



Deep learning of elastic 3D shapes for cryo electron microscopy analysis of continuous conformational changes of biomolecules

Ilyes Hamitouche, Slavica Jonic

► To cite this version:

Ilyes Hamitouche, Slavica Jonic. Deep learning of elastic 3D shapes for cryo electron microscopy analysis of continuous conformational changes of biomolecules. 29th European Signal Processing Conference, EUSIPCO 2021, Aug 2021, Dublin, Ireland. pp.1251-1255, <10.23919/EUSIPCO54536.2021.9616013>. <hal-03266630>

HAL Id: hal-03266630

<https://hal.science/hal-03266630v1>

Submitted on 22 Jun 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

This is the author's version of an article accepted for publication in Proc. EUSIPCO 2021. The final published version will be available on IEEE Xplore and Eurasip website as Open Access.

Deep learning of elastic 3D shapes for cryo electron microscopy analysis of continuous conformational changes of biomolecules

Ilyes Hamitouche and Slavica Jonić*

IMPMC - UMR 7590 CNRS, Sorbonne Université, MNHN, Paris, France

ilyes.hamitouche@upmc.fr, slavica.jonic@upmc.fr (corresponding)*

Abstract—Cryo electron microscopy (cryo-EM) allows high-resolution 3D reconstruction of biomolecular structures from highly noisy 2D parallel-beam projection images containing tens of thousands of copies of the same macromolecular complex but at different random orientations and positions. However, biomolecular complexes are not rigid but flexible entities that change their conformations gradually (continuous transition with many intermediate states) to accomplish biological functions (e.g., DNA replication, protein synthesis, etc.). The determination of the full distribution of conformations (conformational space or landscape) from cryo-EM images is challenging but could provide insights into working mechanisms of the complexes. In this paper, we present a method for conformational space determination, which uses deep learning in combination with cryo-EM image analysis and normal mode analysis (molecular mechanics simulation), where the amplitudes of normal modes are used as parameters of the elastic 3D shapes of complexes (the parameters determining the conformation). We show the performance of this new method using synthetic cryo-EM data.

Index Terms—Deep Learning, Cryo-EM, Elastic 3D-to-2D alignment, Normal Modes, Molecular Dynamics

I. INTRODUCTION

Biomolecular complexes are dynamical objects. They switch among different conformations in response to the environment to accomplish different functional roles. Single-particle cryo electron microscopy (cryo-EM) allows biomolecular structure determination at near-atomic resolution from vitrified homogeneous samples containing many copies of the complex (known as single particles) at random unknown 3D orientations and positions. For conformationally heterogeneous samples, such high resolution of reconstruction can be achieved in many cases but only after many rounds of computationally demanding 2D and 3D classifications to disentangle the 3D shape heterogeneity from 3D orientational and translational heterogeneity [1]. Typically, these approaches result in throwing away all particle

images that do not coincide with the class of images yielding the highest-resolution 3D reconstruction. However, elucidation of conformational variability of biomolecular complexes is the key to understanding their biological functions (e.g., DNA replication, protein synthesis, virus maturation, etc.) and to novel drug discovery [2]. The particle images thrown away in the process of high-resolution 3D reconstruction could carry important biomedical information. Besides, splitting the set of images into a number of classes defined prior to conformational variability analysis is suited to discrete conformational changes (e.g., two-state heterogeneity of binding and unbinding of a complex with another molecule) but, generally, it is suboptimal. Indeed, biomolecular complexes generally adopt gradual transitions with a large unknown number of intermediate conformational states (continuous conformational changes), which yields a particularly challenging type of heterogeneity [3].

The majority of the existing methods are based on classification into a small number of discrete classes, which is often defined using a prior knowledge about the number of expected conformations [4], [5]. The other group of methods deal with continuous conformational heterogeneity by determining the full conformational distribution (also called conformational space, conformational landscape, or conformational manifold) [6], [7], [8], based on which groups of images with similar conformations are made and 3D reconstructions from these groups computed. The development of methods for continuous conformational heterogeneity analysis is currently an active field of research (for more information, the reader is referred to the recent review in [9]). This article presents an original approach to this problem, which is based on deep learning.

Currently, the major applications of deep learning in cryo-EM are automated particle picking [10] and pruning [11]. Recently, a deep learning approach has been reported tackling

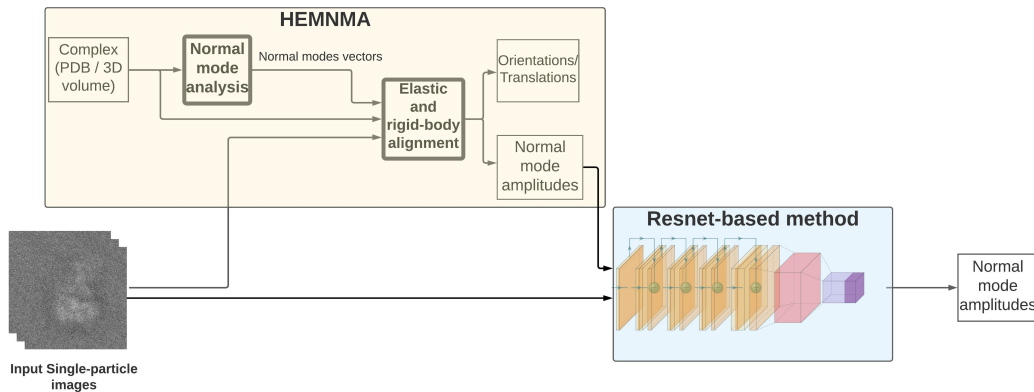


Fig. 1: Pipeline of the approach that combines HEMNMA with deep learning based on Resnet. Only two of HEMNMA modules (the first two HEMNMA steps) are used here, namely “Normal Mode Analysis” and “Elastic and rigid-body alignment” (the blocks with highlighted edges). The inputs to HEMNMA are single particle images and a 3D reference (an atomic structure or a density volume as the volumes that can be obtained by 3D reconstruction from 2D images). The outputs of “Elastic and rigid-body alignment” step are the estimated normal mode amplitudes and the estimated image orientations and translations. For training, only the HEMNMA-estimated normal-mode amplitudes and the input single particle images are used. The neural network is trained to infer normal-mode amplitudes for new input single particle images that were not seen during the training.

the problem of continuous conformational variability using a variational autoencoder but assuming that the orientations and translations of images are already known i.e., determined through alignment with a consensus density map by standard cryo-EM approaches [12]. This is a typical assumption in the majority of other continuous conformational variability analysis methods [7] and its consequence is that these methods are suited to the cases where the conformational heterogeneity is smaller than the orientational and translational heterogeneity so that the orientations and translations can be accurately determined by assuming conformational homogeneity. This limits their applications and requires a prior reduction of the conformational heterogeneity in the given set of images using standard cryo-EM approaches based on multiple runs of 2D and 3D classifications and subclassifications.

Contrary to other approaches, the method known as HEMNMA determines simultaneously all unknown parameters (orientation, position, and conformation), but at the price of high computational cost [6], [13]. HEMNMA software is freely available as part of ContinuousFlex plugin [13] for Scipion V3 open-source software package [14], extensively used in cryo-EM. In this article, we present a hybrid approach that combines HEMNMA with deep learning to speed up the determination of the conformational manifold. In a preliminary work (not presented here), we have compared convolutional autoencoders [15] and residual neural networks (Resnet) [16]. Here, we present our approach using Resnet 34 model, which showed the highest accuracy for this application over the other tested models. The results of our hybrid approach are shown using synthetic cryo-EM data.

II. METHODS

In this section, we present our hybrid approach combining HEMNMA with deep learning (Fig. 1).

A. HEMNMA

HEMNMA is a method that combines cryo-EM image analysis with molecular mechanics simulation by normal mode analysis (NMA). NMA based on the so-called elastic network model [17] of the structure is a simple and fast method to calculate vibrational modes and has been successfully used to predict biologically relevant motions [18]. NMA uses a quadratic approximation of the potential energy function of the molecular system around the given reference structure. The reference structure can be represented with atoms (such as the structures publicly available in the protein data bank PDB at www.rcsb.org) or with pseudoatoms (e.g., 3D Gaussian functions with which the preliminary 3D reconstruction from cryo-EM images can be represented [19]). In the elastic network model, atoms or pseudoatoms are considered to be connected by elastic springs [17] and normal modes are obtained by diagonalizing the Hessian matrix (the matrix of second derivatives of the potential energy function). Normal modes (vectors) and their (squared) frequencies are eigenvectors and eigenvalues of the Hessian matrix, respectively. Low-frequency modes describe collective motions while high-frequency modes describe local motions. Several studies have shown that low-frequency modes correspond to functionally relevant biomolecular motions and that conformational transitions globally follow one or a few low-frequency normal modes [20]. It has been demonstrated that low-frequency normal modes are directly linked to the shape of the molecular system [21].

Atoms (or pseudoatoms) are displaced, to form a new conformation, using a linear combination of normal modes. The coefficients of the linear combination are contributions of different normal modes to the displacement (normal-mode amplitudes). The elements of a normal-mode vector provide information on the direction of the displacement of each atom with this normal mode (the displacement is in angstroms, Å,

which are standard atomic-coordinate units). The length of the vector is determined by the number of atomic coordinates (3 times the number of atoms). The amplitudes of normal modes have no units. NMA allows calculating normal-mode vectors but not the amplitudes of normal modes. The normal-mode amplitudes determine the 3D shape of the molecular system. HEMNMA allows calculating the normal-mode amplitudes for each single particle image (i.e., the 3D shape of the biomolecular complex in each image), which is done simultaneously with the calculation of orientations and translations. The normal-mode amplitudes (elastic parameters) and the image orientation and translation (rigid-body parameters) are iteratively adjusted by optimization of a measure of similarity (alignment) between the image and the projection of the reference structure being elastically modified using normal modes, until the best alignment is achieved. The first six lowest-frequency modes (modes 1-6) are related to combinations of rigid-body motions and are not used because we deal with rotations and translations using rigid-body alignment that HEMNMA performs simultaneously with the elastic alignment.

B. Approach combining HEMNMA with deep learning

The proposed pipeline consists of HEMNMA image analysis followed by deep learning using a convolutional neural network (Fig. 1). The network has two blocks. The first block consists of a Resnet-based model (Resnet 34) that serves to extract features. The second block consists of a 4-layer multilayer perceptron that serves to predict normal mode amplitudes.

HEMNMA is first used to calculate normal modes of the given reference structure (an atomic structure or a density volume), by normal mode analysis. Then, the normal modes, the reference structure, and a set of single-particle images are used within the elastic and rigid body alignment method of HEMNMA to estimate (simultaneously) normal-mode amplitudes, orientations and translations of the images.

The neural network is trained using the HEMNMA-estimated normal-mode amplitudes and the single-particle images. At the inference step, the neural network infers the normal-mode amplitudes from a new set of single particle images.

We use images of size 256×256 pixels for training, validation, and inference. We use their resized versions to 128×128 pixels for HEMNMA estimation, which reduces noise and makes HEMNMA faster.

III. EXPERIMENTS

In this section, we describe our procedure for data synthesis and the experiments performed with these data.

A. Data

We synthesized 2D single particle images of an atomic structure of Adenylate Kinase chain A (AK) from the PDB database (code PDB:4AKE, Fig. 2).

We calculated normal modes of the AK structure and selected three lowest-frequency non-rigid-body modes (modes 7-9). To synthesize one image, AK atoms were first displaced along modes 7-9 with random displacement amplitudes and converted




Fig. 2: Atomic structure of AK from PDB database used to synthesize images and examples of synthetic images. (a) Three views of AK. (b) Left: ideal synthetic image of one view of AK. Right: synthetic image shown on the left side but with applied noise (SNR=0.1) and CTF (defocus of $-0.5 \mu\text{m}$).

into a density map, which was then projected onto an image plane at random orientation in 3D space defined by two out-of-plane rotations (the views covering the entire 3D sphere). Random in-plane rotations (between 0° and 360°) and in-plane shifts (between -5 and $+5$ pixels in x and y directions) were also generated for some tests (different data sets were generated as explained below). Then, the obtained ideal projection image was modified by applying noise (SNR=0.1) and the contrast transfer function (CTF) of the simulated cryo-electron microscope with the defocus of $-0.5 \mu\text{m}$ (the noise and defocus levels corresponding to realistic conditions). This process was repeated to synthesize the desired number of images. In this process, the random amplitudes q_7 - q_9 for modes 7-9, respectively, were generated as follows:

$$\begin{aligned} q_7(r) &= -200 \cdot r, & q_8(r) &= 200 \cdot \sin(\pi \cdot r), \\ q_9(r) &= 200 \cdot \cos(\pi \cdot r), \end{aligned} \quad (1)$$

where r is a random variable, uniformly distributed between 0 and 1, meaning that each image was synthesized using a different random value of r .

To investigate the limitations of our method, the following types of images were generated without and with imperfections (noise and CTF) using the above procedure (random normal-mode amplitudes q_7 - q_9 and random out-of-plane rotations were used in each set): 1) no in-plane rotation, no in-plane shift; 2) random in-plane rotation, no shift; 3) random in-plane shift, no in-plane rotation; and 4) random in-plane rotation and random in-plane shift. Each of these eight sets of synthetic images consisted of 20,000 images (size of 256×256 pixels).

B. Results

In this section, we describe the following two types of experiments: 1) training of our network using synthetic images and normal-mode amplitudes with which the images were generated (ground-truth normal-mode amplitudes); and 2) training of our network using synthetic images and normal-mode amplitudes estimated with HEMNMA from the same images (HEMNMA-estimated normal-mode amplitudes).

TABLE I: Mean absolute error and standard deviation of the inferred normal-mode amplitudes q_7 - q_9 , after training with ground-truth amplitudes. The errors are calculated with respect to the ground-truth amplitudes. The mean (*mean*) and the standard deviation (*std*) are indicated by $mean \pm std$. The global error is the mean absolute error over q_7 - q_9 . In all eight cases, out-of-plane rotations are random. The last row is a typical case in cryo-EM where all rotations and shifts are random.

Cases	Noise	Defocus	Global error	q_7 error	q_8 error	q_9 error
No in-plane rotation, no shift	No	No	2.32	1.49 ± 2.06	3.10 ± 4.49	2.35 ± 2.90
No in-plane rotation, no shift	SNR 0.1	-0.5 μm	5.79	3.62 ± 4.91	7.27 ± 10.69	6.49 ± 9.09
No in-plane rotation, random shift	No	No	4.83	3.02 ± 4.55	6.14 ± 10.11	5.32 ± 8.61
No in-plane rotation, random shift	SNR 0.1	-0.5 μm	7.95	4.90 ± 6.72	9.85 ± 14.72	9.10 ± 13.32
Random in-plane rotation, no shift	No	No	16.86	10.32 ± 15.38	19.53 ± 29.60	20.74 ± 33.43
Random in-plane rotation, no shift	SNR 0.1	-0.5 μm	19.62	12.00 ± 17.64	22.32 ± 31.95	24.53 ± 39.03
Random in-plane rotation and shift	No	No	23.51	14.67 ± 21.14	24.44 ± 34.00	31.42 ± 49.10
Random in-plane rotation and shift	SNR 0.1	-0.5 μm	27.60	17.14 ± 23.65	29.12 ± 39.42	36.55 ± 54.46

TABLE II: Mean absolute error and standard deviation of the inferred normal-mode amplitudes q_7 - q_9 , for images affected by noise and CTF, with random in-plane and out-of-plane rotations and random in-plane shifts, after training with HEMNMA-estimated amplitudes. The errors are calculated with respect to HEMNMA-estimated amplitudes and ground-truth amplitudes. The errors of HEMNMA-estimated amplitudes with respect to ground-truth amplitudes are also provided. The mean (*mean*) and the standard deviation (*std*) are indicated by $mean \pm std$. The global error is the mean absolute error over q_7 - q_9 . See also Figure 3.

ERRORS	GLOBAL	q_7	q_8	q_9
INFERRED vs. HEMNMA-ESTIMATED	19.03	12.35 ± 16.47	19.58 ± 26.15	25.16 ± 34.60
INFERRED vs. GROUND-TRUTH	20.22	12.61 ± 16.84	20.94 ± 27.37	27.13 ± 36.82
HEMNMA-ESTIMATED vs. GROUND-TRUTH	6.58	5.75 ± 8.42	6.23 ± 7.20	7.76 ± 7.23

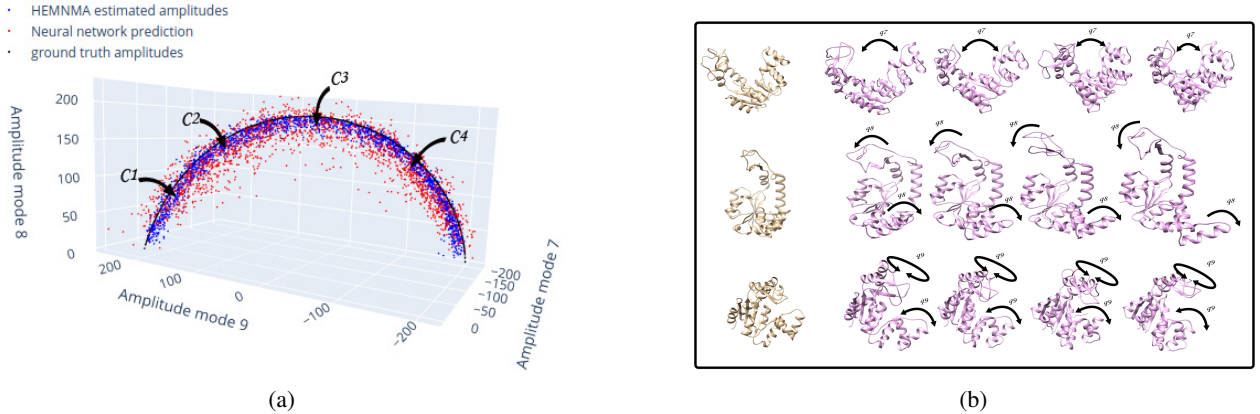


Fig. 3: Results of inferring normal-mode amplitudes for images affected by noise and CTF, with random in-plane and out-of-plane rotations and random in-plane shifts, after training with HEMNMA-estimated normal-mode amplitudes. (a) Overlap between inferred (red), ground-truth (black), and HEMNMA-estimated (blue) amplitudes (each point corresponds to an image and the molecular conformation inside it), with 4 selected points corresponding to 4 conformations (C1-C4) on a hypothetical trajectory through the points. (b) Three views of a sequence of 4 conformations (magenta) depicted in (a) by C1-C4 (one row per view to show, with arrows, a different motion for each of the three normal modes). See also Table II.

From each of the eight sets of synthesized images (Section III-A), one half of images (10,000) were used for the network training with ground-truth normal-mode amplitudes. They were split into 6,000 images for training, 2,000 images for validation, and 2,000 images for inference. The amplitude inference errors with respect to the ground-truth amplitudes are shown in Table I. One can note that the mean absolute error over q_7 - q_9 (“global error”) increases as the randomness of parameters increases and that it is slightly higher for images with noise and CTF than for those without. It should be noted that the training was done

without data augmentation (e.g., generating multiple copies of the same image by random in-plane rotations) and that such strategies may be investigated in the future to improve training and reduce the inference errors.

The case with noise, CTF, and all parameters random (the last row in Table I) is typical in practice and was used for the network training with HEMNMA-estimated amplitudes. From the eight synthesized sets of images, we took the set of 20,000 images affected by noise and CTF, with random in-plane and out-of-plane rotations and random shifts, and used HEMNMA

to estimate the normal-mode amplitudes for this set of images. Then, we analyzed the HEMNMA-estimated amplitudes and removed some outliers based on the Mahalanobis distance measure. From the remaining images (18,055 images), we used 14,000 images for training, 2,000 images for validation, and 2,000 images for inference. The amplitude inference errors with respect to the HEMNMA-estimated and ground-truth amplitudes in this case (training with HEMNMA-estimated amplitudes) are shown in Table II, together with the errors of the HEMNMA estimation with respect to the ground-truth.

The smaller errors in the second row of Table II than in the last row of Table I can be explained by a larger number of images used for training in Table II than in Table I (14,000 and 6,000 images, respectively). In Table II, the mean absolute errors of the inferred amplitudes with respect to the ground-truth amplitudes correspond to an average atomic-coordinate error of around 0.9 Å, as determined by the root-mean-square deviation (RMSD) between the reference atomic coordinates and these coordinates displaced (with normal modes) using the mean absolute errors as the normal-mode amplitudes. In the RMSD terms, the standard deviations of the same errors (more precisely, $mean \pm std$ in the second row of Table II) correspond to the minimum and maximum atomic-coordinate errors of around 0.3 Å and 2 Å, respectively. A 3D plot of the inferred, ground-truth, and HEMNMA-estimated amplitudes is shown in Figure 3a. Four points were selected along a hypothetical parabolic trajectory through the point cloud in the resulting conformational space and the normal-mode amplitudes corresponding to each of the four points were applied onto the reference atomic structure to displace it with normal modes. Figure 3b shows three views of the sequence of the four obtained conformations, describing motions that coincide well with the synthesized motions.

IV. CONCLUSION

We presented a hybrid method to accelerate determination of conformational landscapes of biomolecular complexes from cryo-EM images. It combines a convolutional neural network with HEMNMA cryo-EM image analysis method, which on its turn combines normal mode analysis (a simple and fast molecular mechanics simulation method) with cryo-EM image analysis. The amplitudes of normal modes determine the 3D shape (conformation) of the complex and are calculated from cryo-EM images using the proposed hybrid approach in which HEMNMA is followed by the neural network. In the neural network step, a residual neural network (Resnet 34) extracts features and it is followed by a 4-layer multilayer perceptron that predicts the amplitudes of normal modes. In the presented experiments, using synthetic cryo-EM images with realistically high noise and other imperfections, the atomic-coordinate errors corresponding to the normal-mode amplitude inference errors are below 2 Å and they are on average around 0.9 Å. In this work, we showed how the obtained conformational space can be interpreted in terms of the reference-structure displacement along the trajectories identified in this space (the reference structure is required by HEMNMA and can

be an atomic structure or a 3D cryo-EM map representation with 3D Gaussian functions called pseudoatoms). In the future work, we will show how the proposed method can be used to obtain 3D reconstructions from images corresponding to similar conformations in this space, which is already possible with HEMNMA alone. Also, we will perform tests with larger synthetic data sets (for both training and test) and tests with experimental cryo-EM data of various complexes.

ACKNOWLEDGMENT

We acknowledge the support of the ANR (ANR-19-CE11-0008-01 and ANR-20-CE11-0020-03 to SJ) and the access to the HPC resources of CINES and IDRIS granted by GENCI (2019-A0070710998, AP010712190, AD011012188 to S.J.).

REFERENCES

- [1] T. Bendory, A. Bartsaghi, and A. Singer. Single-particle cryo-electron microscopy: Mathematical theory, computational challenges, and opportunities. *IEEE Signal Processing Magazine*, 37(2):58–76, Mar 2020.
- [2] J.-P. Renaud, A. Chari, C. Ciferri, W.-T. Liu, H.-W. Rémigy, H. Stark, and C. Wiesmann. Cryo-em in drug discovery: achievements, limitations and prospects. *Nature Reviews Drug Discovery*, 17:471–492, 2018.
- [3] S. Jonić. Computational methods for analyzing conformational variability of macromolecular complexes from cryo-electron microscopy images. *Current Opinion in Structural Biology*, 43:114 – 121, 2017.
- [4] P. A. Penczek, J. Frank, and C. M. T. Spahn. A method of focused classification, based on the bootstrap 3d variance analysis, and its application to ef-g-dependent translocation. *Journal of structural biology*, 154(2):184–194, 2006.
- [5] S. Scheres, H. Gao, M. Valle, G. Herman, P. P. Eggermont, J. Frank, and J. Carazo. Disentangling conformational states of macromolecules in 3d-em through likelihood optimization. *Nature Methods*, 4:27–29, 2007.
- [6] Q. Jin, C. O. S. Sorzano, J. M. De La Rosa-Trevín, J. R. Bilbao-Castro, R. Núñez-Ramírez, O. Llorca, F. Tama, and S. Jonić. Iterative elastic 3d-to-2d alignment method using normal modes for studying structural dynamics of large macromolecular complexes. *Structure*, 22(3):496–506, 2014.
- [7] A. Dashti, P. Schwander, R. Langlois, R. Fung, W. Li, A. Hosseinzadeh, H. Y. Liao, J. Pallesen, G. Sharma, V. A. Stupina, A. E. Simon, J. D. Dinman, J. Frank, and A. Ourmazd. Trajectories of the ribosome as a brownian nanomachine. *Proceedings of the National Academy of Sciences*, 111(49):17492–17497, 2014.
- [8] J. Andén and A. Singer. Structural variability from noisy tomographic projections. *SIAM Journal on Imaging Sciences*, 11(2):1441–1492, 2018.
- [9] C. O. S. Sorzano, A. Jiménez, J. Mota, J. L. Vilas, D. Maluenda, M. Martínez, E. Ramírez-Aportela, T. Majtner, J. Segura, R. Sánchez-García, Y. Rancel, L. del Caño, P. Conesa, R. Melero, S. Jonić, J. Vargas, F. Cazals, Z. Freyberg, J. Krieger, I. Bahar, R. Marabini, and J. M. Carazo. Survey of the analysis of continuous conformational variability of biological macromolecules by electron microscopy. *Acta Crystallographica Section F*, 75(1):19–32, 2019.
- [10] T. Wagner, F. Merino, M. Stabrin, T. Moriya, C. Antoni, A. Apelbaum, P. Hagel, O. Sitsel, T. Raisch, D. Prumbaum, D. Quentin, D. Roderer, S. Tacke, B. Siebolds, E. Schubert, T. Shaikh, P. Lill, C. Gatsogiannis, and S. Raunser. Sphire-cryolo is a fast and accurate fully automated particle picker for cryo-em. *Communications Biology*, 2, 2019.
- [11] R. Sanchez-Garcia, J. Segura, D. Maluenda, J. Carazo, and C. O. Sorzano. Deep consensus, a deep learning-based approach for particle pruning in cryo-electron microscopy. *IUCrJ*, 5:854 – 865, 2018.
- [12] E. D. Zhong, T. Bepler, B. Berger, and J. H. Davis. Cryodrgn: reconstruction of heterogeneous cryo-em structures using neural networks. *Nature methods*, 18(2):176–185, 2021.
- [13] M. Harastani, C. O. S. Sorzano, and S. Jonić. Hybrid electron microscopy normal mode analysis with scipion. *Protein Science*, 29(1):223–236, 2020.

- [14] J. M. De la Rosa-Trevín, A. Quintana, L. Del Cano, A. Zaldivar, I. Foche, J. Gutierrez, J. Gomez-Blanco, J. Burguet-Castell, J. Cuenca-Alba, V. Abrishami, J. Vargas, J. Otón, G. Sharov, J. L. Vilas, J. Navas, P. Conesa, M. Kazemi, R. Marabini, C. O. S. Sorzano, and J. M. Carazo. Scipion: A software framework toward integration, reproducibility and validation in 3d electron microscopy. *Journal of structural biology*, 195(1):93–99, 2016.
- [15] L. Le, A. Patterson, and M. White. Supervised autoencoders: Improving generalization performance with unsupervised regularizers. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018*, volume 31, pages 107–117. Curran Associates, Inc., 2018.
- [16] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [17] M. M. Tirion. Large amplitude elastic motions in proteins from a single-parameter, atomic analysis. *Physical Review Letters*, 77(9):1905–1908, 1996.
- [18] J. R. López-Blanco and P. Chacón. New generation of elastic network models. *Current Opinion in Structural Biology*, 37:46 – 53, 2016.
- [19] S. Jonić and C. O. S. Sorzano. Coarse-graining of volumes for modeling of structure and dynamics in electron microscopy: Algorithm to automatically control accuracy of approximation. *IEEE Journal of Selected Topics in Signal Processing*, 10(1):161–173, 2016.
- [20] F. Tama and C. L. Brooks III. Symmetry, form, and shape: guiding principles for robustness in macromolecular machines. *Annual Review of Biophysics and Biomolecular Structure*, 35:115–133, 2006.
- [21] Mingyang Lu and Jianpeng Ma. The role of shape in determining molecular motions. *Biophysical Journal*, 89(4):2395–2401, 2005.