# Data-driven regularisation lowers the size barrier of cryo-EM structure determination

Dari Kimanius[1,2*], Kiarash Jamali[1], Max E Wilkinson[3,4,5], Sofia Lövestam[1], Vaithish Velazhahan[1,6], Takanori Nakane[7], Sjors H.W. Scheres[1*]

[1] MRC Laboratory of Molecular Biology, Francis Crick Avenue, Cambridge, CB2 0QH, UK.
[2] CZ Imaging Institute, 3400 Bridge Parkway, Redwood City, CA 94065, USA.
[3] Broad Institute of MIT and Harvard, Cambridge, MA 02142, USA.
[4] McGovern Institute for Brain Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA.
[5] Howard Hughes Medical Institute, Cambridge, MA 02139, USA.
[6] School of Medicine, Stanford University, 291 Campus Drive, Li Ka Shing Building, Stanford, CA 94305-5101, USA.
[7] Institute for Protein Research, Osaka University, Osaka, 3-2 Yamadaoka, Suita-shi, Osaka 565-0871, Japan.

*Corresponding author(s). E-mail(s): dari@mrc-lmb.cam.ac.uk; scheres@mrc-lmb.cam.ac.uk;

**Abstract**

Macromolecular structure determination by electron cryo-microscopy (cryo-EM) is limited by the alignment of noisy images of individual particles. Because smaller particles have weaker signals, alignment errors impose size limitations on its applicability. Here, we explore how image alignment is improved by the application of deep-learning to exploit prior knowledge about biological macromolecular structures that would otherwise be difficult to express mathematically. We train a denoising convolutional neural network on pairs of half-set reconstructions from the electron microscopy data bank (EMDB) and use this denoiser as an alternative to a commonly used smoothness prior. We demonstrate that this approach, which we call Blush regularisation, yields better reconstructions than existing

1

algorithms, in particular for data with low signal-to-noise ratios. The reconstruction of a protein-nucleic acid complex with a molecular weight of 40 kDa, which was previously intractable, illustrates that regularisation through denoising will expand the applicability of cryo-EM structure determination for a wide range of biological macromolecules.

# Introduction

Despite rapid progress in the past decade [1], many biological macromolecules of interest are still too small to allow reliable cryo-EM structure determination. To limit the damage that electrons cause to the biological structures of interest, cryo-EM images are taken under low-dose conditions, which leads to high levels of experimental noise. The noise in the images impedes their alignment, resulting in an ill-posed optimisation problem, where many (noisy or artifactual) reconstructions are equally likely given the data. Nevertheless, the correct solution may still be identified through the incorporation of prior knowledge in the reconstruction process. Most cryo-EM structures are calculated using explicit regularisation of a likelihood function in Fourier space, which assumes cryo-EM reconstructions are smooth in real space [2–4].

Although we know a lot more about the structures of biological macromolecules than that their density varies smoothly, it has been difficult to incorporate richer sources of prior knowledge in the optimisation. Denoising convolutional neural networks provide a mechanism to incorporate complicated prior knowledge into an iterative optimisation process [5]. By training a denoising network on simulated pairs of noisy and ground-truth images, we previously provided a proof-of-principle that prior knowledge about protein structures can be exploited to improve cryo-EM structure determination [6]. However, we also observed problems with overfitting and the hallucination of protein-like features in the resulting reconstructions. Moreover, because experimental cryo-EM structures often comprise regions of well-ordered proteins and nucleic acid domains alongside less structured regions, including for example membrane patches or flexible domains, it was not clear how one would generate ground-truth pairs for experimental cryo-EM data.

Here, we demonstrate how denoising convolutional neural networks, trained and deployed in an application-specific manner, can improve cryo-EM structure determination with experimental data. Similar to the noise2noise approach [7], we trained a denoiser using pairs of noisy reconstructions that are generated as part of standard procedures in cryo-EM structure determination [8]. This approach improves cryo-EM structure determination of a wide variety of biological macromolecules and reduces current size limitations of cryo-EM structure determination.

# Method

## Rationale

The noise2noise framework [7] facilitates the training of a denoising convolutional neural network in the absence of explicit access to ground-truth images. Instead, it relies on pairs of noisy images to extract information about their shared signal. Here, we present an application-specific approach that incorporates this aspect from the noise2noise framework. We trained a denoiser on a set of 422 pairs of noisy half-maps that we downloaded from the electron microscopy data bank (EMDB) [9]. Only entries with reported resolutions higher than 4 Å that had both unfiltered half-maps deposited where selected.

We tailored data augmentation and training of the denoiser to integrate with the iterative expectation-maximisation algorithm for cryo-EM reconstruction. All pairs of half-maps, $x_i^{(k)} \in \mathbb{R}^N$, with $k \in \{0, 1\}$, were re-scaled to a uniform voxel size of 1.5 Å, and augmented by generating new pairs $y_i^k, \bar{y}_i^k \in \mathbb{R}^N$:

$$y_i^{(k)} = \mathrm{H}_{C,A} \left[ x_i^{(1-k)} + e^{(1-k)} \right], \tag{1}$$

$$\bar{y}_i^{(k)} = \mathrm{H}_{\bar{C},A} \left[ x_i^{(k)} \odot M_i + \mathrm{h} \left( x_i^{(k)} \right) \odot (1 - M_i) \right], \tag{2}$$

where, $e \in \mathbb{R}^N$ is random coloured noise, $M_i \in [0, 1]^N$ is a smooth mask encapsulating the molecules of interest, $\odot$ represents voxel-wise multiplication, and h(.) is a low-pass filter to 15 Å. $\mathrm{H}_{C,A}[.]$ applies an anisotropic Gaussian filter with covariance matrix $C$, an affine transform $A$ that includes rotation and translation, a crop to a cube of $64^3$ voxels, and a voxel-value standardisation. Data augmentation was achieved through random assignments of $C$, $\bar{C}$, $A$, $e$, and $r$.

By using a range of resolution cutoffs for $C$ and $\bar{C}$, the denoiser explicitly learns to handle maps with varying resolutions. This is needed for its application inside the iterative expectation-maximisation algorithm, which typically starts at relatively low resolutions and gradually progresses to higher resolutions. Although using a lower resolution cut-off for $C$ than for $\bar{C}$ could have produced a network that enhances the resolution of the half-maps, similar to deblurring networks [10], we opted not to do so in order to minimize the risk of hallucinating high-resolution features.

Using different degrees of anisotropy in $C$ and $\bar{C}$, the denoiser learns to deal with the artifacts that arise from non-uniform orientational distributions, while random orientations and affine transformations in $A$ lead to invariance with respect to rotations, translations, and intensity scale. By applying h(.) on $\bar{y}_i^k$, the denoiser learns to enforce fewer features in the solvent region. By filling solvent regions with a 15 Å low-pass filtered version of the map, as opposed to a straightforward voxel-wise multiplication with the mask $M_i$, higher density values in regions with disordered molecules, such as detergent micelles, are maintained.

3

## Training the denoiser

Our denoiser $f_\theta$ consists of a U-net with approximately 13 million trainable parameters $\theta$ (Figure 1). It is trained using residual learning [11] and with a dropout rate of 50% [12]. Instance normalization [13] is used to handle small mini-batches $\mathcal{B}$, with $b = 8$ samples from the training dataset, during training. We minimise the following loss:

$$\mathcal{L} = \frac{1}{2b} \sum_{i \in \mathcal{B}} \sum_{k \in \{0,1\}} \left\| \bar{y}_i^{(k)} - f_\theta \left( \mathrm{R}_r \left[ f_\theta, y_i^{(k)} \right] \right) \right\|^2, \tag{3}$$

where, $\mathrm{R}_r[f_\theta, y]$ returns the output of the denoiser $f_\theta$ after recursively calling it $r \in \{0, \ldots, 5\}$ times with $y_i^k$ as the initial input. This enables the denoiser to recognize and suppress artifacts brought about by its repeated usage, thereby limiting the amplification of artifacts in the reconstruction that are introduced by the denoiser during subsequent iterations of the expectation-maximisation algorithm [6].

Training for 950,000 steps took 6 days on a single Nvidia A100 GPU.

## Iterative denoising with spectral trailing

Our pre-trained denoiser expresses prior knowledge about cryo-EM reconstructions. Its iterative use within the expectation-maximisation algorithm that underlies cryo-EM structure determination is called *Blush regularisation*. Blush regularisation has been implemented in the open-source software RELION-5. It can be used for 3D classification, multi-body refinement, and 3D auto-refinement jobs, including particles with point-group or helical symmetry.

At every iteration, the denoiser is applied separately to each intermediate half-map reconstruction, where it replaces the 3D Wiener filter that results from regularisation in Fourier space [2, 3]. Although one effect of the denoiser is that it tends to dampen Fourier components at higher spatial frequencies, the amount by which it does so is not well defined. Therefore, we employ a heuristic, here referred to as *spectral trailing*, to prevent overfitting in 3D auto-refinement and multi-body refinement. First, we calculate the Fourier shell correlation (FSC) between two independently refined half-maps before the application of the denoiser and determine the resolution $\rho$ where the solvent-corrected FSC drops below 0.143. We then apply the denoiser to both half-maps and subsequently apply a low-pass filter at a resolution that is 2 Fourier shells lower than $\rho$. If $\rho$ exceeds the Nyquist frequency of the denoiser, here set to 3 Å, the remaining Fourier shells at higher frequencies are populated with the reconstruction from the standard regularisation in Fourier space. The resulting denoised, low-pass filtered maps are then used as references for alignment in the next iteration. The denoiser is not applied to the output of the final refinement step.

For 3D classification, where data is not separated into independent half-sets, the filtered map from the regularised likelihood approach is used as input for the denoiser. No additional low-pass filtering is applied. In this job type, the denoiser is also applied in the last iteration.

4
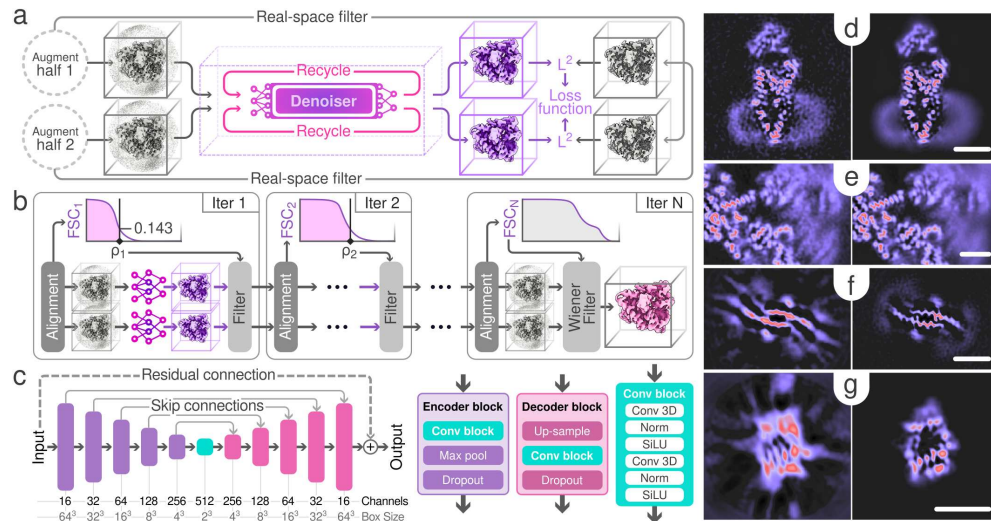
**Fig. 1**: Schematic illustration of Blush and slices of example volumes. **(a)** Training procedure, showing two passes for both half-maps and recycling ($R_r$) of the denoiser output (in pink). **(b)** Iterative reconstruction with spectral trailing. Each half-map is reconstructed separately. At each iteration, the Fourier Shell Correlation (FSC) is used to estimate a cutoff frequency $\rho$, which is subsequently used to low-pass filter the denoiser output. The final output does not pass through the denoiser but is subject to a Wiener filter, similar to baseline reconstruction. **(c)** Denoiser U-net architecture, consisting of five consecutive encoder blocks, a convolution block, followed by five consecutive decoder blocks. **(d-e)** Slices through maps before (left) and after (right) a single application of the denoiser to the final iteration of the reconstruction for PfCRT **(d)** and the Spliceosome **(e)**. **(f-g)** Showing slices through maps of baseline reconstruction (left) and Blush regularisation (right) of the FIA **(f)** and a 40kDa protein-nucleic acid complex **(g)**. The scale bar indicates 30 Å in all slices.

# Results

## Blush improves reconstruction without overfitting

We first tested Blush regularisation on a cryo-EM data set (EMPIAR-10330) [14] of the *Plasmodium falciparum* Chloroquine Resistance Transporter (PfCRT) [15]. This data set has been used as a standard to demonstrate the performance of several approaches to reduce overfitting in cryo-EM refinement [16, 17]. Standard refinement using regularized likelihood optimisation in RELION, which we refer to as the baseline, yields an overall resolution of 3.8 Å for this data set.

Application of Blush regularisation (Figure 2) yielded an overall resolution estimate of 3.4 Å. Spectral trailing was applied with a cut-off at 3.5 Å, beyond which no information from the denoiser was used. Compared to the baseline reconstruction, local resolution improved for most regions of the map, with a corresponding increase

5
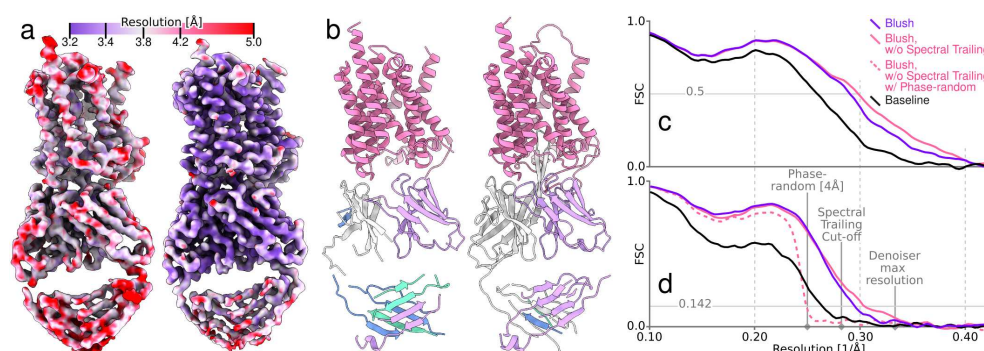
**Fig. 2**: Single-particle reconstruction of the PfCRT dataset. **(a)** Maps coloured by local resolution, comparing baseline (left) and Blush (right). **(b)** Automated atomic modeling by ModelAngelo for the baseline (left) and Blush (right) maps. Colored by chain. **(c)** FSCs between the masked maps and deposited model (ID: 6UKJ). **(d)** Solvent-corrected half-map FSCs. Both plots show FSCs for Blush (purple), Blush without spectral trailing (pink) and baseline (black). Dashed (pink) line show solvent-corrected half-map FSC for Blush without spectral trailing when applied to data with phase randomisation beyond 4 Å resolution.

in visible side-chain densities. The improvement in resolution as measured by half-map FSC was confirmed by FSCs between both maps and the atomic model that was deposited for this dataset (PDB-ID 6UKJ). Throughout this paper, FSCs between the map and atomic model are calculated using Servalcat [18]. In addition, we also assessed the relative quality of both maps by application of our automated model-building software ModelAngelo [19], which generated a model with 84% completeness in the baseline map and 97% completeness in the Blush map. Model completeness is defined as the percentage of residues that match the reference model with a C$\alpha$-distance of 3 Å or less.

To assess the potential for overfitting by the denoiser, we also performed a phase-randomisation test [20]. We applied Blush regularisation without spectral trailing for refinement of the PfCRT dataset with phase-randomisation beyond 4 Å. Even though spectral trailing was not used, no overfitting was observed. Switching off spectral trailing led to a marginal improvement in the quality of reconstruction, as quantified by the FSC between the map and the atomic model (Fig 2d). These results indicate that the denoiser can prevent overfitting for this dataset, even without spectral trailing. In the general ase, we still recommend running Blush regularisation with spectral trailing because the gains of switching it off are small and overfitting may be more prominent for other datasets. Consequently, in the following sections, we only present results obtained using spectral trailing.

# Blush expands the applicability of cryo-EM structure determination

We subsequently assessed the broader applicability of Blush regularisation by applying it to four different types of structures and refinement methods.

First, we tested Blush regularisation on a small membrane protein, Ste2, which is a dimeric G-protein coupled receptor (GPCR) [21] (Figure 3, Extended Data Table 1). The full-length monomeric Ste2 has a molecular weight of 47.85 kDa, which includes a long disordered C-terminal tail that comprises 125 amino acids. The total mass of the ordered dimeric Ste2 that contributes to alignment is roughly 67 kDa, most of which lies embedded in a detergent micelle.

The data set was acquired on a similar complex to PDB entry 7QB9 reported in [21], but with different biochemical conditions that impact the stability of the structure. Alignment of images of Ste2 is difficult because not many protein features extend from the smooth detergent micelle. Baseline reconstruction yielded a map with an overall resolution of 3.8 Å, with limited densities for side chains. Application of Blush regularisation led to a structure with an overall resolution of 3.4 Å. Spectral trailing ensured that no information from the denoiser was inserted beyond 3.7 Å resolution. Compared to the baseline reconstruction, the density of the trans-membrane helices is improved. Loops at the top and bottom of the structure are still relatively poorly resolved, probably due to molecular flexibility. In agreement with the visibility of improved side-chain densities and local resolution estimates, the completeness of models built by ModelAngelo in these maps improved from 19% to 43%.
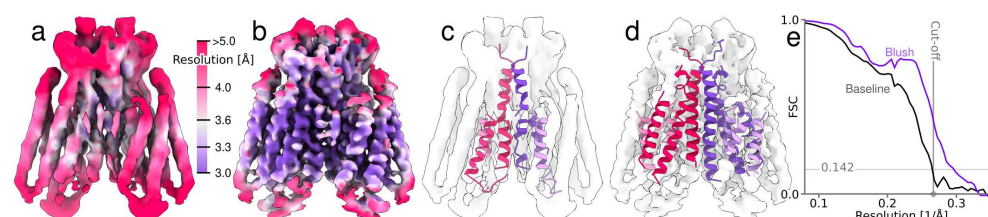


**Fig. 3**: Single-particle reconstruction of the Ste2 dataset. Reconstructions coloured by local resolution, comparing baseline **(a)** and Blush **(b)**. Automated atomic modeling by ModelAngelo, using the baseline **(c)** and Blush **(d)** maps. **(e)** Solvent-corrected half-map FSCs.

Second, we tested Blush regularisation in multi-body refinement [22]. Multi-body refinement uses partial signal subtraction to align independently moving domains of a larger complex. Reconstructions from subtracted images were not part of the training set of the denoiser. Moreover, signal subtraction reduces the amount of signal in each image, placing stringent limitations on the minimal size of domains that can be aligned in multi-body refinement. We applied Blush in multi-body refinement to a publicly available data set (EMPIAR-10180) of the *Saccharomyces cerevisiae* pre-catalytic spliceosomal B complex [23] (Figure 4). Using four bodies, for the core, the foot, the helicase and the SF3b regions, Blush regularisation improved all domains compared
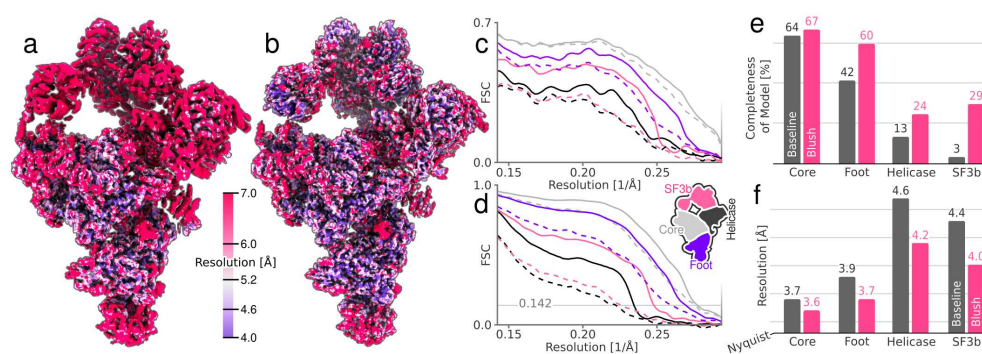
**Fig. 4**: Multibody reconstruction of the spliceosome dataset. Combined maps of the individual bodies, coloured by local resolution, comparing baseline **(a)** and Blush **(b)**. **(c)** FSCs between the masked maps of each body and the corresponding region in the deposited model (ID: 5NRL). **(d)** Solvent-corrected half-map FSCs for the individual bodies. In **(c-d)** dashed and solid lines correspond to baseline and Blush maps, respectively. FSCs are shown for each body: core (grey), foot (black), helicase (purple) and SF3b (red). **(e)** Completeness of atomic models built by ModelAngelo for each body, using baseline (grey) and Blush (red) maps. **(f)** Gold-standard half-map resolutions of each body for baseline (grey) and Blush (red) maps.

to baseline multi-body refinement, as measured by local resolution, half-map FSCs and FSCs with the reference atomic model (PDB-ID 5NRL). The improvements in resolution were largest in the helicase and SF3b regions, which are the most flexible and thus the hardest to reconstruct. The improvements in resolution were reflected by automated model building in ModelAngelo, which increased model completeness of the entire complex from 32% to 48%. In particular, the model completeness for the SF3b region was improved from 3% to 29%.

Third, we assessed the performance of Blush regularisation for a biological assembly that differed from the types of structures the denoiser was trained on: the first intermediate amyloid (FIA) that forms during the *in vitro* assembly of recombinant tau (residues 297-391) [24]. This data set is also publicly available (EMPIAR-11720). Unlike any of the structures in the training set, the FIA has helical symmetry. It is an amyloid filament, with parallel $\beta$-strands repeating every 4.7 Å in the direction of the helical axis. Besides deviating from the types of structures in the training set, the FIA is also one of the smallest amyloid structures solved to date, with only 15 ordered residues in each of two opposing $\beta$-sheets. Baseline helical refinement yielded a 5.0 Å map, in which the density for $\beta$-strands along the helical axis was not separated, and no atomic model could be built. Blush regularisation improved the resolution to 2.8 Å, and ModelAngelo built all 15 ordered residues in the resulting map.

Fourth, we applied Blush to a protein-nucleic acid complex with a combined molecular weight of 40 kDa (Figure 6; Extended Data Table 1). Using multiple different classification and refinement strategies in baseline RELION and CryoSPARC, we were unable to obtain a reliable reconstruction (results not shown). Even though an initial model generated using the standard VDAM algorithm in RELION [25] suffered from
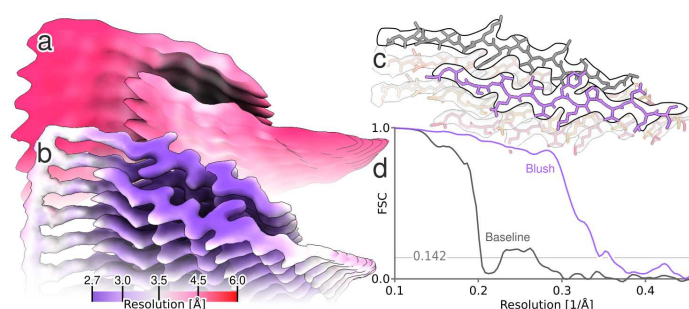
**Fig. 5**: Helical reconstruction of the FIA, coloured by local resolution, for the baseline **(a)** and Blush **(b)** maps. **(c)** Automated atomic modeling by ModelAngelo, comprising tau residues 302-316. **(d)** Solvent-corrected half-map FSCs of the reconstructed maps.

anisotropy, a first 3D classification with Blush regularisation resulted in one class with recognizable protein features. Refinement of the corresponding class yielded a better initial model for a second 3D classification, from which a single class was selected for subsequent CTF refinement [26] and particle polishing [26]. A 3D classification was performed without alignment, followed by a final 3D refinement. All 3D classifications with alignment and 3D refinements used Blush regularisation. The final map achieved a resolution of 2.5 Å, with ModelAngelo successfully building 97% of the protein sequence, and 33 out of 42 nucleotides.

# Discussion

Our results demonstrate that denoising convolutional neural networks can be used as a source of extended prior information in the reconstruction of experimental cryo-EM images. A previous approach that used noise2noise, implemented in the M software [27], exclusively trained a denoising network on half-maps from the dataset to which it was also applied. As a result, unlike our approach, it did not express general prior knowledge about cryo-EM reconstructions. When we tried to express prior knowledge about protein structures by training a denoiser on pairs of noisy and ground-truth maps that were calculated from atomic models, we observed problems with overfitting and hallucinations [6]. Similar problems may also explain why the application of the DeepEMhancer neural network [28] inside the iterative reconstruction algorithm of RELION had to be restricted to only a few iterations at the end of refinement [29]. The approach described in this paper reduces the risk of hallucinating protein-like features in the reconstruction by using a neural network that is trained on experimental cryo-EM half-maps only, i.e. without using atomic models or the geometrical restraints that are used to describe them.

Instead of forcing the map to resemble densities derived from atomic models, our denoiser is trained to introduce more subtle modifications to cryo-EM maps, such as smoothing out density in solvent regions or within detergent micelles. The network also removes artifacts that are commonly encountered in difficult cryo-EM refinements, e.g. anisotropic densities that result from uneven angular distributions, or radially

9
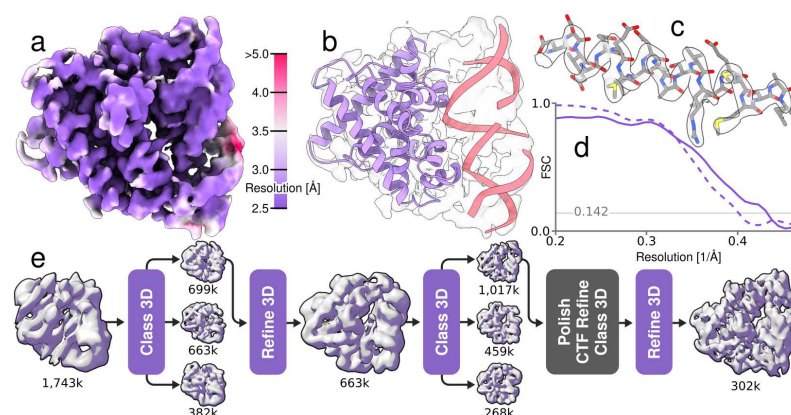
**Fig. 6**: Single-particle reconstruction of the protein-nucleic acid complex with a molecular weight of 40 kDa. **(a)** Blush regularisation local resolution. **(b)** Automated atomic model assignment with ModelAngelo. **(c)** Detailed view of an alpha-helical segment in the reconstructed map and refined atomic model. **(d)** Solid line shows FSC to a reference atomic model and the dashed line shows half-map FSC. Solvent-corrected resolution is 2.8 Å using a spectral trailing cut-off at 3.0 Å. **(e)** Processing pipeline from initial model to final reconstruction. Numbers indicate the number of particles assigned to each map. Purple squares indicate reconstructions using Blush regularization.

extending, streaky features that are often observed in overfitted maps (Figs 1f-g). Our findings illustrate that, even though the effect of a single application of the denoiser is relatively small, its cumulative impact over multiple iterations enhances the performance of cryo-EM structure determination across a diverse range of test cases. As machine-learning methods get better at extracting knowledge from large datasets, it may be tempting to exploit more knowledge about the structures of biological macro-molecules in the reconstruction process. However, doing so may ultimately also take away one of the most powerful ways of assessing whether a reconstruction is correct: the presence of expected features in the map. We thus anticipate that the cryo-EM community will continue to explore the question of how much prior knowledge should inform the reconstruction process and how much should be kept aside for validation.

Within the framework of Blush regularisation, the denoiser replaces the exponential prior that traditionally constrains the power of Fourier-space components in the baseline algorithm. As a result, the FSC between independently refined subsets no longer influences the weighting of Fourier components in intermediate reconstructions. Instead, this FSC is used to determine a resolution cutoff $\rho$, beyond which the Fourier components of the two denoised half-maps are set to zero. Because Fourier components near the resolution estimate of the final map will not have been affected by the denoiser, protein-like features at these frequencies cannot be the result of hallucination or overfitting by the denoiser. Future investigations will explore alternative approaches that do not rely on heuristics, like choosing a specific $\rho$, which arises from the observation that the dampening effect of the network is not well defined in Fourier

10

space. Future adaptation of the VDAM algorithm [25] will also allow the use of Blush regularisation for initial model generation.

In all our tests, Blush regularisation outperformed or equaled the baseline implementation in RELION, with the differences being largest for cases where the baseline approach tends to overfit the data. Consequently, Blush regularisation will be most useful for refinements of data sets with low signal-to-noise ratios, like those of small complexes or complexes embedded in thick ice layers, multi-body refinements involving relatively small bodies, and refinements of maps exhibiting pronounced variations in local resolution. For example, Blush regularisation allowed reconstruction of an amyloid with only 30 residues in its ordered core and a protein-nucleic acid complex with a molecular weight of 40 kDa. Although nucleic acids result in higher signal-to-noise ratios than proteins, 40 kDa approaches predictions of the minimum size protein that is amenable to cryo-EM structure determination [30, 31]. These results illustrate that the use of prior knowledge through denoising convolutional neural networks expands the applicability of cryo-EM structure determination.

## Author contributions

DK designed and implemented Blush regularization, ran most experiments and analysed the results. KJ contributed to data preprocessing. MEW contributed and analysed the dataset of the 40kDa protein-nucleic acid complex, and contributed to analysis of the PfCRT dataset. SL contributed the FIA dataset. VV contributed the Ste2 dataset. SHWS supervised the project and contributed to RELION integration. All authors contributed to the writing of the manuscript.

## Acknowledgements

11

# Competing interests

The authors declare no competing interests.

# Code availability

Blush regularisation has been implemented in the open-source software RELION-5, which is distributed for free under the GPLv2 license and can be downloaded from https://github.com/3dem/relion.

# References

[1] Kühlbrandt, W.: The resolution revolution. Science **343**(6178), 1443–1444 (2014)

[2] Scheres, S.H.: A bayesian view on cryo-em structure determination. Journal of molecular biology **415**(2), 406–418 (2012)

[3] Scheres, S.H.: Relion: implementation of a bayesian approach to cryo-em structure determination. Journal of structural biology **180**(3), 519–530 (2012)

[4] Punjani, A., Rubinstein, J.L., Fleet, D.J., Brubaker, M.A.: cryosparc: algorithms for rapid unsupervised cryo-em structure determination. Nature methods **14**(3), 290–296 (2017)

[5] Romano, Y., Elad, M., Milanfar, P.: The little engine that could: Regularization by denoising (red). SIAM Journal on Imaging Sciences **10**(4), 1804–1844 (2017)

[6] Kimanius, D., Zickert, G., Nakane, T., Adler, J., Lunz, S., Schönlieb, C.-B., Öktem, O., Scheres, S.H.: Exploiting prior knowledge about biological macromolecules in cryo-em structure determination. IUCrJ **8**(1), 60–75 (2021)

[7] Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., Aila, T.: Noise2noise: Learning image restoration without clean data. arXiv preprint arXiv:1803.04189 (2018)

[8] Scheres, S.H., Chen, S.: Prevention of overfitting in cryo-em structure determination. Nature methods **9**(9), 853–854 (2012)

[9] Lawson, C.L., Patwardhan, A., Baker, M.L., Hryc, C., Garcia, E.S., Hudson, B.P., Lagerstedt, I., Ludtke, S.J., Pintilie, G., Sala, R., *et al.*: Emdatabank unified data resource for 3dem. Nucleic acids research **44**(D1), 396–403 (2016)

[10] Albluwi, F., Krylov, V.A., Dahyot, R.: Image deblurring and super-resolution using deep convolutional neural networks. In: 2018 IEEE 28th International Workshop on Machine Learning for Signal Processing (MLSP), pp. 1–6 (2018). IEEE

[11] Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. IEEE transactions on image processing **26**(7), 3142–3155 (2017)

[12] Hinton, G.E., Srivastava, N., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.R.: Improving neural networks by preventing co-adaptation of feature detectors. arXiv preprint arXiv:1207.0580 (2012)

[13] Ulyanov, D., Vedaldi, A., Lempitsky, V.: Instance normalization: The missing ingredient for fast stylization. arXiv preprint arXiv:1607.08022 (2016)

[14] Iudin, A., Korir, P.K., Somasundharam, S., Weyand, S., Cattavitello, C., Fonseca, N., Salih, O., Kleywegt, G.J., Patwardhan, A.: Empiar: the electron microscopy public image archive. Nucleic Acids Research **51**(D1), 1503–1511 (2023)

[15] Kim, J., Tan, Y.Z., Wicht, K.J., Erramilli, S.K., Dhingra, S.K., Okombo, J., Vendome, J., Hagenah, L.M., Giacometti, S.I., Warren, A.L., *et al.*: Structure and drug resistance of the plasmodium falciparum transporter pfcrt. Nature **576**(7786), 315–320 (2019)

[16] Ramlaul, K., Palmer, C.M., Nakane, T., Aylett, C.H.: Mitigating local over-fitting during single particle reconstruction with sidesplitter. Journal of structural biology **211**(2), 107545 (2020)

[17] Punjani, A., Zhang, H., Fleet, D.J.: Non-uniform refinement: adaptive regularization improves single-particle cryo-em reconstruction. Nature methods **17**(12), 1214–1221 (2020)

[18] Yamashita, K., Palmer, C.M., Burnley, T., Murshudov, G.N.: Cryo-em single-particle structure refinement and map calculation using servalcat. Acta Crystallographica Section D: Structural Biology **77**(10), 1282–1291 (2021)

[19] Jamali, K., Kimanius, D., Scheres, S.H.: A graph neural network approach to automated model building in cryo-EM maps. In: The Eleventh International Conference on Learning Representations (2023)

[20] Chen, S., McMullan, G., Faruqi, A.R., Murshudov, G.N., Short, J.M., Scheres, S.H., Henderson, R.: High-resolution noise substitution to measure overfitting and validate resolution in 3d structure determination by single particle electron cryomicroscopy. Ultramicroscopy **135**, 24–35 (2013)

[21] Velazhahan, V., Ma, N., Vaidehi, N., Tate, C.G.: Activation mechanism of the class d fungal gpcr dimer ste2. Nature **603**(7902), 743–748 (2022)

[22] Nakane, T., Kimanius, D., Lindahl, E., Scheres, S.H.: Characterisation of molecular motions in cryo-em single-particle data by multi-body refinement in relion. elife **7**, 36861 (2018)

[23] Plaschka, C., Lin, P.-C., Nagai, K.: Structure of a pre-catalytic spliceosome. Nature **546**(7660), 617–621 (2017)

[24] Lövestam, S., Li, D., Wagstaff, J.L., Kotecha, A., Kimanius, D., McLaughlin, S.H., Murzin, A.G., Freund, S.M., Goedert, M., Scheres, S.H.: Disease-specific tau filaments assemble via polymorphic intermediates. bioRxiv, 2023–07 (2023)

[25] Kimanius, D., Dong, L., Sharov, G., Nakane, T., Scheres, S.H.: New tools for automated cryo-em single-particle analysis in relion-4.0. Biochemical Journal **478**(24), 4169–4185 (2021)

[26] Zivanov, J., Nakane, T., Forsberg, B.O., Kimanius, D., Hagen, W.J., Lindahl, E., Scheres, S.H.: New tools for automated high-resolution cryo-em structure determination in relion-3. elife **7**, 42166 (2018)

[27] Tegunov, D., Xue, L., Dienemann, C., Cramer, P., Mahamid, J.: Multi-particle cryo-em refinement with m visualizes ribosome-antibiotic complex at 3.5 å in cells. Nature Methods **18**(2), 186–193 (2021)

[28] Sanchez-Garcia, R., Gomez-Blanco, J., Cuervo, A., Carazo, J.M., Sorzano, C.O.S., Vargas, J.: Deepemhancer: a deep learning solution for cryo-em volume post-processing. Communications biology **4**(1), 874 (2021)

[29] Ramírez-Aportela, E., Carazo, J.M., Sorzano, C.O.S.: Higher resolution in cryo-em by the combination of macromolecular prior knowledge and image-processing tools. IUCrJ **9**(5) (2022)

[30] Henderson, R.: The potential and limitations of neutrons, electrons and x-rays for atomic resolution microscopy of unstained biological molecules. Quarterly reviews of biophysics **28**(2), 171–193 (1995)

[31] Dickerson, J.L., Lu, P.-H., Hristov, D., Dunin-Borkowski, R.E., Russo, C.J.: Imaging biological macromolecules in thick specimens: The role of inelastic scattering in cryoem. Ultramicroscopy **237**, 113510 (2022)

14

# Extended Data

**Extended Data Table 1**: Details on cryo-EM data collection and processing of the Ste2 and 40 kDa complex datasets.

| | Ste2 | 40 kDa complex |
|---|---|---|
| **Data collection** | | |
| Microscope | Thermo Scientific Krios G3i | Thermo Scientific Krios G3i |
| Camera | Falcon 4i | Gatan K3 |
| Magnification | 96,000x | 165,000x |
| Voltage (kV) | 300 | 300 |
| Pixel size (Å) | 0.824 | 0.531 |
| Total electron exposure ($e^-/\text{Å}^2$) | 56 | 72 |
| Exposure rate (e-/pixel/sec) | 7.9 | 15 |
| Defocus range ($\mu$m) | 0.8 - 2.4 | 0.8 - 1.5 |
| Automation software | EPU | EPU |
| Energy filter slit width | N/A | 20 eV |
| Micrographs collected (no.) | 13,254 | 11,232 |
| | | |
| **3D reconstruction** | | |
| Number of frames for motion correction | 70 | 40 |
| Total extracted particles (no.) | 2,608,181 | 6,449,613 |
| Final particle images (no.) | 147,797 | 301,833 |
| Symmetry imposed | C2 | C1 |
| Resolution (FSC 0.143; unmasked/masked) | 3.9 / 3.4 | 2.8 / 2.6 |
| Sharpening B-factor ($\text{Å}^2$) | -160 | -90 |

15