
CESPED: a new benchmark for supervised particle pose estimation in Cryo-EM.

Ruben Sanchez-Garcia^{1,2}, Michael Saur², Javier Vargas³, Carl Poelking², Charlotte M Deane¹

¹Department of Statistics, University of Oxford, Oxford OX1 3LB, UK
{ruben.sanchez-garcia, deane}@stats.ox.ac.uk

²Astex Pharmaceuticals, Cambridge CB4 0QA, UK
{michael.saur, carl.poelking}@astx.com

³Departamento de Optica, Universidad Complutense de Madrid, Madrid 28040, Spain
jvargas@fis.ucm.es

Abstract

Cryo-EM is a powerful tool for understanding macromolecular structures, yet current methods for structure reconstruction are slow and computationally demanding. To accelerate research on pose estimation, we present CESPED, a new dataset specifically designed for Supervised Pose Estimation in Cryo-EM. Alongside CESPED, we provide a PyTorch package to simplify Cryo-EM data handling and model evaluation. We evaluate the performance of a baseline model, Image2Sphere, on CESPED, showing promising results but also highlighting the need for further advancements in this area.

1 Introduction

1.1 Background

Determining the structure of macromolecules is crucial for understanding biological processes and diseases. Cryogenic Electron Microscopy (Cryo-EM) Single-Particle Analysis (SPA) has become a pivotal method in this area. In SPA, the 3D structure of a macromolecule is reconstructed from thousands of noisy 2D image projections, called particles. These particles adopt random orientations during the imaging process that need to be estimated in order to recover the 3D structure.

Traditional algorithms for pose estimation are based on exhaustive comparison of experimental images and reference projections [3, 6, 24, 27, 29]. While they tend to be relatively robust and accurate, they are computationally intensive and slow. In response, deep learning (DL) alternatives have begun to emerge. Some unsupervised DL methods [7, 15, 14] mirror traditional techniques by maintaining a 3D volume representation to compute 2D projections. Unlike traditional methods, which compare each experimental particle against all images in an $SO(3)$ projection gallery with up to millions of members, these DL methods use an encoder to predict the particle orientation based on the particle image alone. Then, the projection for the predicted orientation is differentially rendered and compared to the experimental particle with a loss function used to simultaneously update the parameters of the encoder and the representation of the volume. Other Supervised DL methods [10, 16] offer an even simpler alternative, consisting only of an encoder module that predicts the particle orientation directly from its image. While these approaches offer remarkable efficiency, they require labelled data (pre-aligned particles), thus limiting their applicability in de novo situations. However, since pose estimation is primarily driven by low to mid-resolution frequencies [28], these models could generalize to new experiments with the same sample in different conditions or even with similar samples, accelerating image processing for the subsequent experiments. Since ligand

binding does not generally modify the overall shape of the protein, supervised approaches can be especially valuable in drug discovery, where pre-aligned data for target proteins is often available.

The performance of DL pose estimation models still lags behind current methods in terms of quality and reliability. This gap can be partly attributed to the unique characteristics of Cryo-EM data, which differ from the natural images DL architectures were designed for; and, importantly, to the lack of a standardized benchmark that would allow for a direct comparison of methods in order to stimulate progress, much as ImageNet [5] did for image classification.

1.2 Main contributions.

In this study, we provide an accessible entry point for a wider scientific audience to engage with the challenges of SPA in Cryo-EM. Toward this goal:

- We compiled CESPED (Cryo-EM Supervised Pose Estimation Dataset), an easy-to-use benchmark specifically designed for Supervised Pose Estimation in Cryo-EM.
- We implemented a PyTorch-based [21] package to handle Cryo-EM particle data and to easily compute Cryo-EM quality metrics.
- We trained and evaluated the Image2Sphere model [13], originally developed for real-world pose estimation, on our benchmark. This illustrates the utility of our benchmark and sheds light on the transferability of real-world pose estimation models to the Cryo-EM domain.

1.3 Related work

1.3.1 Traditional Cryo-EM Pose Estimation.

When sample homogeneity can be assumed, the simplest approach to the pose estimation problem is the projection matching algorithm, which consists of T iterations of two steps. First, in the alignment phase, the pose $(R, t)_i$ of each experimental image x_i is set to be the same as the one of the most similar 2D projection of the reference volume V^t , $(R, t)_i = \arg \min_{(R, t) \in SO(3) \times \mathbb{R}^2} \|x_i - g_i * P_{(R, t)} V^t\|^2$ where $P_{(R, t)}$ is the projector operator and g_i is the point spread function of the microscope. Then, in the reconstruction phase, a new volume (in reciprocal space) is computed from the estimated poses as $\hat{V}^{(t+1)} = \left(\sum_{i=0}^N P_{(R, t)_i}^{-1} \hat{g}_i \hat{x}_i \right) / \left(\sum_{i=0}^N P_{(R, t)_i}^{-1} \hat{g}_i^2 + C \right)$ with C being a constant depending on the Signal to Noise Ratio (SNR) and N the number of particles. This iterative process continues until convergence. More advanced methods build on this approach, for example, Relion [27] employs a Bayesian probabilistic model with a prior for the map, while cryoSPARC [24] accelerates the Bayesian methods via branch-and-bound search and gradient descent optimization. See [1] for a review.

1.3.2 Deep learning for pose estimation.

Traditional pose estimation methods were rooted in matching 2D images to 3D objects, often by finding the correspondence between landmarks. Direct pose estimation was first proposed in PoseNet [12], in which a Convolutional Neural Network (CNN) directly regresses the absolute pose of an object represented as quaternions and xyz shifts. Later works improved on PoseNet by refining architectures [18], losses [11] or via multi-task learning [31].

Handling symmetry and ambiguities, such as those caused by occlusions, presents challenges. This issue has been tackled by either breaking the symmetry when labelling the data [30] or applying loss functions that can handle known symmetries [31]. Alternatively, probabilistic models offer approaches that either classify poses within a discretised space or explicitly learn the parameters of probability distributions [4, 17, 19, 20, 23].

In the context of Cryo-EM, only two supervised methods have been proposed to perform direct pose estimation given pre-aligned particles. DeepAlign [10], a set of CNNs that perform binary classification over a discretisation of S^2 , and the approach of Lian et al. [16] who implemented a CNN to perform direct regression of quaternions. However, due to its limitations, especially for symmetric data, Lian et al. finally adopted a hybrid model with a differentiable projector as in some Cryo-EM Unsupervised estimations. Neither of the two has been used in practical scenarios.

2 Methods

2.1 Benchmark Compilation and pre-processing

The EMPIAR database [9] was searched to identify entries containing at least 200K aligned particles for a given macromolecule. Then, to ensure that all particle poses were estimated consistently, the data and metadata files were downloaded and re-processed using Relion [27]. Those entries for which the reconstructed volume showed resolution values close to the reported ones in the literature were selected for the benchmark. Finally, for consistency, all the images were downsampled to 1.5 Å/pixel, with different image dimensions in each entry as the macromolecules vary in size (see Appendix A).

We perform on-the-fly per-image normalisation using the standard Cryo-EM noise whitening procedure so that the background has mean 0 and std 1. We also correct the contrast inversion caused by the defocus via phase flipping [22]. Finally, since the macromolecule typically represents only between 25% to 50% of the whole particle image, the images are cropped so that neighbouring particles are not included. Data labels are represented as rotation matrices and then converted into grid indices by finding the closest rotation matrix in the $SO(3)_{grid}$. For the cases in which the macromolecule exhibits point symmetry, the labels are expanded as $L_i = \{g_j R_i | g_j \in G\}$ where G is the set of rotation matrices given a point symmetry group (e.g., $C1$). As a result, the labels consist of a vector with $|G|$ non-zero values and $|SO(3)_{grid}| - |G|$ zeros.

2.2 Baseline model

We adapted the state-of-the-art Image2Sphere model [13]. Image2Sphere is a hybrid architecture that uses a ResNet to produce a 2D feature map of the input image, which is then orthographically projected onto a 3D hemisphere and expanded in spherical harmonics. Then, equivariant group convolutions are applied, first with global support on the S^2 sphere and finally as a refinement step, on $SO(3)$. The output of the model is a probability distribution over a discretised grid of rotation matrices. Supervised Cryo-EM methods for pose estimation were not considered for this work due to the lack of publicly available code [16] or their GUI requirements [10].

2.3 Evaluation metrics

The most employed metric in pose estimation is the mean angular error, $MAE = \sum_{i=0}^N \text{angError}_i$ where $\text{angError}_i = \min_{g_j \in G} \arccos \left((\text{trace}(g_j \cdot \text{trueR}_i \cdot \text{predR}_i^T) - 1)/2 \right)$, G is the set of rotation matrices given a point symmetry group, and trueR_i and predR_i are the ground-truth and predicted rotation matrices. However, due to the uncertainty in the estimated poses [10], we propose additional metrics. The first one is the confidence-weighted mean-angular-error, $wMAE = \left(\sum_{i=0}^N \text{conf}_i \cdot \text{angError}_i \right) / \left(\sum_{i=0}^N \text{conf}_i \right)$, which weights the angError_i by conf_i , the confidence in the ground-truth pose, measured as Relion’s `rlnMaxValueProbDistribution`. While $wMAE$ is still sensitive to ground-truth and confidence estimation errors, due to its simplicity, we used it as the criterion for hyperparameter tuning.

We also directly assess the quality of the reconstructed volume given the predicted poses by comparing it with the ground-truth volume generated from the original poses. We compute the real space Pearson’s Correlation Coefficient (PCC), and the Fourier Shell Correlation Resolution at threshold 0.5 ($FSC_{0.5}$) and 0.143 ($FSC_{0.143}$), which corresponds to the highest frequency at which the two maps agree with a SNR of 1 and 0.5 respectively [25]. Finally, in order to decouple the different quality levels of the different benchmark entries, we report the differences of the metrics with respect to the ground-truth levels, $\Delta PCC = PCC(GT_0, GT_1) - PCC(GT, V)$ and $\Delta FSC = FSC(GT_0, GT_1) - FSC(GT, V)$, where GT is the ground-truth map, GT_i is the ground-truth map reconstructed with the i -th half of the data, and V is the predicted 3D volume (see Figure 1).

2.4 Training

Each benchmark entry was trained independently with the same hyperparameters (see Appendix B). Due to the uncertainty in the estimated orientations, we employed a weighted cross-entropy loss

using the pose reliability estimate of each particle as the per-image weight

$$L = \frac{1}{N} \sum_{i=0}^N \sum_{c=0}^C -\text{conf}_i \cdot P(R_{c,i}) \log(Q(R_{c,i}))$$

where $Q(R_{c,i})$ is the predicted probability for the rotation matrix with grid index c and $P(R_{c,i})$ is $1/|G|$ when any of the ground-truth matrices is c and zero otherwise.

2.5 Evaluation protocol

Due to the uncertainty in the ground-truth labels and the fact that what matters to the Cryo-EM practitioners is the quality of the reconstructed volume, we devised an evaluation protocol based on the Cryo-EM Gold Standard [8], which is a per-entry 2-fold cross-validation strategy in which the poses of each half of the data are independently estimated and used to reconstruct two volumes (half-maps). For benchmarking supervised methods, it involves training an independent model for each dataset half to infer the poses of the other half of the dataset. After that, the final 3D volume is computed by reconstructing the two half-maps and averaging them. The final averaged map can then be compared with the ground-truth map obtained from the original poses (see Figure 1).

Since Image2Sphere predicts only rotation matrices but not image shifts, when reconstructing the volumes, we employed the ground-truth translations. This could result in an overoptimistic estimation of performance, however, since the effect of the translations is tightly coupled with the accuracy of the angular estimation, this overestimation should be small. We leave for future work the full inference of both the rotations and translations. Finally, to avoid overfitting to the validation set, we performed hyperparameter tuning only over one half of the data using wMANE as metric.

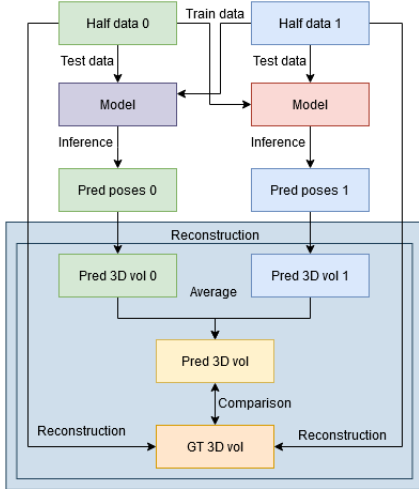


Figure 1: Evaluation protocol based on the Cryo-EM Gold Standard

3 Results and discussion

3.1 Benchmark, ParticlesDataset class and evaluation tool

Our benchmark consists of a diverse set of 7 macromolecules, with an average number of 300K particles including soluble and membrane macromolecules, symmetric and asymmetric complexes, and resolutions ranging from 5 Å to 3.2 Å (see Appendix A). For each particle in the dataset, we provide its image and estimated pose together with an estimate of the reliability of the poses. The benchmark can be automatically downloaded from Zenodo [32] using our cespied Python package.

The package includes a `ParticlesDataset` class, which implements the PyTorch Dataset API for seamless integration. It also offers optional yet recommended pre-processing steps commonly adopted in Cryo-EM and specialized data augmentation techniques, like affine transformations that adjust both the image and its corresponding pose (see Appendix B). Additionally, the package offers an automatic evaluation pipeline that only requires as inputs the predicted poses (grey box in Figure 1). For ease of use, a Singularity image definition file is included, eliminating the need for installing Cryo-EM-specific software like Relion. This design enables those without Cryo-EM experience to utilize the cespied benchmark and package as effortlessly as they would with standard datasets such as MNIST. Usage examples can be found in Appendix C.

3.2 Performance of the baseline model on the benchmark

Table 1 summarizes the results of the Image2Sphere [13] model on our benchmark, with per-entry results in Appendix D. While the wMANE is $\sim 24^\circ$, for the best cases the error is as small as 9° . The ΔPCC for the worse cases is > 0.1 , highlighting that, for some entries, the reconstructed volumes are

Table 1: Image2Sphere results on CESPED. MAnE and wMAnE measure angular errors in predicted poses. Δ PCC and Δ FSCR measure the reduction in quality of the predicted volumes compared to the ground-truth. The mean and std of the metrics are computed over the seven benchmark entries.

	MAnE (°)	wMAnE (°)	Δ PCC	Δ FSCR _{0.5} (Å)	Δ FSCR _{0.143} (Å)
mean (std)	29.2 (13.6)	24.2 (14.1)	0.056 (0.035)	3.6 (0.5)	1.5 (0.7)

far from the ground-truth solution. Still, as before, for a few instances, the results are much better, with Δ PCC < 0.03 . In terms of prediction vs ground-truth FSCR_{0.5}, most maps are in the 8-6 Å range, with Δ FSCR_{0.5} of 3.6 Å. However, the FSCR_{0.143} values between 4-5 Å, indicate better correlation at lower signal levels. This visually translates into a relatively well-resolved central part of the map that becomes blurrier away from the centre (see Appendix E). For the top-performing cases, a simple and fast local refinement of the predicted poses is sufficient to obtain high-resolution reconstructions comparable to ground-truth volumes, at a computational cost threefold less than global refinement (Appendix F). Since the Image2Sphere model inference takes only minutes, far less than the hours needed for traditional refinement, further improvements could reduce computing times by an order of magnitude if local refinement is no longer needed.

Appendix H presents an example where a model initially trained on the CESPED target 10280, is able to predict the poses of another dataset of the same protein in different conditions. While the results are of lower quality, they still exhibit a significant level of generalizability, offering promising potential to speed up image processing in this common scenario.

Given the inherent difficulties of Cryo-EM data, the fact that a generic pose-estimation model can produce meaningful results in some examples without major modifications suggests that equivariant architectures can be useful for the Cryo-EM data domain.

3.3 Challenges and future directions

Cryo-EM particle images are fundamentally different from the kinds of images encountered in other fields. One of the most critical challenges is their poor signal-to-noise ratio, which can be as low as 0.01 [1]. While some methods have tried to mitigate this issue by applying filtering techniques [14] or using CNNs with larger kernel sizes [2, 10, 26], these solutions are not entirely effective.

Symmetry presents another complex facet of Cryo-EM data. On the one hand, exploiting symmetry can drastically reduce the computational requirements for pose estimation. On the other, it can prevent simple models from learning. The unique combination of rotationally equivariant convolutions with the probabilistic estimation of poses makes the Image2Sphere model an ideal candidate to exploit this feature. However, the hybrid S^2 / SO(3) formalism means that separation of rotational degrees of freedom from translational in-plane shifts is not easily achieved within this framework. A significant area for future work lies in leveraging rotational equivariance and translational equivariance for the joint estimation of the rotational and translational components of the poses (e.g. SE(3) equivariance).

In this work we have considered only the case of homogeneous refinement, which assumes that all particles are projections from a single macromolecule in a unique conformation. However, this is not always the case and our benchmark could potentially be extended to deal with such examples. Models would then need to perform conformation classification alongside pose estimation.

4 Conclusions

Pose estimation is one of the most critical steps of the Cryo-EM processing pipeline, and while current algorithms are relatively robust and reliable, they are also computationally slow. Deep learning holds the promise of overcoming these challenges, but achieving this potential hinges on improvements in accuracy and reliability, for which systematic benchmarking is required. In this study, we introduce a benchmark specifically designed for Supervised Pose Inference of Cryo-EM particles, along with a suite of code utilities to assist machine learning practitioners unfamiliar with Cryo-EM. We also present a real-world image pose prediction model applied to our benchmark, demonstrating promising preliminary results on a subset of the data. This preliminary success suggests that addressing Cryo-EM-specific challenges, such as high noise levels and label inaccuracies, could

lead to even better performance. The improvements in models for this benchmark will not only pave the way for more effective Supervised Pose Prediction models, but are also likely to give rise to innovative approaches to closely related challenges like Unsupervised Pose Estimation and Heterogeneity Analysis. Ultimately, those advancements could serve as a catalyst for even further developments, leading to a new paradigm in Cryo-EM image processing.

5 Availability

CESPED dataset and code can be found at <https://github.com/rsanchezgarc/cesped>.

6 Acknowledgments

Ruben Sanchez-Garcia is funded by an Astex Pharmaceuticals Sustaining Innovation Post-Doctoral Award. Javier Vargas is supported by the Spanish Ministry of Science and Innovation, program Proyectos de Generación de Conocimiento 2022 with reference PID2022-137548OB-I00.

References

- [1] Tamir Bendory, Alberto Bartesaghi, and Amit Singer. Single-particle cryo-electron microscopy: Mathematical theory, computational challenges, and opportunities. *IEEE signal processing magazine*, 37(2):58, mar 2020. ISSN 15580792. doi: 10.1109/MSP.2019.2957822. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7213211/>.
- [2] Tristan Bepler, Andrew Morin, Micah Rapp, Julia Brasch, Lawrence Shapiro, Alex J. Noble, and Bonnie Berger. Positive-unlabeled convolutional neural networks for particle picking in cryo-electron micrographs. *Nature Methods*, 16(11):1153–1160, 2019. doi: 10.1038/s41592-019-0575-8. URL <https://www.nature.com/articles/s41592-019-0575-8>.
- [3] J. M. De la Rosa-Trevín, J. Otón, R. Marabini, A. Zaldívar, J. Vargas, J. M. Carazo, and C. O.S. Sorzano. Xmipp 3.0: An improved software suite for image processing in electron microscopy. *Journal of Structural Biology*, 184(2):321–328, nov 2013. ISSN 10478477. doi: 10.1016/j.jsb.2013.09.015. URL <http://www.ncbi.nlm.nih.gov/pubmed/24075951>.
- [4] Haowen Deng, Mai Bui, Nassir Navab, Leonidas Guibas, Slobodan Ilic, and Tolga Birdal. Deep Bingham Networks: Dealing with Uncertainty and Ambiguity in Pose Estimation. *International Journal of Computer Vision*, 130(7):1627–1654, jul 2022. ISSN 15731405. doi: 10.1007/S11263-022-01612-W/METRICS. URL <https://link.springer.com/article/10.1007/s11263-022-01612-w>.
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255. IEEE, jun 2009. ISBN 978-1-4244-3992-8. doi: 10.1109/CVPR.2009.5206848. URL <https://ieeexplore.ieee.org/document/5206848/>.
- [6] Timothy Grant, Alexis Rohou, and Nikolaus Grigorieff. CisTEM, user-friendly software for single-particle image processing. *eLife*, 7, mar 2018. ISSN 2050084X. doi: 10.7554/eLife.35383. URL <https://elifesciences.org/articles/35383>.
- [7] Harshit Gupta, Michael T McCann, Laurene Donati, and Michael Unser. CryoGAN: A New Reconstruction Paradigm for Single-Particle Cryo-EM Via Deep Adversarial Learning. *IEEE Transactions on Computational Imaging*, 7:759–774, mar 2021. ISSN 23339403. doi: 10.1109/TCI.2021.3096491. URL <https://www.biorxiv.org/content/10.1101/2020.03.20.001016v2>.
- [8] Richard Henderson, Andrej Sali, Matthew L. Baker, Bridget Carragher, Batsal Devkota, Kenneth H. Downing, Edward H. Egelman, Zukang Feng, Joachim Frank, Nikolaus Grigorieff, Wen Jiang, Steven J. Ludtke, Ohad Medalia, Pawel A. Penczek, Peter B. Rosenthal, Michael G. Rossmann, Michael F. Schmid, Gunnar F. Schröder, Alasdair C. Steven, David L. Stokes, John D. Westbrook, Willy Wriggers, Huanwang Yang, Jasmine Young, Helen M. Berman, Wah Chiu, Gerard J. Kleywegt, and Catherine L. Lawson. Outcome of the first electron microscopy

- validation task force meeting. In *Structure*, volume 20, pages 205–214. Elsevier, feb 2012. doi: 10.1016/j.str.2011.12.014. URL <http://www.cell.com/article/S0969212612000147/fulltext>.
- [9] Andrii Iudin, Paul K Korir, José Salavert-Torres, Gerard J Kleywegt, and Ardan Patwardhan. EMPIAR: A public archive for raw electron microscopy image data, may 2016. ISSN 15487105. URL <http://www.nature.com/articles/nmeth.3806>.
 - [10] A. Jiménez-Moreno, D. Štěrlák, J. Filipovič, J. M. Carazo, and C. O.S. Sorzano. DeepAlign, a 3D alignment method based on regionalized deep learning for Cryo-EM. *Journal of Structural Biology*, 213(2):107712, jun 2021. ISSN 10958657. doi: 10.1016/j.jsb.2021.107712. URL <https://www.sciencedirect.com/science/article/pii/S1047847721000174>.
 - [11] Alex Kendall and Roberto Cipolla. Geometric Loss Functions for Camera Pose Regression with Deep Learning. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017-Janua:6555–6564, jul 2017. ISSN 1063-6919. doi: 10.1109/CVPR.2017.694. URL <https://arxiv.org/abs/1704.00390>.
 - [12] Alex Kendall, Matthew Grimes, and Roberto Cipolla. PoseNet: A convolutional network for real-time 6-dof camera relocalization. In *Proceedings of the IEEE International Conference on Computer Vision*, volume 2015 Inter, pages 2938–2946, may 2015. ISBN 9781467383912. doi: 10.1109/ICCV.2015.336. URL <http://arxiv.org/abs/1505.07427>.
 - [13] David M. Klee, Ondrej Biza, Robert Platt, and Robin Walters. Image to Sphere: Learning Equivariant Features for Efficient Pose Prediction. *International Conference on Learning Representations*, feb 2023. URL <http://arxiv.org/abs/2302.13926>.
 - [14] Axel Levy, Frédéric Poitevin, Julien Martel, Youssef Nashed, Ariana Peck, Nina Miolane, Daniel Ratner, Mike Dunne, and Gordon Wetzstein. CryoAI: Amortized Inference of Poses for Ab Initio Reconstruction of 3D Molecular Volumes from Real Cryo-EM Images. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 13681 LNCS, pages 540–557, mar 2022. ISBN 9783031198021. doi: 10.1007/978-3-031-19803-8_32. URL <https://arxiv.org/abs/2203.08138v4>.
 - [15] Axel Levy, Gordon Wetzstein, Julien Martel, Frederic Poitevin, and Ellen D. Zhong. Amortized Inference for Heterogeneous Reconstruction in Cryo-EM. *Advances in neural information processing systems*, 35:13038, dec 2022. ISSN 10495258. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10392957/>.
 - [16] Ruyi Lian, Bingyao Huang, Liguang Wang, Qun Liu, Yuewei Lin, and Haibin Ling. End-to-end orientation estimation from 2D cryo-EM images. *Acta Crystallographica Section D: Structural Biology*, 78(2):174–186, jan 2022. ISSN 20597983. doi: 10.1107/S2059798321011761. URL <https://scripts.iucr.org/cgi-bin/paper?S2059798321011761>.
 - [17] Siddharth Mahendran, Haider Ali, and Rene Vidal. A mixed classification-regression framework for 3D pose estimation from 2D images. In *British Machine Vision Conference 2018, BMVC 2018*, 2019.
 - [18] Iaroslav Melekhov, Juha Ylioinas, Juho Kannala, and Esa Rahtu. Image-Based Localization Using Hourglass Networks. *Proceedings - 2017 IEEE International Conference on Computer Vision Workshops, ICCVW 2017*, 2018-Janua:870–877, jul 2017. doi: 10.1109/ICCVW.2017.107. URL <https://arxiv.org/abs/1703.07971>.
 - [19] David Mohlin, Gérald Bianchi Tobii Danderyd, and Josephine Sullivan. Probabilistic Orientation Estimation with Matrix Fisher Distributions. *Advances in Neural Information Processing Systems*, 33:4884–4893, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/33cc2b872>.
 - [20] Kieran Murphy, Carlos Esteves, Varun Jampani, Srikumar Ramalingam, and Ameesh Makadia. Implicit-PDF: Non-Parametric Representation of Probability Distributions on the Rotation Manifold. In *Proceedings of Machine Learning Research*, volume 139, pages 7882–7893. PMLR, jul 2021. ISBN 9781713845065. URL <https://proceedings.mlr.press/v139/murphy21a.html>.

- [21] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, volume 32, 2019. doi: 10.5555/3454287.3455008. URL <https://dl.acm.org/doi/10.5555/3454287.3455008>.
- [22] Pawel A. Penczek. Image restoration in cryo-electron microscopy. In *Methods in Enzymology*, volume 482, pages 35–72. NIH Public Access, 2010. doi: 10.1016/S0076-6879(10)82002-6. URL [/pmc/articles/PMC3166661/](http://pmc/articles/PMC3166661/).
- [23] Sergey Prokudin, Peter Gehler, and Sebastian Nowozin. Pose estimation with uncertainty quantification. In *Proceedings of the European conference on computer vision*, pages 534–551. Springer Verlag, may 2018. ISBN 9783030012397. doi: 10.1007/978-3-030-01240-3_33. URL <https://arxiv.org/abs/1805.03430v1>.
- [24] Ali Punjani, John L. Rubinstein, David J. Fleet, and Marcus A. Brubaker. cryoSPARC: algorithms for rapid unsupervised cryo-EM structure determination. *Nature Methods* 2017 14:3, 14(3):290–296, feb 2017. ISSN 1548-7105. doi: 10.1038/nmeth.4169. URL <https://www.nature.com/articles/nmeth.4169>.
- [25] Peter B. Rosenthal and Richard Henderson. Optimal determination of particle orientation, absolute hand, and contrast loss in single-particle electron cryomicroscopy. *Journal of Molecular Biology*, 333(4):721–745, oct 2003. ISSN 00222836. doi: 10.1016/j.jmb.2003.07.013.
- [26] Ruben Sanchez-Garcia, Joan Segura, David Maluenda, Jose Maria Carazo, and Carlos Oscar Sorzano Sorzano. Deep Consensus, a deep learning-based approach for particle pruning in cryo-electron microscopy Ruben. *IUCrJ*, 5(Pt 6):854–865, nov 2018. ISSN 20522525. doi: 10.1107/S2052252518014392. URL <http://scripts.iucr.org/cgi-bin/paper?S2052252518014392>.
- [27] Sjors H.W. Scheres. RELION: Implementation of a Bayesian approach to cryo-EM structure determination. *Journal of Structural Biology*, 180(3):519–530, dec 2012. ISSN 10478477. doi: 10.1016/j.jsb.2012.09.006. URL <http://www.ncbi.nlm.nih.gov/pubmed/23000701>.
- [28] Sjors H.W. Scheres and Shaoxia Chen. Prevention of overfitting in cryo-EM structure determination. *Nature Methods* 2012 9:9, 9(9):853–854, jul 2012. ISSN 1548-7105. doi: 10.1038/nmeth.2115. URL <https://www.nature.com/articles/nmeth.2115>.
- [29] Guang Tang, Liwei Peng, Philip R. Baldwin, Deepinder S. Mann, Wen Jiang, Ian Rees, and Steven J. Ludtke. EMAN2: An extensible image processing suite for electron microscopy. *Journal of Structural Biology*, 157(1):38–46, jan 2007. ISSN 10478477. doi: 10.1016/j.jsb.2006.05.009. URL <http://www.ncbi.nlm.nih.gov/pubmed/16859925>.
- [30] Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond PASCAL: A benchmark for 3D object detection in the wild. *2014 IEEE Winter Conference on Applications of Computer Vision, WACV 2014*, pages 75–82, 2014. doi: 10.1109/WACV.2014.6836101. URL https://cvgl.stanford.edu/papers/xiang_wacv14.pdf.
- [31] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. PoseCNN: A Convolutional Neural Network for 6D Object Pose Estimation in Cluttered Scenes. In *Robotics: Science and Systems*, 2018. ISBN 9780992374747. doi: 10.15607/RSS.2018.XIV.019. URL <https://arxiv.org/abs/1711.00199>.
- [32] Zenodo. Zenodo: Research. Shared., 2013. URL <https://zenodo.org/>.

APPENDIXES

Ruben Sanchez-Garcia^{1,2}, Michael Saur², Javier Vargas³, Carl Poelking², Charlotte M Deane¹

¹Department of Statistics, University of Oxford, Oxford OX1 3LB, UK
{ruben.sanchez-garcia, deane}@stats.ox.ac.uk

²Astex Pharmaceuticals, Cambridge CB4 0QA, UK
{michael.saur, carl.poelking}@astx.com

³Departamento de Optica, Universidad Complutense de Madrid, Madrid 28040, Spain
jvargas@fis.ucm.es

1 Appendix A. Benchmark composition.

Supplementary Table 1. CESPED benchmark entries.

EMPIAR ID	Composition	Symmetry	Image pixels	$\text{FSCR}_{0.143}$ (Å)	Masked $\text{FSCR}_{0.143}$ (Å)	# particles
10166	Human 26S proteasome bound to the chemotherapeutic Oprozomib	C1	284	5.0	3.9	238631
10786	Substance P-Neurokinin Receptor G protein complexes (SP-NK1R-miniGs399)	C1	184	3.3	3.0*	288659
10280	Calcium-bound TMEM16F in nanodisc with supplement of PIP2	C2	182	3.6	3.0*	459504
11120	M22 bound TSHR Gs 7TM G protein	C1	232	3.4	3.0*	244973
10647	PKM2 in complex with L-threonine	D2	222	3.7	3.3	234956
10409	Replicating SARS-CoV-2 polymerase (Map 1)	C1	240	3.3	3.0*	406001
10374	Human ABCG2 transporter with inhibitor MZ29 and 5D3-Fab	C2	216	3.7	3.0*	323681

- * Nyquist Frequency at 1.5 Å/pixel; Resolution is estimated at the usual threshold 0.143.
- Reported $\text{FSCR}_{0.143}$ values were obtained directly from the relion_refine logs while Masked $\text{FSCR}_{0.143}$ values were collected from the relion_postprocess logs.

2 Appendix B. Image2Sphere and training hyperparameters

Our Image2Sphere model follows the implementation of Klee et al. [3] with the following configuration:

- Feature extractor: ResNet152 [2] with default parameters as implemented in torchvision using imageNet weights. The input images are resized to 256 pixels before being fed, giving a feature map of shape 2048x8x8. Since the input images only contain one channel, but the ResNet expects 3 channels, two additional channels were added by applying a Gaussian filter with sigma 1 and 2 to the input image.
- Image projector to S2: Default orthographic projector with HEALPix [1] grid order 3 ($\sim 7.5^\circ$), where only 50% of the grid points are sampled. The feature map is projected from 2048 channels to 512 using a 1x1 Conv2d and then converted to spherical harmonics with $l_{max} = 8$.
- S2 convolution: 512 filters with global support on a HEALPix grid of order 3.
- SO(3) convolution: 16 filters with local support ($\max_beta = \pi/8$, $\max_gamma = 2\pi$, $n_alpha = 8$, $n_beta = 3$).
- Probability distribution discretization: HEALPix grid of order 4 ($\sim 3.7^\circ$).

Training was conducted using RAdam [4] as optimizer with an initial learning rate of $1e-3$. A weight decay of $1e-5$ was employed. The learning rate was halved each time the validation loss stagnated during 10 epochs. The training was stopped when the number of epochs reached 400 or the validation loss did not improve for 12 epochs.

Data augmentation was conducted with the following composed transformations:

- Random shift from -5% to 5% with probability 0.5.
- Random rotation from -20° to 20° with probability 0.5.
- Random 90° rotation with probability 1.
- Uniform noise addition with a random scale from 0 to 2 with probability 0.2.
- Gaussian noise addition with a random standard deviation from 0 to 0.5 with probability 0.2.
- Random zoom-in of size 0% to 5% with probability 0.2.
- Random erasing of patches of size 0% to 2% with probability 0.1.

Notice that rotation transformations require adjustments in the ground-truth labels.

3 Appendix C. ParticlesDataset usage example

Dataset instantiation only requires providing the name of the target (a string like "10280") and the half-set number (0 or 1) (Listing 1). ParticlesDatasets can be directly used as datasets in PyTorch DataLoader(s).

Listing 1: Example of how to load and use a CESPED benchmark entry in a training loop

```
import torch
from cespded.particlesDataset import ParticlesDataset

listOfEntries = ParticlesDataset.getCESPEDEntries()
targetName, halfset = listOfEntries[0] #We will work with the first example
dataset = ParticlesDataset(targetName, halfset)
dl = DataLoader(dataset, batch_size=32, num_workers=4)
for batch in dl:
    iid, img, (rotMat, xyShiftAngs, confidence), metadata = batch

    #iid is the id of the particle (a string)
    #img is a batch of Bx1xNxN images
    #rotMat is a batch of rotation matrices Bx3x3
    #xyShiftAngs is a batch of image shifts in Angstroms Bx2
    #confidence is a batch of numbers between 0 and 1, Bx1
    #metata is a dict of names:values with particle information

    predRot = model(img)
    loss = loss_function(rotMat, predRot)
    loss.backward()
    optimizer.step()
```

```
optimizer.zero_grad()
```

ParticlesDataset objects can also be used to update the metadata with newly predicted poses and to save the results in Relion star format, commonly used in Cryo-EM software (Listing 2).

Listing 2: Example of how to save predictions for usage in Cryo-EM packages and evaluation

```
for iid, pred_rotmats, maxprob in predictions:
    #iid is the list of ids of the particles (string)
    #pred_rotmats is a batch of predicted rotation matrices Bx3x3
    #maxprob is a batch of numbers, between 0 and 1, Bx1
    #that indicates the confidence in the prediction (e.g., softmax values)
    n_preds = pred_rotmats.shape[0]
    dataset.updateMd(ids=iid, angles=pred_rotmats,
                    shifts=torch.zeros(n_preds, 2),
                    confidence=maxprob,
                    angles_format="rotmat")
dataset.saveMd("predictions.star") #Save dataset as an starfile
```

Once the predictions are computed for the two halves of the benchmark entry, evaluation can be automatically computed by providing the starfiles of both predictions via a command line tool (Listing 3) or a function. While you can use your local installation of Relion, we also provide a Singularity definition file so that you do not need to manually install it. See instructions at <https://github.com/rsanchezgarc/cesped>,

Listing 3: Evaluation script execution

```
python -m cespded.evaluateEntry --predictionType SO3 --targetName 10280
--half0PredsFname particles_preds_0.star
--half1PredsFname particles_preds_1.star
--n_cpus 12 --outdir evaluation/
```

Listing 4: Evaluation function example

```
from cespded.evaluateEntry import evaluate
evaluation_metrics = evaluate(targetName="10280",
                             half0PredsFname="particles_preds_0.star",
                             half1PredsFname="particles_preds_1.star",
                             predictionType="SO3", #Literal["S2", "SO3", "SO3xR2"],
                             usePredConfidence=True,
                             n_cpus=4,
                             outdir="output/directory")
```

4 Appendix D. Image2Sphere per-entry results

This section contains per-entry statistics for the Image2Sphere model predictions using the evaluation protocol proposed in the main text. These metrics compare the ground-truth map (GT) against the map reconstructed from the predicted poses (V), namely, $PCC(GT, V)$ and $FSCR_t(GT, V)$, where t denotes the threshold 0.5 or 0.143 and

$$PCC(X, Y) = \frac{\sum_{i=1}^n (X_i - \bar{X})(Y_i - \bar{Y})}{\sqrt{\sum_{i=1}^n (X_i - \bar{X})^2} \sqrt{\sum_{i=1}^n (Y_i - \bar{Y})^2}}$$

where n is the number of voxels and X and Y are two volumes, and $FSCR_t$ is the first frequency k such that $FSCR_k(X, Y) < 0.143$ with

$$FSCR_k(X, Y) = \frac{\sum_{\mathbf{r} \in \text{shell}(k)} \hat{X}(\mathbf{r}) \cdot \hat{Y}^*(\mathbf{r})}{\sqrt{\left(\sum_{\mathbf{r} \in \text{shell}(k)} |\hat{X}(\mathbf{r})|^2 \right) \cdot \left(\sum_{\mathbf{r} \in \text{shell}(k)} |\hat{Y}(\mathbf{r})|^2 \right)}}$$

where $\hat{X}(\mathbf{r})$ and $\hat{Y}(\mathbf{r})$ represent the Fourier transforms of the two volumes at frequency \mathbf{r} , $\text{shell}(k)$ is the shell of frequency k and $\hat{Y}^*(\mathbf{r})$ is the complex conjugate of $\hat{Y}(\mathbf{r})$.

In addition, we also report half-to-half map metrics $PCC(GT_0, GT_1)$, $PCC(V_0, V_1)$, $FSCR_t(GT_0, GT_1)$ and $FSCR_t(V_0, V_1)$, where 0 and 1 denote the dataset half.

Supplementary Table 2. Per-entry CESPED benchmark results using an Image2Sphere model.


EMPIAR ID	MAnE (°)	wMAnE (°)	PCC(V_0, V_1)	PCC(GT, V)	FSCR _{0.143} (V_0, V_1) (Å)	FSCR _{0.5} (V_0, V_1) (Å)	FSCR _{0.143} (GT, V) (Å)	FSCR _{0.5} (GT, V) (Å)	FSCR _{0.143} (GT_0, GT_1) (Å)	FSCR _{0.5} (GT_0, GT_1) (Å)	PCC(GT_0, GT_1)
10166	15.7	9.1	0.986	0.974	5.1	6.8	6.2	8.1	4.4	4.8	0.992
10786	32.6	29.5	0.957	0.925	3.8	4.3	3.4	7.6	3.1	3.5	0.974
10280	17.8	14.9	0.954	0.949	4.2	5.0	4.6	7.3	3.6	3.9	0.981
11120	44.7	41.1	0.989	0.863	4.1	4.6	6.0	8.3	3.2	3.7	0.965
10647	13.3	10.6	0.941	0.918	4.0	4.5	4.6	7.3	3.5	3.7	0.977
10409	45.3	39.2	0.960	0.884	3.5	4.0	4.0	8.3	3.0	3.3	0.988
10374	35.0	24.8	0.986	0.968	3.7	4.2	4.3	7.0	3.0	3.6	0.993

- MAnE: Mean Angular Error; wMAnE: weighted Mean Angular Error; PCC(V_0, V_1): Reconstructed half to half Pearson's Correlation Coefficient; PCC(GT, V): Reconstructed to ground-truth Pearson's Correlation Coefficient; FSCR_{0.143}(V_0, V_1): Reconstructed half to half FSC resolution at threshold 0.143, and FSCR_{0.5}(V_0, V_1) at threshold 0.5; FSCR_{0.143}(GT, V): Reconstructed to ground-truth resolution at threshold 0.143, and FSCR_{0.5}(GT, V) at threshold 0.5; FSCR_{0.143}(GT_0, GT_1): Ground-truth half to half FSC resolution at threshold 0.143, and FSCR_{0.5}(GT_0, GT_1) at threshold 0.5. PCC(GT_0, GT_1): Ground-truth half to half Pearson's Correlation Coefficient.
- All reported resolutions were obtained using manually computed masks that are available at <https://zenodo.org/record/8392782>.


5 Appendix E. Reconstructed volumes

This appendix shows the volumes reconstructed for the best performing examples of the Image2Sphere model on our benchmark. As it is shown in all three cases, the quality of the central region of the protein is quite close to the one of the ground-truth. However, the density for the regions that are at the edges of the macromolecule are much worse. This is in line to what could be expected if


there were some degree of inaccuracy in the angular estimation, as the magnitude of the errors in the volume is proportional to both the angular error and the radius of the macromolecule.



Supplementary Figure 1. Ground-truth reconstruction for EMPIAR-10280 (grey) and reconstruction using the angles predicted with the Image2Sphere model (yellow).




Supplementary Figure 2. Ground-truth reconstruction for EMPIAR-10374 (grey) and reconstruction using the angles predicted with the Image2Sphere model (yellow).



Supplementary Figure 3. Ground-truth reconstruction for EMPIAR-10786 (grey) and reconstruction using the angles predicted with the Image2Sphere model (yellow).

6 Appendix F. Locally refined solution


In this section, we illustrate the usefulness of our approach by showing the effect of classical local refinement on the Image2Sphere results for the benchmark entry 10374. In this case, the Image2Sphere model predicted poses with a wMANE of 24.8° that lead to a reconstructed map with $\text{FSCR}_{0.143}(V_0, V_1)$ of 3.7 \AA . When the predicted poses are used as priors for a local refinement in Relion with `-sigma_angle 2`, the refined map achieved a $\text{PCC}(GT, V)$ of 0.997 compared to the original 0.968, showing that the refined map is much more similar to the ground truth map. Indeed, as it can be appreciated in Supplementary Figure 4, after the local refinement, not only the quality of the core of the protein is comparable to the quality of the ground truth, but also the quality of the distant parts of the maps is much better, almost as good as in the ground truth. Equally important, since we are limiting the angular search to the neighbourhood around the predicted poses ($\pm 6^\circ$), the number of image comparisons carried out by Relion is much smaller, resulting in a three-fold speed-up in computational time, even when including the time required for pose inference using the Image2Sphere model.



Supplementary Figure 4. Ground-truth reconstruction for EMPIAR-10374 (grey) and reconstruction using the angles predicted with the Image2Sphere model and locally refined using Relion with priors (yellow).

7 Appendix H. Example of generalizability across dataset

In this section, we present results from an Image2Sphere model trained on the EMPIAR-10280 protein (CESPED entry 10280), used to predict poses of the same protein under different conditions in another dataset (EMPIAR-10278). Supplementary Figure 5 displays the reconstructed volumes for the EMPIAR-10278 dataset using inferred (yellow) and actual poses (grey), alongside the volume reconstructed for EMPIAR-10280 using inferred poses (blue). As expected, the EMPIAR-10278 map reconstructed with original poses shows superior quality compared to the others. Similarly, the EMPIAR-10280 map generated from the model trained on EMPIAR-10280 half-datasets, exhibits better quality than the EMPIAR-10278 map inferred using the same model, reflecting the differences between the two datasets despite containing the same protein. Despite these quality differences, the model's capacity for generalization across datasets is evident through visual inspection, as the overall shape of the protein and several key secondary structure elements are clearly recognizable. This suggests that further improvements in the model could lead to the desired goal of training the model once and then inferring the poses of similar datasets at much faster speeds.



Supplementary Figure 5. Ground-truth reconstruction for EMPIAR-10278 (grey) and reconstruction using the angles predicted with an Image2Sphere model trained on EMPIAR 10280 (yellow), and reconstruction for the EMPIAR 10280 using the angles predicted with the same Image2Sphere model trained on EMPIAR 10280 (blue).

References

- [1] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelman. HEALPix – a Framework for High Resolution Discretization, and Fast Analysis of Data Distributed on the Sphere. *The Astrophysical Journal*, 622(2):759–771, sep 2004. doi: 10.1086/427976. URL <http://arxiv.org/abs/astro-ph/0409513><http://dx.doi.org/10.1086/427976>.
- [2] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2016-Decem:770–778, dec 2016. ISSN 10636919. doi: 10.1109/CVPR.2016.90.
- [3] David M. Klee, Ondrej Biza, Robert Platt, and Robin Walters. Image to Sphere: Learning Equivariant Features for Efficient Pose Prediction. *International Conference on Learning Representations*, feb 2023. URL <http://arxiv.org/abs/2302.13926>.
- [4] Yi Liu, Hao Yuan, Lei Cai, and Shuiwang Ji. Deep Learning of High-Order Interactions for Protein Interface Prediction. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 20:679–687, jul 2020. doi: 10.1145/3394486.3403110. URL <http://arxiv.org/abs/2007.09334><http://dx.doi.org/10.1145/3394486.3403110>.