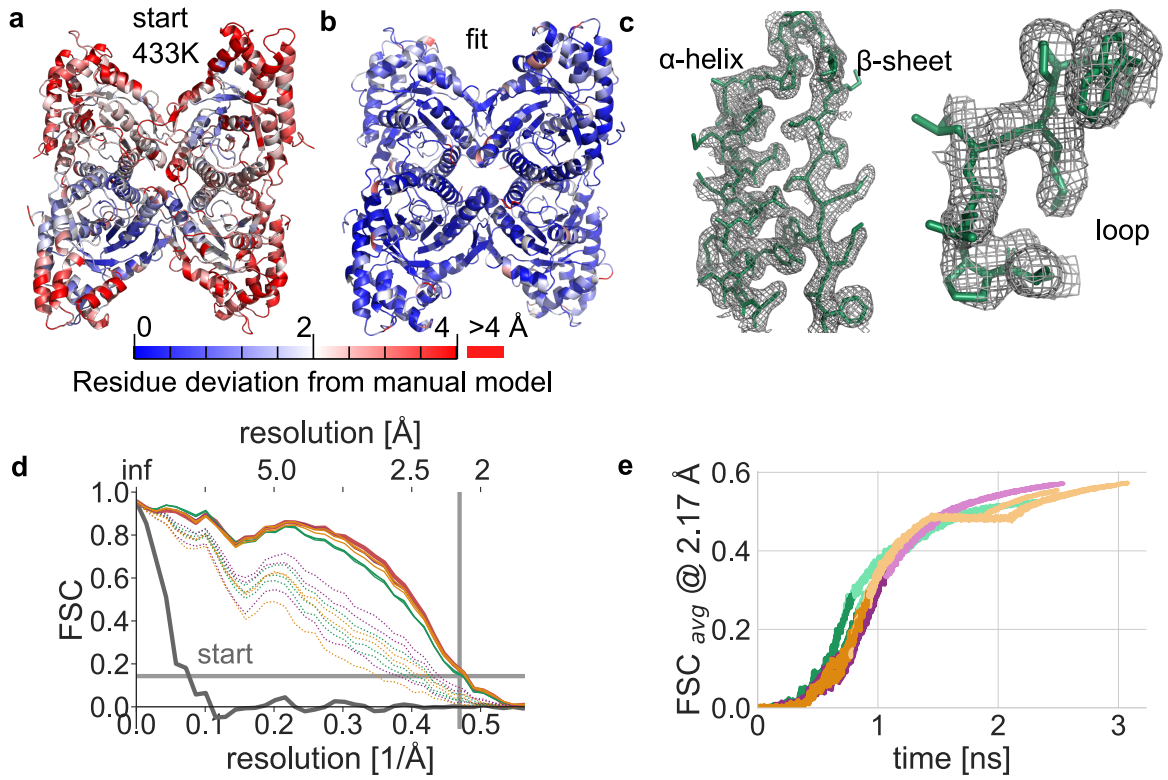

Figure 6. Cryo-EM data rectifies model distortions with density guided simulations. **a** Distorted starting model RSDM with respect to manually built model (PDB id 6V20).**b** Final model structure after refinement into a cryo-EM density (EMD-21023:halfmap1) using adaptive force scaling and relative-entropy similarity score. **c** Close-up of structural features of the final simulation model (green lines) and cryo-EM density (gray mesh). **d** Fourier shell correlation of starting structure (gray line) as well as refinement results in the last simulation frame (solid lines) and the best fitting frame with a potential energy running average below 3σ deviations of an equilibrium simulation (dotted lines). The reported cryo-EM map resolution and 0.143 FSC value are indicated with grey lines. **e** Un-weighted FSC average over the course of a refinement simulation. Light colors indicate exceeding the 3σ running average potential energy criterion.

features, while the other potentials show gradual improvements in fit.

H. Model rectification by combining force-field and cryo-EM data

To further assess the performance, we repeat the refinement when starting from initial model structures that have been distorted by heating with partially unfolded secondary

structure elements (Fig. 6, as described in Methods). Figure 6b shows the final relative-entropy based model of the refinement procedure that achieved 1.13 Å heavy-atom RMSD from the manually built model. Structural details at map resolution match in secondary structure elements. In contrast to refinement of the undistorted X-ray structure, the relative-entropy based potential gains less from the long-rangedness of the potential and the rapid alignment of large-scale features, because structural rearrangements were needed on all length-scales. The adaptive force scaling protocol alleviates differences between density based potentials in refinement speed and allows for refinement with good structural agreement in less than 3 ns.

Here, the three-sigma potential energy criterion works as an indicator that an imbalance between force-field contributions and cryo-EM data is needed to reach high FSC (Fig. 6d). The adaptive force scaling protocol allows the modeling to be more steered by cryo-EM data and reach model structures that would not have been accessible by modeling using stereochemical information from the force-field alone.

III. DISCUSSION

While defining a purely empirical similarity measure can sometimes suffice to fit structures to cryo-EM densities, connecting the similarity measure to the underlying measurement process of the target density enables derivation of a natural similarity measure. From very few assumptions, this results in the density-based potential derived via Bayes' approach that coincides with the relative-entropy between a model density generated from model/simulation atom coordinates and the target cryo-EM density.

The newly defined potential has favourable features like long-rangedness and low ruggedness, which avoids false local minima during refinement and allows rapid alignment of large-scale features. It performs superior to established refinement in the zero noise setting with synthetic density maps. The noise-content in current cryo-EM densities is likely still too high to be handled with the current minimalistic model assumptions, but as the quality of cryo-EM and other low-to-medium resolution techniques continues to rapidly improve, we believe there will be even more advantages to models that do not depend on smoothing. In addition, the adaptive force-scaling provides a surpisingly simple way to tackle the inaccessibility of the balance between force-field and density based forces within our model assumptions. It

allows for parameter-free refinement between one and two orders of magnitude faster than currently established protocols.

Another illustration of the usefulness of the Bayesian framework for handling force-field vs. fitting forces is how it enables us to deduce a close-to-optimal model density spread for refinement, and even more so that this value is not identical to the common practice of setting it equal to the experimental resolution. While many of these factors could still be tuned manually, removing them as free parameters means fewer arbitrary settings that avoids over- or underfitting, which will be even more important when trying to combine e.g. multiple sources of experimental data. For trial structure refinement against recent cryo-EM data, we show that we achieve excellent fits independent of initial model quality. For larger structural changes, structures still have to be allowed larger conformational changes at the cost of higher potential energies. It is an interesting question whether this could be explained by limited sampling, in which case the relative-entropy potential that is considerably softer close to the minimum might make it possible to combine refinement with ensemble simulation techniques such as Markov state models (MSM).

One limitation of the current formulation is that it does not explicitly take the local resolution information from the cryo-EM map into account. In practice the local resolution will still influence the fit since low local resolution will correspond to smoother regions of the map, and lower-magnitude gradients will lead to lower-magnitude fitting forces in those regions. However, a formally more correct way to address the problem is likely to treat the target electron density as a statistical distribution with a variance that is spatially resolved - this is something we intend to pursue in the future to see whether it can further improve local fits.

The algorithms proposed in this work are freely available, integrated and maintained as part of GROMACS[29]. Overall, three independent building blocks are provided to aid modeling of cryo-EM data that each may be individually implemented in current modelling tools: A new refinement potential, a new criterion for how to calculate the model density, both based on reasoning via Bayes' approach, and adaptive force scaling to gently and automatically bias stereochemistry and cryo-em data influence. The implementation also provides tools to monitor the refinement process. Although it can still be difficult for any automated method to compete with manual model building by an experienced structural biologist, we believe these methods provide new ways to extract as as much structural

information as possible from cryo-EM densities at minimal human and computational cost, which is particularly attractive e.g. for fully automated model building.

IV. METHODS

A. Calculating density-based forces

For ease of implementation and computational efficiency the derivative of Eq. (4) is decomposed into a similarity measure derivative and a simulated density model derivative, summed over all density voxels v

$$\mathbf{F}_{\text{density}} = k \sum_{v} \partial_{\rho_v^{\text{s}}} S(\rho, \rho^{\text{s}}) \cdot \nabla_{\mathbf{r}} \rho_v^{\text{s}}(\mathbf{r}) \,. \tag{9}$$

Though the convolution in eq. (9) might be evaluated with possible performance benefits in Fourier space, the more straightforward real-space approach has been implemented.

The forward model ρ^s is calculated using fast Gaussian spreading as used in [30], the integral over the three-dimensional Gaussian function over a voxel is approximated by its function value at the voxel center $\mathbf{v}$ at little information loss (Fig. 2 in Supporting Information S1 Appendix). Amplitudes of the Gaussian functions[23] have been approximated with unity for all atoms except hydrogen. The explicit terms that follow for $S(\rho, \rho^{\text{s}})$ and $\nabla_{\mathbf{r}} \rho_v^{\text{s}}(\mathbf{r})$ are stated in the Supporting Information S1 Appendix.

B. Multiple time-stepping for density-based forces

For computational efficiency, density-based forces are applied only every N_{fit} steps. The applied force is scaled by N_{fit} to approximate the same effective Hamiltonian as when applying the forces every step, while maintaining time-reversibility and energy conservation[31, 32]. The maximal time-step should not exceed the fastest oscillation period of any atom within the map potential divided by π. This oscillation period depends on the choice of reference density, the similarity measure and the force constant and has thus been estimated heuristically.

C. Adaptive force scaling

Adaptive force constant scaling decreases the force-constant when similarity increases by a factor $1 + \alpha$, with $\alpha > 0$ and reversely increases it by a factor $1 + 2\alpha$ when similarity decreases. The larger increase than decrease factor enforces an increase in similarity over time.

To avoid spurious fast changes in force-constant, similarity decrease and increase is determined by comparing similarity scores of an exponential moving average. The simulation time scale is coupled to the adaptive force scaling protocol by setting $\alpha = \frac{N_{\text{fit}}\Delta t}{\tau}$, where Δt is the smallest time increment step of the simulation and τ determines the time-scale of the coupling.

This adaptive force scaling protocol ensures a growing influence of the density data in the course of the simulation, eventually dominating over the force-field. Simulations with adaptive force scaling are ended when overall forces on the system are too large to be compatible with the integration time-step.

D. Comparing refined structures to manually built models

Root mean square deviations (RMSD) of all heavy atom coordinates (excluding hydrogen atoms) used absolute positions without super-position as structures, because the cryo-EM density provides the absolute frame of reference. This is an upper bound to RMSD values between refined and manually built models calculated with rotational and translational alignment.

E. Comparing refined structures to cryo-EM densities

Fourier shell correlation curves and un-weighted Fourier shell correlation averages[26] have been calculated at 4 ps intervals from structures during the trajectories by generating densities from the model structures using a Gaussian σ of 0.45 Å, corresponding to a resolution as defined in EMAN2[25] to 2Å.

F. Map and model preparation before refinement

Noise-free density maps at 2Å simulated resolution on a 1 Å voxel grid have been generated from atomic models using "molmap" as provided by chimera[33].

Non-pertinent parts of half-map one deposited at EMD-21023 were discarded, resulting in a 192^3 map. The starting models were roughly aligned after visual inspection.

G. Generation of a distorted model

To generate a distorted starting model, the aldolase protein X-ray was heated to $433\frac{1}{3}$ K over a period of 4 ns. During heating, pressure was controlled with the Berendsen barostat, favoring simulation stability over thermodynamic considerations. To disentangle effects from decreasing the temperature and fitting, the distorted structure was subjected to 5 ns of equilibration at 300 K before starting the density guided simulations with the same protocol as described above.

H. Molecular dynamics simulation

All simulations were carried out with GROMACS2020.2[29] and the CHARMM27 force-field[34, 35] in an NPT and NVT ensembles with neutralized all-atom systems in 150 mMol NaCl solution. Temperature was regulated with the velocity-rescaling thermostat at a coupling frequency of 0 ps to ensure rapid dissipation of excess energy from density based potential, when structures are very dissimilar from the cryo-EM density, i.e., far from equilibrium. Pressure was controlled with the Parinello-Rahman barostat for aldolase simulations, helix simulations were carried out at constant volume.

Aldolase was roughly aligned to the density by placing its center of geometry in the center of the cryo-EM density box.

Forces from density guided simulations are applied every $N_{\text{fit}} = 10$ steps according to the protocol described above[32]. All simulations with adaptive force-scaling used a coupling constant of $\tau = 4\,\text{ps}$, balancing time to result with time for structures to relax.

The Gaussian spread width for aldolase simulations has been determined using a lower bound on the highest estimated local resolution of 1.83 Å.

Periodic boundary conditions are treated as described in the Supporting Information S1 Appendix. All simulation setup parameters and workflows have been made available.

V. DATA AVAILABILITY

Simulation starting structures, generated densities, setup parameters, complete workflow setups via Makefiles and Python scripts to generate Fig. 2 as well as data for Fig. 3 and per-residue RMSD are available via Zenodo (https://doi.org/10.5281/zenodo.4556616). The code to perform density-guided molecular dynamics simulations is maintained within GROMACS and publicly available in release 2021 and later, as well as in the repository at https://http://gitlab.com/gromacs/gromacs. Fourier shell correlation analysis of trajectories has been implemented on top of the GROMACS codebase following conventions in EMAN2[25] and is available at https://gitlab.com/gromacs/gromacs/-/commits/fscavg. Python scripts to generate Fig. 2, as well as data for Fig. 3 and per-residue RMSD are available via Zenodo (https://doi.org/10.5281/zenodo.4556616).

VI. REFERENCES

[1] Cheng, Y. Single-particle cryo-em—how did it get here and where will it go. *Science* **361**, 876–880 (2018).

[2] Yip, K. M., Fischer, N., Paknia, E., Chari, A. & Stark, H. Atomic-resolution protein structure determination by cryo-em. *Nature* **587**, 157–161 (2020).

[3] Nakane, T. *et al.* Single-particle cryo-em at atomic resolution. *Nature* **587**, 152–156 (2020).

[4] Scheres, S. H. Relion: implementation of a bayesian approach to cryo-em structure determination. *Journal of structural biology* **180**, 519–530 (2012).

[5] Alnabati, E. & Kihara, D. Advances in structure modeling methods for cryo-electron microscopy maps. *Molecules* **25**, 82 (2020).

[6] Bock, L. V. *et al.* Energy barriers and driving forces in trna translocation through the ribosome. *Nature structural & molecular biology* **20**, 1390–1396 (2013).

[7] Eswar, N. *et al.* Comparative protein structure modeling using modeller. *Current protocols in bioinformatics* **15**, 5–6 (2006).

[8] Topf, M. *et al.* Protein structure fitting and refinement guided by cryo-em density. *Structure* **16**, 295–307 (2008).

[9] Kim, D. N. *et al.* Cryo_fit: Democratization of flexible fitting for cryo-em. *Journal of structural biology* **208**, 1–6 (2019).

[10] Liebschner, D. *et al.* Macromolecular structure determination using x-rays, neutrons and electrons: recent developments in phenix. *Acta Crystallographica Section D: Structural Biology* **75**, 861–877 (2019).

[11] Igaev, M., Kutzner, C., Bock, L. V., Vaiana, A. C. & Grubmüller, H. Automated cryo-em structure refinement using correlation-driven molecular dynamics. *Elife* **8**, e43542 (2019).

[12] Trabuco, L. G., Villa, E., Mitra, K., Frank, J. & Schulten, K. Flexible fitting of atomic structures into electron microscopy maps using molecular dynamics. *Structure* **16**, 673–683 (2008).

[13] Ahmed, A., Whitford, P. C., Sanbonmatsu, K. Y. & Tama, F. Consensus among flexible fitting approaches improves the interpretation of cryo-em data. *Journal of structural biology* **177**, 561–570 (2012).

[14] Wang, Z. & Schröder, G. F. Real-space refinement with direx: From global fitting to side-chain improvements. *Biopolymers* **97**, 687–697 (2012).

[15] Singharoy, A. *et al.* Molecular dynamics-based refinement and validation for sub-5 å cryo-electron microscopy maps. *Elife* **5**, e16105 (2016).

[16] Miyashita, O., Kobayashi, C., Mori, T., Sugita, Y. & Tama, F. Flexible fitting to cryo-em density map using ensemble molecular dynamics simulations. *Journal of computational chemistry* **38**, 1447–1461 (2017).

[17] Habeck, M. Bayesian modeling of biomolecular assemblies with cryo-em maps. *Frontiers in molecular biosciences* **4**, 15 (2017).

[18] Bonomi, M. *et al.* Bayesian weighing of electron cryo-microscopy data for integrative structural modeling. *Structure* **27**, 175–188 (2019).

[19] Cossio, P. *et al.* Bioem: Gpu-accelerated computing of bayesian inference of electron microscopy images. *Computer Physics Communications* **210**, 163–171 (2017).

[20] Páll, S. *et al.* Heterogeneous parallelization and acceleration of molecular dynamics simulations in gromacs. *J. Chem. Phys.* **153**, 134110 (2020).

[21] Jaynes, E. T. Information theory and statistical mechanics. *Physical review* **106**, 620 (1957).

[22] Cheng, Y., Grigorieff, N., Penczek, P. A. & Walz, T. A primer to single-particle cryo-electron microscopy. *Cell* **161**, 438–449 (2015).

[23] Kirkland, E. J. *Advanced computing in electron microscopy* (Springer, 1998).

[24] Hummer, G. & Köfinger, J. Bayesian ensemble refinement by replica simulations and reweighting. *The Journal of chemical physics* **143**, 12B634_1 (2015).

[25] Tang, G. *et al.* Eman2: an extensible image processing suite for electron microscopy. *Journal of structural biology* **157**, 38–46 (2007).

[26] Brown, A. *et al.* Tools for macromolecular model building and refinement into electron cryo-microscopy reconstructions. *Acta Crystallographica Section D: Biological Crystallography* **71**, 136–153 (2015).

[27] Choi, K. H. *et al.* Structure of a fructose-1, 6-bis (phosphate) aldolase liganded to its natural substrate in a cleavage-defective mutant at 2.3 å. *Biochemistry* **38**, 12655–12664 (1999).

[28] Wu, M., Lander, G. C. & Herzik Jr, M. A. Sub-2 angstrom resolution structure determination using single-particle cryo-em at 200 kev. *Journal of Structural Biology: X* **4**, 100020 (2020).

[29] Páll, S. *et al.* Heterogeneous parallelization and acceleration of molecular dynamics simulations in gromacs. *The Journal of Chemical Physics* **153**, 134110 (2020).

[30] Greengard, L. & Lee, J.-Y. Accelerating the nonuniform fast fourier transform. *SIAM review* **46**, 443–454 (2004).

[31] Leimkuhler, B. J., Reich, S. & Skeel, R. D. Integration methods for molecular dynamics. In *Mathematical Approaches to biomolecular structure and dynamics*, 161–185 (Springer, 1996).

[32] Ferrarotti, M. J., Bottaro, S., Pérez-Villa, A. & Bussi, G. Accurate multiple time step in biased molecular simulations. *Journal of chemical theory and computation* **11**, 139–146 (2015).

[33] Pettersen, E. F. *et al.* Ucsf chimera—a visualization system for exploratory research and analysis. *Journal of computational chemistry* **25**, 1605–1612 (2004).

[34] Mackerell Jr, A. D., Feig, M. & Brooks III, C. L. Extending the treatment of backbone energetics in protein force fields: Limitations of gas-phase quantum mechanics in reproducing protein conformational distributions in molecular dynamics simulations. *Journal of computational chemistry* **25**, 1400–1415 (2004).

[35] Bjelkmar, P., Larsson, P., Cuendet, M. A., Hess, B. & Lindahl, E. Implementation of the charmm force field in gromacs: analysis of protein stability effects from correction maps, virtual interaction sites, and water models. *Journal of chemical theory and computation* **6**, 459–466 (2010).

VII. ACKNOWLEDGEMENTS

This work was supported by grants from the Swedish Research Council (2017-04641, 2018-06479, 2019-02433), the Swedish e-Science Research Centre, the BioExcel Center of Excellence (EU 823830), the Knut and Alice Wallenberg Foundation (1484505) and the Carl Trygger Foundation (CTS-15:298). Computational resources were provided by SNIC, the Swedish National Infrastructure for Computing. We would like to thank Rebecca J. Howard and Marta Carroni for insightful discussions of the manuscript.

A. Contributions

C.B. initiated the project, derived the theoretical results and implemented the algorithms. L.Y. carried out the molecular dynamics simulations. L.Y. and C.B. analysed the molecular dynamics simulations. E.L. supervised the research. C.B., L.Y. and E.L. prepared figures and wrote the publication.