



Non-uniform refinement: adaptive regularization improves single-particle cryo-EM reconstruction

Ali Punjani^{1,2,3} , Haowei Zhang¹ and David J. Fleet^{1,2}

Cryogenic electron microscopy (cryo-EM) is widely used to study biological macromolecules that comprise regions with disorder, flexibility or partial occupancy. For example, membrane proteins are often kept in solution with detergent micelles and lipid nanodiscs that are locally disordered. Such spatial variability negatively impacts computational three-dimensional (3D) reconstruction with existing iterative refinement algorithms that assume rigidity. We introduce non-uniform refinement, an algorithm based on cross-validation optimization, which automatically regularizes 3D density maps during refinement to account for spatial variability. Unlike common shift-invariant regularizers, non-uniform refinement systematically removes noise from disordered regions, while retaining signal useful for aligning particle images, yielding dramatically improved resolution and 3D map quality in many cases. We obtain high-resolution reconstructions for multiple membrane proteins as small as 100 kDa, demonstrating increased effectiveness of cryo-EM for this class of targets critical in structural biology and drug discovery. Non-uniform refinement is implemented in the cryoSPARC software package.

Single-particle cryo-EM has transformed rapidly into a mainstream technique in biological research¹. Cryo-EM images individual protein particles, rather than crystals and has therefore been particularly useful for structural studies of integral membrane proteins, which are difficult to crystallize². These molecules are critical for drug discovery, targeted by more than half of drugs today³. Membrane proteins pose challenges in cryo-EM sample preparation, imaging and computational 3D reconstruction, as they are often of small size, appear in multiple conformations, have flexible subunits and are embedded in a detergent micelle or lipid nanodisc². These characteristics cause strong spatial variation in structural properties, such as rigidity and disorder, across the target molecule's 3D density. Traditional cryo-EM reconstruction algorithms, however, are based on the simplifying assumption of a uniform, rigid particle.

We develop an algorithm that incorporates such domain knowledge in a principled way, improving 3D reconstruction quality and allowing single-particle cryo-EM to achieve higher-resolution structures of membrane proteins. This expands the range of proteins that can be effectively studied and is especially important for structure-based drug design^{4,5}. We begin by formulating a cross-validation (CV) regularization framework for single-particle cryo-EM refinement and use it to account for the spatial variability in resolution and disorder found in a typical molecular complex. The framework incorporates general domain knowledge about protein molecules, without specific knowledge of any particular molecule and critically, without need for manual user input. Through this framework we derive a new algorithm called non-uniform refinement, which automatically accounts for structural variability, while ensuring that key statistical properties for validation are maintained to mitigate the risk of over-fitting during 3D reconstruction.

With a graphics processing unit-accelerated implementation of non-uniform refinement in the cryoSPARC software package⁶, we demonstrate improvements in resolution and map quality for a range of membrane proteins. We show results on a 48-kDa membrane protein in lipid nanodisc with a Fab bound, a 180-kDa membrane protein complex with a large detergent micelle and a 245-kDa

sodium channel complex with flexible domains. Non-uniform refinement is reliable and automatic, requiring no change in parameters between datasets and is without reliance on hand-made spatial masks or manual labels.

Iterative refinement and regularization. In standard cryo-EM 3D structure determination^{6–8}, a generative model describes the formation of two-dimensional (2D) electron microscope images from a target 3D protein density (Coulomb potential). According to the model, the target density is rotated, translated and projected along the direction of the electron beam. The 2D projection is modulated by a microscope contrast transfer function (CTF) and corrupted by additive noise. The goal of reconstruction is to infer the 3D density map from particle images, without knowledge of latent 3D pose variables, that is, the orientation and position of the particle in each image. Iterative refinement methods formulate inference as a form of maximum likelihood or maximum a posteriori optimization^{6,9–11}. Such algorithms can be viewed as a form of block-coordinate descent or expectation-maximization¹², each iteration comprising an E-step, estimating the pose of each particle image, given the 3D structure, and an M-step, regularized 3D density estimation given the latent poses.

Like many inverse problems with noisy, partial observations, the quality of cryo-EM map reconstruction depends heavily on regularization. Regularization methods, widely used in computer science and statistics, leverage prior domain knowledge to penalize unnecessary model complexity and avoid over-fitting. In cryo-EM refinement, regularization is needed to mitigate the effects of imaging and sample noise so that protein signal alone is present in the inferred 3D density and so accumulated noise does not contaminate latent pose estimates.

Existing refinement algorithms use an explicit regularizer in the form of a shift-invariant linear filter, typically obtained from Fourier shell correlation (FSC)^{6,10,13–17}. Such filters smooth the 3D structure using the same kernel, and hence the same degree of smoothing, at all locations. As FSC captures the average resolution of the map, such filters presumably under- and over-regularize

¹Department of Computer Sciences, University of Toronto, Toronto, Ontario, Canada. ²Vector Institute, Toronto, Ontario, Canada. ³Structura Biotechnology Inc., Toronto, Ontario, Canada. ✉e-mail: alipunjani@cs.toronto.edu; fleet@cs.toronto.edu




Fig. 1 | A 3D map from uniform refinement reveals spatial variations in resolution in a prototypical membrane protein (TRPA1 ion channel, EMPIAR-10024). Following the default FSC-based regularization and B-factor sharpening in cryoSPARC, the density map has been thresholded at a relatively low value to clearly visualize regions with substantial levels of noise. Color depicts local resolution²² as a proxy for local structure properties. Red indicates higher resolution (e.g. the core inner region), yellow indicates moderate resolutions (e.g. the solvent-facing region) and blue indicates poor resolution (e.g. the disordered detergent micelle and the flexible tail at the bottom).

different regions, allowing noise accumulation in some regions and a loss of resolvable detail in others. This effect should be pronounced with membrane proteins that have highly non-uniform rigidity and disorder across the molecule. As a motivating example, Fig. 1 shows a reconstruction of the TRPA1 membrane protein¹⁸ with a relatively low density threshold to help visualize regions with substantial noise levels, which indicate over-fitting (e.g. the disordered micelle and the flexible tail at and bottom of the protein). We hypothesize that under-fitting occurs in the core region where over-regularization attenuates useful signal. As such, accumulated noise and attenuated signal will degrade pose estimates during refinement, limiting final structure quality. For inference problems of this type, the amount and form of regularization depends on regularization parameters. Correctly optimizing these parameters is often critical, but care must be taken to ensure that the optimization itself is not also prone to over-fitting.

Results

We next outline the formulation of an adaptive form of regularization and with it, a new refinement algorithm called non-uniform refinement. We discuss its properties and demonstrate its application on several membrane protein datasets.

Adaptive cross-validation regularization. With the aim of incorporating spatial non-uniformity into cryo-EM reconstruction, we formulate a family of regularizers denoted r_θ , with parameters $\theta(x)$ that depend on spatial position x . Given a 3D density map $m(x)$, the regularization operator, evaluated at x , is defined by

$$(r_\theta \circ m)(x) = \sum_{\xi} h(\xi; \theta(x)) m(\xi - x), \quad (1)$$

where $h(x; \psi)$ is symmetric smoothing kernel, the spatial scale of which is determined by parameter ψ .

This family provides greater flexibility than shift-invariant regularizers, but in exchange, requires making the correct choice of a

new set of parameters, $\theta(x)$. We formulate the selection of the regularization parameters as an optimization subproblem during refinement, for which we adopt a twofold CV objective^{19,20}. The data are first randomly partitioned into two halves. On each of two trials, one part is considered as training data and the other is treated as held-out validation data. To find the regularizer parameters θ , one minimizes the sum of the per-trial validation errors (e), measuring consistency between the model and the validation data; that is,

$$E(\theta) = e(r_\theta \circ m_1, m_2) + e(r_\theta \circ m_2, m_1) \quad (2)$$

where m_1 and m_2 are reconstructions from the two folds of the data. Similar objectives have been used for image de-noising²¹. We also introduce constraints on the parameters $\theta(x)$ to control degrees of freedom. The optimization problem is solved using a discretized search algorithm (Methods).

The resulting CV regularizer automatically identifies regions of a protein density with differing structural properties, optimally removing noise from each region. Fig. 2 illustrates the difference between uniform filtering (FSC-based) and the CV-optimal adaptive regularizer in non-uniform refinement. Shift-invariant regularization smooths all regions equally, while the adaptive regularizer preferentially removes noise from disordered regions, retaining high-resolution signal in well-structured regions that is essential for 2D–3D alignment.

Non-uniform refinement algorithm. The non-uniform refinement algorithm takes as input a set of particle images and a low-resolution ab initio 3D map. The data are randomly partitioned into two halves, each of which is used to independently estimate a 3D half-map. This ‘gold-standard’ refinement¹⁷ allows use of FSC for evaluating map quality and for comparison with existing algorithms. The key to non-uniform refinement is the adaptive CV regularization, applied at each iteration of 3D refinement. The regularizer parameters $\theta(x)$ are estimated independently for each half-map (Methods), adhering to the ‘gold-standard’ assumptions. In contrast, in conventional uniform refinement the two half-maps are not strictly independent as the regularization parameters, determined by FSC, are shared by both half-maps. Finally, the optimization and application of the adaptive regularizer cause non-uniform refinement to be approximately two times slower than uniform refinement in cryoSPARC.

Validation with membrane protein datasets. We experimentally compared non-uniform refinement with conventional uniform refinement with three membrane protein datasets, all processed in cryoSPARC (see Supplementary Information for an additional membrane protein with no soluble region). Both algorithms were given the same ab initio structures. Except for the regularizer, all parameters and stages of data analysis, including the 2D–3D alignment and back-projection, were identical in uniform and non-uniform refinements. The default use of spatial solvent masking during 2D–3D alignment in cryoSPARC was used for both algorithms (Supplementary Information). No manual masks were used and no masking was used to identify or separate micelle or nano-disc regions. The same non-uniform refinement default parameters (other than symmetry) were used for all datasets.

When computing gold-standard FSC during the analysis of the reconstructed 3D density maps, for each dataset we used the same mask for uniform and non-uniform refinement. The masks were tested using phase randomization¹³ to avoid FSC bias. Also, the same B-factor was used to sharpen both uniform and non-uniform refinement maps for each dataset. This consistency helped to ensure that visible differences in 3D structure were due solely to algorithmic differences. To color maps using local resolution, we used a straightforward implementation of Blocres²² for resolution estimation. No local filtering or local sharpening was used for visualization.




Fig. 2 | Images illustrate the difference between shift-invariant regularization and adaptive (CV-optimal) regularization on a membrane protein (PfCRT, EMPIAR-10330). A central slice through a raw reconstruction (M-step) of a half-map after nine iterations (left). It shows substantial levels of noise. The same reconstruction is shown after uniform isotropic filtering (based on FSC between half-maps). It shows uniform noise attenuation throughout the slice (middle). The same half-map reconstruction after adaptive regularization with the optimal CV regularizer shows greater attenuation of noise in the solvent background and the nanodisc region, while better preserving the high-resolution structure in the well-ordered protein region (right).

Uniform and non-uniform refinement density maps were thresholded to contain the same total volume for visual comparison.

We also note that the FSC-based regularizer in conventional uniform refinement is equivalent to a shift-invariant regularizer optimized with CV. That is, one can show that the optimal shift-invariant filter under CV with squared error has a transfer function equivalent to the FSC curve. Thus, the experiments below also capture the differences between adaptive and shift-invariant regularization.

STRA6-CaM: CV regularization yields improved pose estimates and FSC. Zebrafish STRA6-CaM²³ is a 180-kDa C2-symmetric membrane protein complex comprising the 74-kDa STRA6 protein bound to calmodulin (CaM). STRA6 mediates the uptake of retinol in various tissues. We processed a dataset of 28,848 particle images of STRA6-CaM in a lipid nanodisc, courtesy of O. Clarke and F. Mancina (O. Clarke, personal communication). Non-uniform refinement provides a substantial improvement in nominal resolution from 4.0 Å to 3.6 Å (Fig. 3a), indicating an improvement in the average signal-to-noise over the entire structure. However, different regions exhibited different resolution characteristics (Fig. 3c), as is often observed with protein reconstructions²². There was substantial improvement in structural detail in most regions, while peripheral and flexible regions remained at low resolutions. Differences in structure quality, clearly visible in detailed views of α -helices (Fig. 3d), are especially important during atomic model building, where in many cases, protein backbone and side chains can only be traced with confidence in the non-uniform refinement map. Improvements in structure quality coincided with changes in particle alignments (Fig. 3b), approximately 3° on average. While disorder in the lipid nanodisc and nonrigidity of the CaM subunits are problematic for uniform refinement, adaptive regularization in non-uniform refinement reduces the influence of noise on alignments and produces improved map quality, especially in the periphery of the protein.

PfCRT: enabling atomic model building from previously unusable data. The *Plasmodium falciparum* chloroquine resistance transporter (PfCRT) is an asymmetric 48-kDa membrane protein²⁴. Mutations in PfCRT are associated with the emergence of resistance to chloroquine and piperazine as antimalarial treatments. We processed 16,905 particle images of PfCRT in lipid nanodisc with a Fab bound (EMPIAR-10330 (ref. 24)). For PfCRT, the difference in resolution (Fig. 4a) and map quality (Fig. 4c) between uniform and non-uniform refinement is striking.

Using uniform refinement, reaching 6.9 Å, transmembrane α -helices are barely resolvable. Non-uniform refinement recovers

signal up to 3.6 Å and provides a map from which an atomic model can be built with confidence. Transmembrane α -helices can be directly traced, including side chains. In contrast, the uniform refinement map does not show helical pitch and does not separate β -strands in the Fab domain. Indeed, an early version of the non-uniform refinement algorithm was essential for reconstructing the published high-resolution density map and model²⁴.

On the spectrum of proteins studied with cryo-EM, the PfCRT-Fab complex (100 kDa) is small. The lipid nanodisc (~50 kDa) also accounts for a large fraction of total particle molecular weight (see Fig. 2). We hypothesized that disorder in this relatively large nanodisc region leads to over- and under-regularization in uniform refinement. Most particle images exhibited orientation differences >6° between the two algorithms (Fig. 4b), suggesting that a large fraction of particle images were grossly misaligned by uniform refinement.

Na_v1.7 channel: improvement in high-resolution features. Na_v1.7 is a voltage-gated sodium channel found in the human nervous system²⁵. It plays a role in the generation and conduction of action potentials and is targeted by toxins and therapeutic drugs (e.g. for pain relief). We processed data of Na_v1.7 bound to two Fabs, forming a 245-kDa C2-symmetric complex solubilized in detergent²⁵. Following pre-processing (Methods), as described elsewhere²⁵, we detected both active and inactive conformational states of the channel. We obtained reconstructions with resolutions better than the published literature for both, but here we focus on 300,759 particle images of the active state.

Compared to the preceding datasets, the Na_v1.7 complex has a higher molecular weight and, in relative terms, a smaller detergent micelle. However, other regions are disordered or flexible, namely, a central four-helix bundle, peripheral transmembrane domains and the Fabs. For Na_v1.7, uniform refinement reaches 3.4 Å resolution and non-uniform refinement reaches an improved 3.1 Å (Fig. 5a). This result is also an improvement over the published result of 3.6 Å (EMDB-0341)²⁵, where the authors performed all processing in cisTEM¹⁰.

With non-uniform refinement, map quality was clearly improved in central transmembrane regions, while some flexible parts of the structure (Fab domains and four-helix bundle) remained at intermediate resolutions (Fig. 5c). In detailed views of α -helices (Fig. 5d) closer to the periphery of the protein, improvement in map quality and interpretability was readily apparent in the non-uniform refinement map, allowing modeling of side-chain rotamers with confidence. Central α -helices showed less improvement, but map quality




Fig. 3 | Results of uniform and non-uniform refinement from 28,848 particle images of STRA6-CaM in lipid nanodisc (pixel size 1.07 Å, box size 256 pixels, C2 symmetry yielding 57,696 asymmetric units). **a**, FSC curves computed with the same mask for both refinements, showing improvement from 4.0 Å to 3.6 Å between uniform and non-uniform methods. **b**, Histograms of change in particle alignment pose and shift between uniform and non-uniform refinement. **c**, 3D density maps from uniform and non-uniform refinement are filtered using the corresponding FSC curves and sharpened with the same B-factor, -140 Å^2 . No local filtering or sharpening was used and thresholds were set to keep the enclosed volume constant. Map differences were due to algorithmic rather than visualization differences. Map color depicts local resolution (Blocres²²) on a single-color scale and shows how map improvement depends on the region within the map. **d**, Individual α -helical segments from the non-uniform map (purple) and uniform map (gray) illustrate differences in resolvability of backbone and side chains. The left-most α -helix is peripheral, whereas the right-most is central. Docked atomic model is courtesy of O. Clarke.

remained equal or slightly improved, indicating that reconstructions of protein regions without disorder were not harmed by using non-uniform refinement.

Increased data efficiency. Examining map quality as a function of dataset size further helped to explore data efficiency. Figure 6 plots inverse resolution versus the number of particles. Such plots are useful as they relate to image noise and computational extraction of signal¹⁵. Across all datasets, non-uniform refinement reached the same

resolution as uniform refinement with fewer than half the number of particle images. It has also been argued that higher curves indicate more accurate pose estimates²⁶. Notably, the resolution gap between uniform and non-uniform refinement persists over a wide range of dataset sizes. While this resolution gap may decrease for much larger datasets than those considered here, the collection of more data alone may not allow uniform refinement to match the performance of non-uniform refinement for membrane protein targets until resolution is saturated as one nears fundamental limits.




Fig. 4 | Results of uniform and non-uniform refinement from 16,905 particle images of PfCRT in lipid nanodisc with a single Fab bound (pixel size 0.5175 Å, box size 300 pixels, no symmetry). **a**, FSC curves, computed with the same mask, show numerical improvement from 6.9 Å to 3.6 Å. **b**, Histograms of change in particle alignments between uniform and non-uniform refinement. **c**, 3D density maps from uniform and non-uniform refinement, both filtered using the corresponding FSC curve and sharpened with the same B-factor of -100 Å^2 . No local filtering or sharpening is used and thresholds are set to keep the enclosed volume constant. Density differences are thus due to algorithmic rather than visualization differences. Maps are colored by local resolution from Blocres²², all on the same color scale. **d**, Individual α -helical segments and β -strands from the non-uniform map (purple) and uniform map (gray) illustrate localized differences in resolvability between the maps.

Discussion

With adaptive regularization and effective optimization of regularization parameters using CV, non-uniform refinement achieves reconstructions of higher resolution and quality from single-particle cryo-EM data. It is particularly effective for membrane proteins, which exhibit varying levels of disorder, flexibility or occupancy. Here we focused on one specific family of adaptive regularizers (Methods), but it is possible within the same framework to explore other families that may be even more effective. For example, one could look to the extensive de-noising literature for different regularizers.

Spatial variability of structure properties and the existence of over- and under-regularization in uniform refinement are well known. For instance, cryo-EM methods for estimating local map resolution once 3D refinement is complete leverage statistical tests on the coefficients of a windowed Fourier transform²³ or some form of wavelet transform^{27,28}. Although non-uniform refinement does not estimate resolution per se, the regularizer parameter $\theta(x)$ is related to a local frequency band limit. As such it might be viewed as a proxy for local resolution, but with important differences. Notably, $\theta(x)$ is optimized with a CV objective and it does not depend on a specific definition of 'local resolution' nor on an explicit resolution estimator.




Fig. 5 | Results of uniform and non-uniform refinement on 300,759 particle images of Na_v1.7 in a detergent micelle with two Fabs bound (pixel size 1.21 Å, box size 360 pixels, refined with C2 symmetry yielding 601,518 asymmetric units). **a, FSC curves for uniform and non-uniform refinement indicate global numerical resolutions of 3.4 Å to 3.1 Å, respectively. **b**, Histograms of change in particle alignments between uniform and non-uniform refinement. Optimized adaptive regularization yields improved alignments through multiple iterations. **c**, The 3D density maps were filtered using corresponding FSC curves, sharpened with a B-factor of -85 Å^2 . Thresholds were set to keep the enclosed volume constant. No local filtering or sharpening was used. Color depicts local resolution (Blocres²²) on the same color scale. **d**, Individual α -helical segments from the non-uniform map (purple) and uniform map (gray) illustrate localized differences in resolvability in the peripheral transmembrane domain (left, center) and the central core (right). Docked atomic model is PDB 6N4Q²⁵.**

Local resolution estimates^{22,29} or statistical tests³⁰ have also been used to adaptively filter 3D maps, for visualization, to assess local map quality or to aid molecular model building. The family of filters and frequency cutoffs are typically selected to maximize visual quality. They are not optimized for the local resolution estimator nor is consideration given to the number of degrees of freedom in filter parameters or the strict independence of half-maps, all critical for reliable regularization. Thus, while local resolution estimation followed by local filtering is useful for post-processing, its use for regularization during iterative refinement (e.g. EMAN2.2

documentation⁹), can yield over-fitting (Methods). Map artifacts such as spikes or streaks are especially problematic for datasets with junk particles, structured outliers or small difficult-to-align particles. Non-uniform refinement couples an implicit resolution measure to the choice of regularizer, with optimization designed to control model capacity and avoid over-fitting of regularization parameters (Methods).

Another related technique, used in cisTEM¹⁰ and Frealign⁷ and among the first to acknowledge and address under- and over-fitting, entails manual creation of a mask to label a local




Fig. 6 | Plots of inverse resolution (frequency squared) versus number of particles on a log scale, for each of the STRA6-CaM, PfCRT and Na_v1.7 datasets. Each point represents an independent run of uniform (black) or non-uniform (purple) refinement, using a random sample from the entire dataset. For each dataset, the same mask is used to compute the FSC at all sample sizes. Data points for non-uniform refinement have higher resolution value than for uniform refinement at all sample sizes, for all three datasets (STRA6-CaM²³, PfCRT²⁴ and Na_v1.7 (ref. ²⁵)).

region one expects to be disordered (e.g. detergent micelle), followed by low-pass filtering in that region to a pre-set resolution at each refinement iteration. While this shares the motivation for non-uniform refinement, it relies on manual interaction during refinement, often necessitating a tedious trial and error process that can be difficult to replicate.

SIDESPLITTER³¹ is a recently proposed method to mitigate over-fitting during refinement. Based on estimates of signal-to-noise ratios from voxel statistics, it smooths local regions with low signal-to-noise more aggressively than indicated by a global FSC-based resolution. The method does not directly address under-fitting nor does it maintain independence of half-maps or control for the number of degrees of freedom in the local filter model. More generally, the CV framework here is robust to different noise distributions and model mis-specification, which can be problematic for methods with parametric noise models. Empirical results indicate that SIDESPLITTER mitigates over-fitting and shows improvement in map quality relative to uniform refinement.

Bases other than the Fourier basis may also enable local control of reconstruction. Wavelet bases have been used for local resolution estimation, but less commonly for reconstruction³². Kucukelbir et al.³³ proposed the use of an adaptive wavelet basis and a sparsity prior. While similar in spirit to the goal of non-uniform refinement, this method has a single regularizer parameter for the entire 3D map, computed from noise in the corners of particle images. This may not capture variations in noise due to disorder, motion or partial occupancy.

Non-uniform refinement has been successful in helping to resolve several new structures. Examples include multiple conformations of an ABC exporter³⁴, mTORC1 docked on the lysosome³⁵, the respiratory syncytial virus polymerase complex³⁶, a GPCR-G protein-β-arrestin megacomplex³⁷ and the SARS-CoV-2 spike protein³⁸.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41592-020-00990-8>.

Received: 12 February 2020; Accepted: 6 October 2020;
Published online: 30 November 2020

References

- Cheng, Y. Single-particle Cryo-EM at crystallographic resolution. *Cell* **161**, 450–457 (2015).
- Cheng, Y. Membrane protein structural biology in the era of single particle cryo-EM. *Curr. Opin. Struct. Biol.* **52**, 58–63 (2018).
- Overington, J. P., Al-Lazikani, B. & Hopkins, A. L. How many drug targets are there? *Nat. Rev. Drug Discov.* **5**, 993–996 (2006).
- Renaud, J.-P. et al. Cryo-EM in drug discovery: achievements, limitations and prospects. *Nat. Rev. Drug Discov.* **17**, 471–492 (2018).
- Scapin, G., Potter, C. S. & Carragher, B. Cryo-EM for small molecules discovery, design, understanding, and application. *Cell Chem. Biol.* **25**, 1318–1325 (2018).
- Punjani, A., Rubinstein, J. L., Fleet, D. J. & Brubaker, M. A. CryoSPARC: Algorithms for rapid unsupervised cryo-em structure determination. *Nat. Methods* **14**, 290–296 (2017).
- Grigorieff, N. FREELIGN: High resolution refinement of single particle structures. *J. Struct. Biol.* **157**, 117–125 (2007).
- Scheres, S. H. W. & Bayesian, A. view on cryo-em structure determination. *J. Mol. Biol.* **415**, 406–418 (2012).
- Bell, J. M., Chen, M., Baldwin, P. R. & Ludtke, S. J. High resolution single particle refinement in EMAN2.1. *Methods* **100**, 25–34 (2016).
- Grant, T., Rohou, A. & Grigorieff, N. cisTEM, user-friendly software for single-particle image processing. *eLife* **7**, e35383 (2018).
- Zivanov, J. et al. New tools for automated high-resolution cryo-em structure determination in RELION-3. *eLife* **7**, e42166 (2018).
- Dempster, A. P., Laird, N. M. & Rubin, D. B. Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. Ser. B* **39**, 1–38 (1977).
- Chen, S. et al. High-resolution noise substitution to measure overfitting and validate resolution in 3D structure determination by single particle electron cryomicroscopy. *Ultramicroscopy* **135**, 24–35 (2013).
- Haraux, G. & van Heel, M. Exact filters for general geometry three dimensional reconstruction. *Optik* **73**, 146–156 (1986).
- Rosenthal, P. B. & Henderson, R. Optimal determination of particle orientation, absolute hand, and contrast loss in single-particle electron cryomicroscopy. *J. Mol. Biol.* **333**, 721–745 (2003).
- Scheres, S. H. W. RELION: implementation of a Bayesian approach to cryo-EM structure determination. *J. Struct. Biol.* **180**, 519–530 (2012).
- Scheres, S. H. W. & Chen, S. Prevention of overfitting in cryo-EM structure determination. *Nat. methods* **9**, 853–854 (2012).
- Paulsen, C. E., Armache, J. P., Gao, Y., Cheng, Y. & Julius, D. Structure of the TRPA1 ion channel suggests regulatory mechanisms. *Nature* **520**, 511–517 (2015).
- Golub, G. M., Heath, M. & Wahba, G. Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics* **21**, 215–223 (1979).
- Wahba, G. A comparison of GCV and GML for choosing the smoothing parameter in the generalized spline smoothing problem. *Ann. Stat.* **13**, 1378–1402 (1985).
- Lehtinen, J. et al. Noise2Noise: learning image restoration without clean data. *Proc. Mach. Learn. Res.* **7**, 4620–4631 (2018).
- Cardone, G., Heymann, J. B. & Steven, A. C. One number does not fit all: mapping local variations in resolution in cryo-em reconstructions. *J. Struct. Biol.* **184**, 226–236 (2013).
- Chen, Y. et al. Structure of the strA6 receptor for retinol uptake. *Science* **353**, aad8266 (2016).
- Kim, J. et al. Structure and drug resistance of the *Plasmodium falciparum* transporter PfCRT. *Nature* <https://doi.org/10.1038/s41586-019-1795-x> (2019).
- Xu, H. et al. Structural basis of Nav1.7 inhibition by a gating-modifier spider toxin. *Cell* <https://doi.org/10.1016/j.cell.2018.12.018> (2019).
- Stagg, S., Noble, A., Spilman, M. & Chapman, M. Reslog plots as an empirical metric of the quality of cryo-EM reconstructions. *J. Struct. Biol.* **185**, 418–426 (2014).
- Kucukelbir, A., Sigworth, F. J. & Tagare, H. D. Quantifying the local resolution of cryo-EM density maps. *Nat. Methods* **11**, 63–65 (2014).

28. Vilas, J. L. et al. Monores: automatic and accurate estimation of local resolution for electron microscopy maps. *Structure* **26**, 337–344 (2018).
29. Felsberg, M. & Sommer, G. The monogenic signal. *IEEE Trans. Signal Process.* **49**, 3136–3144 (2001).
30. Ramlaul, K., Palmer, C. M. & Aylett, C. H. S. A local agreement filtering algorithm for transmission EM reconstructions. *J. Struct. Biol.* **205**, 30–40 (2019).
31. Ramlaul, K., Palmer, C. M., Nakane, T. & Aylett, C. H. S. Mitigating local over-fitting during single particle reconstruction with SIDESPLITTER. *J. Struct. Biol.* **211**, 107545 (2020).
32. Vonesch, C., Wang, L., Shkolnisky, Y. and Singer, A. Fast wavelet-based single particle reconstruction in cryo-EM. *Proc. IEEE Int. Symp. Biomed. Imaging* <https://doi.org/10.1109/ISBI.2011.5872791> (2011).
33. Kucukelbir, A., Sigworth, F. J. & Tagare, H. D. A Bayesian adaptive basis algorithm for single particle reconstruction. *J. Struct. Biol.* **179**, 56–67 (2012).
34. Hofmann, S. et al. Conformation space of a heterodimeric ABC exporter under turnover conditions. *Nature* **571**, 580–583 (2019).
35. Rogala, K. B. et al. Structural basis for the docking of mTORC1 on the lysosomal surface. *Science* **366**, 468–475 (2019).
36. Gilman, M. S. A. et al. Structure of the respiratory syncytial virus polymerase complex. *Cell* **179**, 193–204 (2019).
37. Nguyen, A. H. et al. Structure of an endosomal signaling GPCR-G protein- β -arrestin megacomplex. *Nat. Struct. Mol. Biol.* **26**, 1123–1131 (2019).
38. Wrapp, D. et al. Cryo-EM structure of the 2019-nCoV spike in the prefusion conformation. *Science* **367**, 1260–1263 (2020).
- Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.
- © The Author(s), under exclusive licence to Springer Nature America, Inc. 2020

Methods

Regularization in iterative refinement. In the standard cryo-EM 3D reconstruction problem setup, the target 3D density map m is typically parameterized as a real-space 3D array with density at each voxel, in a Cartesian grid of box size N , and a corresponding discrete Fourier representation, $\hat{m} = Fm$. The goal of reconstruction is to infer the 3D densities of the voxel array, called model parameters. Representing the 3D density, its 2D projections, the observed images and the noise model in the Fourier domain is common practice for computational efficiency, exploiting the well-known convolution and Fourier-slice theorems^{6–9}. The unobserved pose variables for each image are latent variables.

Algorithm 1. Iterative refinement (expectation-maximization)

Require: Particle image dataset \mathcal{D} and ab initio 3D map	
1:	Use smoothed ab initio 3D map array as the initial model parameters $m^{(0)}$
2:	while not converged do
3:	E-step: Given current estimate of model parameters, $m^{(t-1)}$, from step $t-1$, estimate (via marginalization or maximization) the latent variables: $z^{(t)} \leftarrow f(m^{(t-1)}, \mathcal{D})$
4:	M-step: Given the latent variables $z^{(t)}$, compute raw estimates of the model parameter $\tilde{m}^{(t)}$ (without regularization): $\tilde{m}^{(t)} \leftarrow h(z^{(t)}, \mathcal{D})$
5:	Regularize: Given noisy model parameters $\tilde{m}^{(t)}$, apply the regularization operator, r_θ , with regularization parameters θ : $m^{(t)} \leftarrow r_\theta(\tilde{m}^{(t)})$
6:	end while

Iterative refinement methods (Algorithm 1), which provide state-of-the-art results^{6,9–11}, can be interpreted as variants of block-coordinate descent or the expectation-maximization algorithm¹². In cryo-EM, and more generally in inverse problems with noisy, partial observations, a critical component that modulates the quality of the results is regularization.

One can regularize problems explicitly, using a prior distribution over model parameters or implicitly, by applying a regularization operator to the model parameters during optimization. Iterative refinement methods tend to use implicit regularizers, attenuating noise in the reconstructed map at each iteration. In either case, the separation of signal from noise is the crux of many inference problems.

In the cryo-EM refinement problem, like many latent variable inverse problems, there is an additional interplay between regularization, noise buildup and the estimation of latent variables. Retained noise due to under-regularization will contaminate the estimation of latent variables. This contamination is propagated to subsequent iterations and causes over-fitting.

This paper reconsiders the task of regularization based on the observation that common iterative refinement algorithms often systematically under-fit and over-fit different regions of a 3D structure simultaneously. This causes a loss of resolvable detail in some parts of a structure and the accumulation of noise in others. The reason stems from the use of frequency-space filtering as a form of regularization. Some programs, such as cisTEM¹⁰, use a strict resolution cutoff, beyond which Fourier amplitudes are set to zero before alignment of particle images to the current 3D structure. In RELION¹⁶, regularization was initially formulated with an explicit Gaussian prior on Fourier amplitudes of the 3D structure, with a hand-tuned parameter that controls Fourier amplitude shrinkage. Later versions of RELION¹⁷ and cryoSPARC's homogeneous (uniform) refinement⁶ use a transfer function (or Wiener filter) determined by FSC computed between two half-maps^{13–15}.

Such methods presume a Fourier basis and shift-invariance. Although well suited to stationary processes, they are less well suited to protein structures, which are spatially compact and exhibit non-uniform disorder (e.g. due to motion or variations in signal resolution). FSC, for instance, provides an aggregate measure of resolution. To the extent that FSC under- or over-estimates resolution in different regions, FSC-based shift-invariant filtering will over-smooth some regions, attenuating useful signal and under-smooth others, leaving unwanted noise. To address these issues we introduce a family of adaptive regularizers that can, in many cases, find better 3D structures with improved estimates of the latent poses during refinement.

Cross-validation regularization. We formulate a new regularizer for cryo-EM reconstruction in terms of the minimization of a CV objective^{19,20}. CV is a general principle that is widely used in machine learning and statistics for model selection and parameter estimation with complex models. In CV, observed data are randomly partitioned into a training set and a held-out validation set. Model parameters are inferred using the training data, the quality of which is then assessed by measuring an error function applied to the validation data. In k -fold CV, the observations are partitioned into k parts. In each of k trials, one part is selected as the held-out validation set and the remaining $k-1$ parts comprise the training set. The per-trial validation errors are summed, providing the total CV error. This procedure measures agreement between the optimized model and the

observations without bias due to over-fitting. Rather, over-fitting during training is detected directly as an increase in the validation error. Notably, formulating regularization in a CV setting provides a principled way to design regularization operators that are more complex than the conventional, isotropic frequency-space filters. The CV framework is not restricted to a Fourier basis. One may consider more complex parameterizations, the use of meta-parameters and incorporate cryo-EM domain knowledge.

Given a family of regularizers r_θ with parameters θ , the minimization of CV error to find θ is often applied as an outer loop. This requires the optimization of model parameters m to be repeated many times with different values of θ , a prohibitively expensive cost for problems like cryo-EM. Instead, one can also perform CV optimization as an inner loop, while optimization of model parameters occurs in the outer loop. Regularizer parameters θ are then effectively optimized on-the-fly, preventing under- or over-fitting without requiring multiple 3D refinements to be completed.

To that end, consider the use of twofold CV optimization to select the regularization operator, denoted $r_\theta(m)$ in the regularization step in Algorithm 1 (note that $k > 2$ is also possible). The dataset \mathcal{D} is partitioned into two halves, \mathcal{D}_1 and \mathcal{D}_2 and two (unregularized) refinements are computed, namely m_1 and m_2 . For each, one half of the data is the 'training set' and the other is held out for validation. To find the regularizer parameters θ we wish to minimize the total CV error E , that is,

$$\begin{aligned} \min_{\theta} E(\theta) &= \min_{\theta} e(r_\theta(m_1); \mathcal{D}_2) + e(r_\theta(m_2); \mathcal{D}_1) \\ &= \min_{\theta} \|r_\theta(m_1) - m_2\|^2 + \|r_\theta(m_2) - m_1\|^2 \end{aligned} \quad (3)$$

where e is the negative log likelihood of the validation half-set given the regularized half-map. The second line simplifies this expression by using the raw reconstruction from the opposite half-set as a proxy for the actual observed images. Note that assumptions for 'gold-standard' refinement¹⁷ are not broken in this procedure. With the L2 norm, equation (3) reduces to a sum of per-voxel squared errors, corresponding to white Gaussian noise between the half-set reconstructions. When the choice of θ causes r_θ to remove too little noise from the raw reconstruction, the residual error E will be unnecessarily large. If θ causes r_θ to over-regularize, removing too much structure from the raw reconstruction, then E increases as the structure retained by r_θ no longer cancels corresponding structure in the opposite half-map. As such, minimizing $E(\theta)$ provides the regularizer that optimally separates signal from noise.

We note that similar objectives have been used for general image de-noising²¹ and adapted for cryo-EM images^{39–41}; however, in these methods the aim was to learn a general neural network de-noiser, whereas the goal here is to optimize regularization parameters on a single data sample. It is also worth noting that this formulation can be extended in a straightforward way to compare each half-set reconstruction against images directly (dealing appropriately with the latent pose variables) or to use error functions corresponding to different noise models.

Regularization parameter optimization. The CV formulation in equation (3) provides great flexibility in choosing the family of regularizers r_θ , taking domain knowledge into account. For non-uniform refinement, we wish to accommodate protein structures with spatial variations in disorder, motion and resolution. Accordingly, we define the regularizer to be a space-varying linear filter. The filter's spatial extent is determined by the regularization parameter $\theta(x)$, which varies with spatial position. Here, we write the regularizer in operator form:

$$(r_\theta \circ m)(x) = \sum_{\xi} h(\xi; \theta(x)) m(\xi - x), \quad (4)$$

where $h(x; \psi)$ is symmetric smoothing kernel, the spatial scale of which is specified by ψ . In practice we let $h(x; \psi)$ be an eighth-order Butterworth kernel. The eighth-order kernel provides a middle ground between the relatively poor frequency resolution of the Gaussian kernel and the sharp cutoff of the sinc kernel, which suffers from spatial ringing. We have experimented with different orders of Butterworth filters and found that an eighth-order kernel performs well on a broad range of particles.

When equation (4) is combined with the CV objective for the estimation of $\theta(x)$, one obtains

$$\begin{aligned} \theta^* &= \arg \min_{\theta} E(\theta) \\ &= \arg \min_{\theta} \sum_x | (r_\theta \circ m_1)(x) - m_2(x) |^2 + | (r_\theta \circ m_2)(x) - m_1(x) |^2. \end{aligned} \quad (5)$$

With one regularization parameter at each voxel, that is, $\theta(x)$, this reduces to a large set of decoupled optimization problems, one for each voxel. That is, for voxel x one obtains

$$\theta^*(x) = \arg \min_{\theta(x)} | (r_\theta \circ m_1)(x) - m_2(x) |^2 + | (r_\theta \circ m_2)(x) - m_1(x) |^2 \quad (6)$$

With this decoupling, $\theta(x)$ can transition quickly from voxel to voxel, yielding high spatial resolution. On the other hand, the individual sub-problems in equation

(6) are not well constrained as each parameter is estimated from data at a single location, so the parameter estimates are not useful. In essence, our regularizer design has two competing goals, namely, reliable signal detection and high spatial resolution (respecting boundaries between regions with different properties). Signal detection improves through aggregation of observations (e.g. neighboring voxels), while high spatial resolution prefers minimal aggregation (equation (6)).

To improve signal detection, we further constrain θ to be smooth. That is, although in some regions θ should change quickly (solvent–protein boundaries), in most regions we expect it to change slowly (in solvent and regions of rigid protein mass). Smoothness effectively limits the number of degrees of freedom in θ , which is important to ensure that θ itself does not over-fit during iterative refinement. One can encourage smoothness in θ by explicitly penalizing spatial derivatives of θ in the objective (equation (5)), but this yields a Markov random field problem that is hard to optimize. Alternatively, one can express θ in a low-dimensional basis (e.g. radial basis functions), but this requires prior knowledge of the expected degree of smoothness. Instead, we adopt a simple but effective approach. Assuming that θ is smoothly varying, we treat measurements in the local neighborhood of x as additional constraints on $\theta(x)$. A window function can be used to give more weight to points close to x . We thereby obtain the following least-squares objective:

$$\min_{\theta(x)} \sum_{\xi} w_{\rho}(\xi - x) \left[| (r_{\theta(x)} \circ m_1)(\xi) - m_2(\xi) |^2 + | (r_{\theta(x)} \circ m_2)(\xi) - m_1(\xi) |^2 \right] \quad (7)$$

where $w_{\rho}(x)$ is positive and integrates to 1, with spatial extent ρ . This allows one to estimate θ at each voxel independently, while the overlapping neighborhoods ensure that $\theta(x)$ varies smoothly.

This approach also provides a natural way to allow for variable neighborhood sizes, where $\rho(x)$ depends on location x , so both rigid regions and transition regions are well modeled. Notably, we want $\rho(x)$ to be large enough to reliably estimate the local power of the CV residual error to estimate $\theta(x)$ correctly, but small enough to enable rapid local transitions. A reasonable balance can be specified in terms of the highest frequency with substantial signal power, which is captured by the regularization parameter $\theta(x)$ itself. In particular, for regularizers that are close to optimal, we expect the residual signal to have its power concentrated at wavelengths near $\theta(x)$. In this case a good measure of the local residual power is to aggregate the squared residual over a small number of wavelengths $\theta(x)^{42}$. Thus we can reliably estimate both $\theta(x)$ and $\rho(x)$ as long as $\rho(x)$ is constrained to be a small multiple of $\theta(x)$.

We therefore adopt a simple heuristic, that $\rho(x) > \gamma \theta(x)$ where γ , the adaptive window factor (AWF), is a constant. With this constraint we obtain the final form of the computational problem solved in non-uniform refinement to regularize 3D electron density at each iteration; that is,

$$\begin{aligned} \theta^*(x) = & \arg \min_{\theta(x)} \min_{\rho(x)} \sum_{\xi} w_{\rho(x)}(\xi - x) \\ & \times \left[| (r_{\theta(x)} \circ m_1)(\xi) - m_2(\xi) |^2 \right. \\ & \left. + | (r_{\theta(x)} \circ m_2)(\xi) - m_1(\xi) |^2 \right] \text{ s.t. } \rho(x) > \gamma \theta(x) \end{aligned} \quad (8)$$

We find that as long as $\gamma > 3$ we obtain reliable estimates of the local power of the residual signal. For $\gamma < 2$, estimates of residual power are noisy and optimization of the regularization parameters therefore suffers. The algorithm is relatively insensitive to values of $\gamma > 3$, but there is some loss in the spatial resolution of the adaptive regularizer as γ increases.

Non-uniform refinement algorithm. Algorithm 2. Regularization step for non-uniform refinement

Require Particle image dataset \mathcal{D} with pose estimates z	
1:	Randomly partition \mathcal{D} into halves, \mathcal{D}_1 and \mathcal{D}_2 with corresponding poses z_1 and z_2
2:	Reconstruct \tilde{m}_1 and \tilde{m}_2 , the raw (noisy) 3D maps from each half-set
3:	Estimate regularization parameters θ' by solving equation (8)
4:	Reconstruct a single map from \mathcal{D} , z and apply the optimal regularizer $r_{\theta'}$

Given a set of particle images and a low-resolution ab initio 3D map, non-uniform refinement comprises three main steps, similar to conventional uniform (homogeneous) refinement (Algorithm 1). The data are randomly partitioned into two halves, each of which is used to independently estimate a 3D half-map. This 'gold-standard' refinement¹⁷ allows use of FSC for evaluating map quality, and for comparison with existing algorithms. The alignment of particle images against their respective half-maps, and the reconstruction of the raw 3D density map (the E and M steps in Algorithm 1) are also identical to uniform refinement.

The difference between uniform and non-uniform refinement is in the regularization step. First, in non-uniform refinement, regularization is performed independently in the two half-maps. As such, the estimation of the spatial regularization parameters in Algorithm 2 effectively partitions each half-dataset

into quarter-datasets. We often refer to the raw reconstructions in Algorithm 2 as quarter-maps. The non-uniform refinements on half-maps are therefore entirely independent, satisfying the assumptions of a 'gold-standard' refinement¹⁷. In contrast, conventional uniform refinement uses FSC between half-maps to determine regularization parameters at each iteration, thereby sharing masks and regularization parameters, both of which contaminate final FSC-based assessment because the two half-maps are no longer reconstructed independently.

Most importantly, non-uniform refinement uses equation (8) to define the optimal parameters with which to regularize each half-set reconstruction at each refinement iteration. Figure 2 shows an example of the difference between uniform filtering (FSC-based) and the new CV-optimal regularizer used in non-uniform refinement. Uniform regularization removes signal and noise from all parts of the 3D map equally. Adaptive regularization, on the other hand, removes more noise from disordered regions, while retaining the high-resolution signal in well-structured regions that is critical for aligning 2D particle images in the next iteration.

In practice, for regularization parameter estimation, equation (8) is solved on a discretized parameter space where a relatively simple discrete search method can be used (e.g. as opposed to continuous gradient-based optimization). The algorithm is implemented in Python within the cryoSPARC software platform⁶, with most of the computation implemented on graphics processing unit accelerators. An efficient solution to equation (8) is important in practice because this subproblem is solved twice for each iteration of a non-uniform refinement.

Finally, the tuning parameters for adaptive regularization are interpretable and relatively few in number. They include the order of the Butterworth kernel, the discretization of the parameter space and the scalar relating $\rho(x)$ and $\theta(x)$, called the AWF. In all experiments, we use an eighth-order Butterworth filter and a fixed AWF parameter $\gamma = 3$. We discretize the regularization parameters into 50 possible values, equispaced in the Fourier domain to provide greater sensitivity to small-scale changes at finer resolutions. We find that non-uniform refinement is approximately two times slower than uniform refinement in our current implementation.

Over-fitting of regularizer parameters. As mentioned in the Discussion, local resolution estimates or local statistical tests have commonly been used to adaptively filter 3D maps. While these methods are generally satisfactory as one-time post-processing steps for visualization, in our experience they can lead to severe over-fitting when used iteratively within refinement as a substitute for regularization. (Supplementary Fig. 1 illustrates one example with the Na_{1.7} dataset.) During iterative refinement, small mis-estimations of local resolution at a few locations (due to high estimator variance²³) cause subtle over- or under-fitting, leaving slight density variations. Over multiple iterations of refinement, these errors can produce strong erroneous density that contaminate particle alignments and the local estimation of resolution itself, creating a vicious cycle. A related technique using iterative local resolution and filtering was described briefly in EMAN2.2 documentation⁹ and may suffer the same problem. The resulting artifacts (e.g. streaking and spikey density radiating from the structure) are particularly prevalent in datasets with junk particles, structured outliers or small particles that are already difficult to align. To mitigate these problems, the approach we advocate couples an implicit resolution measure to a particular choice of local regularizer, with optimization explicitly designed to control model capacity and avoid over-fitting of regularizer parameters.

Data pre-processing. Experimental results for STRA6-CaM and PfCRT datasets were computed directly from particle image stacks, with no further pre-processing. The original data for the Na_{1.7} protein comprise 25,084 raw microscope movies (EMPIAR-10261) from a Gatan K2 Summit direct electron detector in counting mode, with a pixel size of 0.849 Å. We processed the dataset through motion correction, CTF estimation, particle picking, 2D classification, ab initio reconstruction and heterogeneous refinement in cryoSPARC v2.11 (ref. ⁶). A total of 738,436 particles were extracted and then curated using 2D classification yielding 431,741 particle images (pixel size 1.21 Å, box size 360 pixels). As described elsewhere²⁵, we detected two discrete conformations corresponding to the active and inactive states of the channel, with 300,759 and 130,982 particles. We obtained reconstructions with resolutions better than the published literature for both states, but for the results in this work we focus solely on the active state (refined with C2 symmetry yielding 601,518 asymmetric units).

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

No new datasets were created in this study. The raw datasets analyzed in this study were either downloaded from the EMPIAR repository (EMPIAR-10024, EMPIAR-10330, EMPIAR-10261) or were provided by authors of other studies, cited in the main text. Density maps and atomic coordinates from EMDB-0341 and Protein Data Bank 6N4Q were used for visualization. All other results and outputs of data analysis in this study are available from the corresponding author on reasonable request.

Code availability

The cryoSPARC software package is freely available in executable form for nonprofit academic use at www.cryosparc.com.

References

39. Bepler, T., Noble, A. & Berger, B. Topaz-Denoise: general deep denoising models for cryoEM. *Nat. Commun.* **11**, 5208 (2020).
40. Tegunov, D. & Cramer, P. Real-time cryo-electron microscopy data preprocessing with warp. *Nat. Methods* **16**, 1146–1152 (2019).
41. Buchholz, T.-O., Jordan, M., Pigino, G. & Jug, F. cryoCARE: content-aware image restoration for cryo-transmission electron microscopy data. Preprint at *arXiv* <https://arxiv.org/abs/1810.05420> (2018).
42. Potamianos, A. & Maragos, P. A comparison of the energy operator and the Hilbert transform approach to signal and speech demodulation. *Signal Process.* **37**, 95–120 (1994).

Acknowledgements

We are extraordinarily grateful to O. Clarke, F. Mancia and Y. Zi Tan for providing valuable cryo-EM data early in this project and for sharing their experience as early adopters of non-uniform refinement and cryoSPARC. We thank the entire team at Structura Biotechnology Inc., which designs, develops and maintains the cryoSPARC software system in which this project was implemented and tested. We thank J. Rubinstein for comments on this manuscript. Resources used in this research were provided, in part, by the Province of Ontario, the Government of Canada through

NSERC (Discovery Grant RGPIN 2015-05630 to D.J.F.) and CIFAR and companies sponsoring the Vector Institute.

Author contributions

A.P. and D.J.F. designed the algorithm, A.P. and H.Z. implemented the software and performed experimental work. A.P. and D.J.F. contributed expertise and supervision. D.J.F. and A.P. contributed to manuscript preparation.

Competing interests

A.P. is CEO of Structura Biotechnology, which builds the cryoSPARC software package. D.J.F. is an advisor to Structura Biotechnology. Novel aspects of the method presented are described in a patent application (WO2019068201A1), with more details available at <https://cryosparc.com/patent-faqs>.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41592-020-00990-8>.

Correspondence and requests for materials should be addressed to A.P. or D.J.F.

Peer review information Allison Doerr was the primary editor on this article and managed its editorial process and peer review in collaboration with the rest of the editorial team.

Reprints and permissions information is available at www.nature.com/reprints.

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see [Authors & Referees](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

- | | |
|-------------------------------------|---|
| n/a | Confirmed |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> The statistical test(s) used AND whether they are one- or two-sided
<i>Only common tests should be described solely by name; describe more complex techniques in the Methods section.</i> |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> A description of all covariates tested |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals) |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
<i>Give P values as exact values whenever suitable.</i> |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes |
| <input checked="" type="checkbox"/> | <input type="checkbox"/> Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated |

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection

We did not perform data collection in this study.

Data analysis

This study is centered around new algorithms that were developed for data analysis. These algorithms are implemented in the cryoSPARC software package, and this implementation was used for validation in the study. CryoSPARC is freely available for academic use, and is available via licensing for commercial use, at www.cryosparc.com. The version of cryoSPARC used in this study will be made available to editors/reviewers according to Nature Research guidelines.

Versions of software package and tools used:

cryoSPARC v2.11

Python 2.7

Chimera 1.4 and ChimeraX 0.9

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

No new datasets were created in this study. The raw datasets analyzed in this study were either downloaded from the EMPIAR repository (EMPIAR-10024, EMPIAR-10330, EMPIAR-10261) or were provided by authors of other studies, cited in the main text.

The following PDB and EMDB datasets were used: EMD-0341, PDB-6N4Q.

All other results and outputs of data analysis in this study are available from the corresponding author on reasonable request.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/documents/nr-reporting-summary-flat.pdf](https://www.nature.com/documents/nr-reporting-summary-flat.pdf)

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	For each dataset analyzed, the sample size of particle images used in the analysis was determined simply based on how many particle images were captured during microscopy for the datasets that were analyzed. These sample sizes are sufficient as the final resolution of reconstructed 3D protein density maps is statistically validated and results are only interpreted at resolutions where the specific sample size of particle images for a given dataset produces sufficient signal-to-noise ratio.
Data exclusions	No data were excluded.
Replication	All results in the study can be replicated exactly by running the described algorithm/software with the same input data.
Randomization	Randomization was used in the statistical validation of achieved reconstruction resolution for each image dataset that was analyzed. Randomization takes the form of randomly splitting a dataset into two halves, so that each can be reconstructed independently. This process is explicitly described in the main text.
Blinding	Blinding during data analysis was not relevant as human input is not used in the analysis of the data. Blinding during data collection was not relevant as no new raw data was collected for this study.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> Antibodies
<input checked="" type="checkbox"/>	<input type="checkbox"/> Eukaryotic cell lines
<input checked="" type="checkbox"/>	<input type="checkbox"/> Palaeontology
<input checked="" type="checkbox"/>	<input type="checkbox"/> Animals and other organisms
<input checked="" type="checkbox"/>	<input type="checkbox"/> Human research participants
<input checked="" type="checkbox"/>	<input type="checkbox"/> Clinical data

Methods

n/a	Involved in the study
<input checked="" type="checkbox"/>	<input type="checkbox"/> ChIP-seq
<input checked="" type="checkbox"/>	<input type="checkbox"/> Flow cytometry
<input checked="" type="checkbox"/>	<input type="checkbox"/> MRI-based neuroimaging