



Deep neural-network prior for orbit recovery from method of moments

Yuehaw Khoo^a, Sounak Paul^{a,*}, Nir Sharon^b

^a Department of Statistics, University of Chicago, Chicago, USA

^b School of Mathematical Sciences, Tel Aviv University, Tel Aviv, Israel

ARTICLE INFO

Keywords:

Amortized learning
Orbit recovery problems
Method of moments
Multireference alignment
3D recovery in cryo-EM
Neural-network

ABSTRACT

Orbit recovery problems are a class of problems that often arise in practice and various forms. In these problems, we aim to estimate an unknown function after being distorted by a group action and observed via a known operator. Typically, the observations are contaminated with a non-trivial level of noise.

Two particular orbit recovery problems of interest in this paper are multireference alignment and single-particle cryo-EM modeling. In order to suppress the noise, we suggest using the method of moments approach for both problems while introducing deep neural network priors. In particular, our neural networks should output the signals and the distribution of group elements, with moments being the input. In the multireference alignment case, we demonstrate the advantage of using the NN to accelerate the convergence for the reconstruction of signals from the moments. Finally, we use our method to reconstruct simulated and biological volumes in the cryo-EM setting.

1. Introduction

Orbit recovery refers to a type of estimation problems that involve incorporating the effect of a group on a data model. The resulting solution is determined up to an arbitrary group action, meaning that the solution forms an orbit. This class of estimation problems is crucial in various fields of science and engineering, ranging from signal processing to structural biology. For instance, medical tomography often collects imaging data that undergoes unknown transformations. Along with pixel-wise noise, each image may experience rotation, translation, flipping, or other group actions in an unknown manner. This article examines two problems in this category and proposes a new approach to solving them.

The first issue we discuss is Multi-Reference Alignment (MRA), which involves estimating a signal from the observation of noisy, circularly shifted copies of it. This model, which has its origins in both signal processing [1] and structural biology [2,3], provides a foundation for exploring the relationship between the group structure, noise levels, and the possibility of recovery [4–6]. The second problem we consider is 3D volume reconstruction in Cryogenic Electron Microscopy (cryo-EM), as discussed in [7]. In cryo-EM, the goal is to retrieve a 3D volume from 2D noisy images that result from rotating the volume and applying a fixed tomographic projection. The outcome is a set of 2D images, which are usually heavily contaminated with noise.

The Method of Moments (MoM) is a classical estimation technique that has been adapted in modern forms to provide a powerful computational tool for solving large-scale problems, especially when dealing with high noise levels. The MoM consists of two stages. First, we compute the observable moments by averaging the low-order statistics of any observation. The second stage involves

* Corresponding author.

E-mail addresses: ykhoo@uchicago.edu (Y. Khoo), paulsounak96@uchicago.edu (S. Paul), nir.sharon@math.tau.ac.il (N. Sharon).

retrieving the required signal from the observable moments by analyzing the relationship between the observable and analytical moments, applying moments-fitting, and deriving the unknown parameters from it. This second stage is the focus of this study. The usage of MoM is advantageous in several ways. Its robustness is derived from the fact that noise is averaged out during the computation of observable moments. Namely, given enough data, the effect of noise can be rendered insignificant. Also, MoM gleans information about the data only through the moments, so it does not require multiple passes over the data set. This is beneficial while dealing with huge data sets, as the moment calculation from the data takes place only in the first stage and in one pass [7,8]. However, this method does have a major drawback. We can lose resolution since we are not using information from all the moments. MoM thus leads to low-dimensional reconstructions. Fortunately, our focus is mainly to recover an *ab-initio* model; hence a low-dimensional reconstruction suffices. In the case of cryo-EM, this *ab-initio* model is used as an initialization for iterative refinement algorithms, where reconstruction enables several possible conformations by further refinement [9].

This paper introduces a new version of the MoM that incorporates a neural-network prior to tackling orbit recovery problems. In particular, we demonstrate the effectiveness of the amortized MoM for the two orbit recovery problems discussed earlier: Multi-Reference Alignment (MRA) and single-particle Cryogenic Electron Microscopy (cryo-EM) modeling. Learning algorithms have recently taken a central role in cryo-EM computational methods: a deep neural network for modeling continuous heterogeneity (3DFlex) [10], *ab initio* neural reconstruction [11,12], and many other parts of the cryo-EM pipeline [13–15], to name a few. Moreover, amortized learning has recently appeared in a study for 3D modeling in cryo-EM [16]. However, noise resilience remains one of the most significant challenges in cryo-EM 3D reconstruction. The proposed “amortized” MoM technique provides a genuine alternative that addresses this challenge effectively while also addressing the additional challenge of ever-growing cryo-EM datasets.

In our method, we treat the group elements of each problem as random variables and consider them as nuisance parameters or latent variables. Rather than estimating them directly, we aim to target their density function along with the unknown signal. Our MoM incorporates neural networks to approximate the signal and distribution to achieve this. We demonstrate that in the case of MRA, a neural network can encode existing algorithms for solving the inverse problem from the moments. Moreover, we propose that the moment inversion process can be significantly improved by initializing the neural networks with those trained in a supervised manner on similar instances of the recovery problem. Our approach to the MRA problem serves as a proof-of-concept, and we extend these techniques to the case of cryo-EM.

The paper is organized as follows. Section 2 describes the problem formulation. Then, in Section 3, we present the method of moments approach for the MRA model and cryo-EM model individually as special cases of our class of estimation problems (2). Next, Section 4 introduces neural network priors for representing the volume and distribution of group elements for both models. Next, Section 5 illustrates the performance of our neural network priors in the reconstruction of various simulated as well as real-world biological volumes. Finally, we conclude with Section 6, including a summary of the next steps in this exciting line of research.

2. Problem formulation: orbit recovery

Let v be an unknown scalar-valued object defined as a function

$$v : \Omega \rightarrow \mathbb{R}, \quad (1)$$

and let G be a group with a well-defined action on v , that is $G \curvearrowright \Omega$. One class of estimation problems we are concerned with consists of the following general formulation. Our goal is to estimate the function v , where we observe N samples,

$$v_j = \mathcal{A}(g_j \circ v) + \varepsilon_j, \quad g_j \sim \rho, \quad j = 1, \dots, N, \quad (2)$$

where $\{\varepsilon_j\}_{j=1}^N$ is a set of i.i.d. random noise terms, $\mathcal{A}$ is a known operator, and $\{g_j\}_{j=1}^N$ is a set of i.i.d. random group elements distributed according to some distribution ρ on G . These are treated as latent variables or nuisance parameters for our problem since the objective is to get v . Note that one can only estimate v up to a group action, since for any estimator $\hat{v}$ and $\{\hat{g}_j\}_{j=1}^N$ for the object and latent group elements respectively, $g \circ \hat{v}$ and $\{\hat{g}_j g^{-1}\}_{j=1}^N$ give another set of equivalent estimators for any fixed $g \in G$. Hence our goal becomes the *orbit recovery* of v .

Customarily, in the lower-noise regime, where the magnitude of ε_j is smaller than the magnitude of v , a solution to (2) is obtained using the following scheme. First $g_{ij} \approx g_i g_j^{-1}$ is estimated from v_i and v_j . Then one recovers the group elements $\{g_j\}_{j=1}^N$ from the set of their ratios $\{g_i g_j^{-1}\}_{i,j=1}^N$, i.e. solving a *synchronization problem over G* . Then with a good estimation for $\{g_j\}_{j=1}^N$, we solve for v in problem (2) via solving a linear system of equations [17].

As the level of noise in the observations increases, the random noise heavily influences the alignment results so that even with the ground truth v given, one would often be fooled to assign wrong group actions with large errors [8,18]. A different approach consists of treating the group elements $\{g_j\}_{j=1}^N$ as nuisance parameters and having the signal be the primary estimation target. Therefore, when considering a high level of noise, we focus on methods that marginalize over the nuisance parameters by treating them as random variables [19]. The estimation of v can be done via maximizing the marginalized posterior distribution that has v being the random variable or using a method of moments with moments formed by averaging v_j 's such that there is no dependency on g_j 's.

3. Method of moments

The method of moments (MoM) is a classical technique to estimate parameters from observed statistics. Two recent models where MoM was already successfully employed are multireference alignment (MRA) [5,20] and cryo-EM recovery [7], where the operator $\mathcal{A}$ of (2) is the identity and a tomographic projection, respectively. The group consists of circular shifts on MRA and 3D rotations in cryo-EM recovery. Then, the m th moment is the expectation of the m -th-order tensor product of the samples with themselves, i.e., $v_j^{\otimes m}$. Interestingly, the minimal number of moments to guarantee uniqueness also determines the sample complexity — the number of samples needed, as a function of noise level, in order to have a consistent estimation, see [5,21,22]. Therefore, when studying (2), the MoM plays a significant role as a baseline for designing computational algorithms and analyzing the sample complexity.

3.1. The MRA model

We begin with the MRA model, (2), where $\mathcal{A}$ is the identity. In this situation, the unknown signal v is defined on a unit, symmetric segment $\mathcal{I} = \left[-\frac{1}{2}, \frac{1}{2}\right]$. Namely, the signal is $v : \mathcal{I} \rightarrow \mathbb{R}$, and we further assume it is a periodic, band-limited function. Let G be the group of circular translations (rotations) on $\mathcal{I}$, whose elements s_j shift v in the following manner,

$$s_j \circ v := v(\cdot - s_j). \quad (3)$$

Here, we interpret the difference as modulo the segment, namely $\cdot - s_j$ is always in $\mathcal{I}$.

We next formulate the MRA problem in the Fourier domain. For convenience, we discuss the case when there is no noise. Let $\hat{v}_j$ be the Fourier transform of v_j , in this case, a shift s_j becomes a phase, i.e.

$$\hat{v}_j(k) = \exp(iks_j) \hat{v}(k), \quad k \in [-\pi, \pi]. \quad (4)$$

The frequency k has a natural bandlimit $|k| \leq \pi$ since the signal v_j is usually provided on n discretized points in $\mathcal{I}$, where n is chosen to satisfy its Nyquist frequency. As for our observation, let K_1 be the set of n equispaced points between $[-\pi, \pi]$. Then, we have

$$\hat{v}_j(k) = \exp(iks_j) \hat{v}(k), \quad k \in K_1. \quad (5)$$

Henceforth, for brevity, we use $\hat{v}_j(K_1) = \exp(iK_1 s_j) \odot \hat{v}(K_1)$ instead of the pointwise notation, where “ $\odot$ ” denotes the Hadamard product.

Finally, in MoM for MRA, we let

$$M_F^1[\hat{v}, \rho](k_1) = \mathbb{E}_\rho \left(\frac{1}{N} \sum_{j=1}^N \hat{v}_j(k_1) \right), \quad M_F^2[\hat{v}, \rho](k_1, k_2) = \mathbb{E}_\rho \left(\frac{1}{N} \sum_{j=1}^N \hat{v}_j(k_1) \hat{v}_j(k_2)^* \right). \quad (6)$$

Here, M_F^1 and M_F^2 are functions of $\hat{v}, \rho$. The goal is to retrieve $\hat{v}$ from unbiased estimators $\hat{M}_F^1, \hat{M}_F^2$ of M_F^1, M_F^2 when having noisy data via matching the moments.

3.2. The cryo-EM model

The problem (2) also serves as a simplified model of single-particle cryo-EM, where the operator $\mathcal{A}$ is a tomographic projection along a fixed axis. Cryo-EM is a prominent method for determining the high-resolution 3-D structure of biological macromolecules from its 2-D noisy projection images [9].

For the cryo-EM model, we denote by $v : \mathbb{R}^3 \rightarrow \mathbb{R}$ the Coulomb potential of the 3D volume we aim to determine, where we assume that v is compactly supported in a ball of radius $\frac{1}{2}$ around the origin, that is inside $\mathcal{I}^3$. We define the composition of R_j with the volume v as

$$R_j \circ v(x, y, z) = v(R_j^T [x \ y \ z]^T), \quad (x, y, z) \in \mathcal{I}^3, \quad (7)$$

viewing R_j as a 3×3 matrix in the right hand side of (7) since $\text{SO}(3) \subset \mathbb{R}^{3 \times 3}$. Let $\mathcal{P} : \mathbb{R}^3 \rightarrow \mathbb{R}^2$ be the operator that projects a 3D volume along the z axis to a 2D image, i.e.

$$\mathcal{P} \circ v(x, y) = \int_{-\infty}^{\infty} v(x, y, z) dz, \quad (x, y) \in \mathcal{I}^2. \quad (8)$$

Then, a standard image formation model in the absence of noise, after filtering the effect of the contrast transfer function (CTF), image cropping, and centering, is (see [9,23]),

$$v_j = \mathcal{P} \circ R_j \circ v, \quad j = 1, \dots, N, \quad (9)$$

where $R_j \in \text{SO}(3)$ are the unknown group elements. To avoid the computationally intensive integration in (8), we reformulate our problem in the Fourier domain. There, we can exploit the Fourier Slice Theorem to speed up computation significantly. We define $\hat{v} : [-\pi, \pi]^3 \rightarrow \mathbb{C}$ as the Fourier transform of v , and $S : [-\pi, \pi]^2 \rightarrow \mathbb{C}$ as the slice operator given as

$$S \circ \hat{v}(k_x, k_y) = \hat{v}(k_x, k_y, 0), \quad (10)$$




Fig. 1. Overview of our MRA pipeline: The encoder ξ_θ takes moments $(\hat{M}_F^1, \hat{M}_F^2)$ as input, and outputs $z_\rho \in \mathbb{R}^n$, approximating a discretized probability density $\rho(X_1)$, and $z_v \in \mathbb{R}^n$ that approximates a discretized Fourier signal $\hat{v}(K_1)$. Next, we use z_ρ and z_v to create $(M_F^1[z_v, z_\rho](K_1), M_F^2[z_v, z_\rho](K_1, K_1))$ via Eq. (25), which we then compare with the inputs to the encoder, i.e., $(\hat{M}^1, \hat{M}^2)$ via the loss function $\mathcal{L}_{\text{recon}}$ (26).

i.e., $S \circ \hat{v}$ is obtained by slicing $\hat{v}$ across the plane given by $z = 0$. Then, the Fourier Slice Theorem states that:

$$\mathcal{F}_{2D} \circ \mathcal{P} \circ R = S \circ R \circ \mathcal{F}_{3D}, \quad (11)$$

where $R \in \text{SO}(3)$, $\mathcal{F}_{2D}$ and $\mathcal{F}_{3D}$ are the 2D and 3D Fourier transformations, respectively. Therefore, in the no-noise setting, the equivalent of (9) becomes,

$$\hat{v}_j(k_x, k_y) = S \circ R_j \circ \hat{v}(k_x, k_y), \quad (k_x, k_y) \in [-\pi, \pi]^2, \quad (12)$$

where $\hat{v}_j$ is the Fourier transform of v_j . Let K_2 be a grid of n^2 equispaced points on $[-\pi, \pi]^2$, flattened as a one-dimensional vector. Now our observations are

$$\hat{v}_j(K_2) = S \circ R_j \circ \hat{v}(K_2), \quad (13)$$

and the associated moments are

$$\begin{aligned} M_F^1[\hat{v}, \rho](k_x, k_y) &= \mathbb{E}_\rho \left(\frac{1}{N} \sum_{j=1}^N \hat{v}_j(k_x, k_y) \right) \\ M_F^2[\hat{v}, \rho](k_x, k_y, k'_x, k'_y) &= \mathbb{E}_\rho \left(\frac{1}{N} \sum_{j=1}^N \hat{v}_j(k_x, k_y) \hat{v}_j(k'_x, k'_y)^* \right). \end{aligned} \quad (14)$$

We aim to retrieve $\hat{v}$ by matching the moments $M_F^1[\hat{v}, \rho](K_2, K_2)$ and $M_F^2[\hat{v}, \rho](K_2, K_2)$ with some unbiased estimators $\hat{M}_F^1, \hat{M}_F^2$ when having noisy data.

4. Neural network priors for method of moments

This section presents neural network (NN) approaches for reconstructing the signal v and distribution ρ in MRA and cryo-EM settings. The general strategy is to view both the signal and distribution as being mapped by a NN from the estimated moments $\hat{M}_F^1$ and $\hat{M}_F^2$, as various previous works have shown that for MRA, $\hat{M}_F^1, \hat{M}_F^2$ are generically sufficient statistics for estimating the unknowns [5], while for cryo-EM, they have enough information for recovering a low-resolution reconstruction [7]. In the MRA case, we design an encoder that can map the empirical moments to discretized signal and density. In the cryo-EM case, we further design an encoder–decoder structure that allows us to take the moments as input and give a continuous representation of a 3D volume.

4.1. NN for MRA

In the MRA case, we want to train a NN that can take the moments as inputs and give the signal and density as outputs, which can further be used to initialize an iterative reconstruction algorithm. More precisely, we define $F \in \mathbb{C}^{n \times n}$ as the matrix representation of a normalized Fourier transform where $F^* F = I_n$, and X_1, K_1 as sets of n equispaced points on $\mathcal{I} = [-\frac{1}{2}, \frac{1}{2}]$ and $[-\pi, \pi]$ respectively. The main component is an encoder, i.e., a neural network ξ_θ , that takes the moments $\hat{M}_F^1, \hat{M}_F^2$ as inputs and outputs $z_\rho, z_v \in \mathbb{R}^n$, where z_ρ approximates a discretized density $\rho(X_1)$ and z_v approximates a discretized signal $\hat{v}(K_1)$.

The encoder $\xi_\theta := (\xi_\theta^v, \xi_\theta^\rho)$ consists of two NN ξ_θ^v and ξ_θ^ρ , which are two 1D convolutional NNs (CNNs) that take $\hat{M}_F^1 \in \mathbb{C}^n, \hat{M}_F^2 \in \mathbb{C}^{n \times n}$ as input vector fields supported on n grid points. Fig. 1 provides an overview of our pipeline for MRA. While the details of the architectures are provided in Appendix, here we provide motivations as for how a CNN has the capability to learn a mapping from the moments $\hat{M}_F^1 \in \mathbb{C}^n, \hat{M}_F^2 \in \mathbb{C}^{n \times n}$ to $\hat{v}(K_1)$. For simplicity, suppose $|\hat{v}(k)| = 1$. Using the definitions in (6) and the fact that translating v by s is equivalent to letting $\hat{v}(k) \rightarrow \hat{v}(k) \exp(iks)$, one can show that

$$M_F^2[\hat{v}, \rho](k_1, k_2) = \hat{v}(k_1) \hat{\rho}(k_1 - k_2) \hat{v}(k_2)^*, \quad (15)$$

as in [5]. In this case, $M_F^2[\hat{v}, \rho](K_1, K_1)$ admits the eigendecomposition

$$\begin{aligned} M_F^2[\hat{v}, \rho](K_1, K_1) &= \text{diag}(\hat{v}(K_1)) F^* (F[\hat{\rho}(k_1 - k_2)]_{k_1, k_2} F^*) F \text{diag}(\hat{v}(K_1)^*) \\ &= [\text{diag}(\hat{v}(K_1))] F^* \text{diag}(\rho(X_1)) [F \text{diag}(\hat{v}(K_1)^*)] \end{aligned} \quad (16)$$

since $[F \text{diag}(\hat{v}(K_1)^*)]$ is an orthogonal matrix (due to the assumption $|\hat{v}(k)| = 1$). From this form, it is clear that the eigenvalues of $M_F^2[\hat{v}, \rho](K_1, K_1)$ are $\rho(X_1)$ and furthermore, the eigenvectors are $F \text{diag}(\hat{v}(K_1)^*)$. Since the spectral information of the second moments contains information concerning the signal and density, if an NN can mimic a spectral method, then it can learn the mapping from moments to the signal and density.

The form of $M_F^2[\hat{v}, \rho](K_1, K_1)$ in (16) suggests that it is a circulant matrix. Therefore if we want to devise a neural network that takes $\hat{M}_F^2 = M_F^2[\hat{v}, \rho](K_1, K_1)$ (when there is no noise) as input and output the eigenvectors $F \text{diag}(\hat{v}(K_1)^*)$, we can have a neural network, composed of 1D convolutional layers, that takes $\hat{M}_F^2$ as a 1D n -dimensional vector field supported on n grid points. For example, to get an eigenvector of $M_F^2[\hat{v}, \rho](K_1, K_1)$, a convolutional layer $l_1 : \mathbb{C}^n \rightarrow \mathbb{C}^n$ can take the form

$$l_1(u) = \frac{\hat{M}_F^2 u}{\|\hat{M}_F^2 u\|_2}. \quad (17)$$

One can think about $\hat{M}_F^2$ as the weights of the convolutional layer l_1 , and the division by $\|\hat{M}_F^2 u\|_2$ as some nonlinearities in the NN. Repeated applications of l_1 , gives an eigenvector of $M_F^2[\hat{v}, \rho](K_1, K_1)$. After obtaining an eigenvector, say $F(\cdot, 1)\hat{v}(K_1(1))$ where $F(\cdot, 1)$ is the first column of F , the NN can simply apply some pointwise nonlinearities layer $l_2 : \mathbb{C}^n \rightarrow \mathbb{C}^n$ that performs

$$l_2(u(i)) = \frac{u(i)}{F(i, 1)}, \quad i \in [n]. \quad (18)$$

Putting these elements together into a deep NN, i.e., $l_2 \circ l_1 \circ \dots \circ l_1$ should give $\hat{v}(K_1(1))$. Similar operations can be carried out for other eigenvectors. We also use a similar structure for ξ_θ^ρ to output z_ρ that approximates $\rho(X_1)$, since it is clear that if $u = l_2 \circ l_1 \circ \dots \circ l_1(\hat{M}_F^2)$ is an eigenvector of $\hat{M}_F^2$, applying another nonlinearity

$$l_3(u) = \langle u, \hat{M}_F^2 u \rangle \quad (19)$$

gives the eigenvalue of $\hat{M}_F^2$ which contains information of $\rho(X_1)$ (as shown in (16)).

4.2. NN for cryo-EM

We make some alterations to our MRA architecture for cryo-EM reconstruction since we need to output a continuous representation of the volume to facilitate computing the moments involving the reconstructed volume. Just as in the case of MRA, we have an encoder ξ_θ^ρ that outputs information regarding the density. More precisely, let $Q \subset \text{SO}(3)$ be a set of quadrature points on $\text{SO}(3)$ and $q = |Q|$. We want $\xi_\theta^\rho : (\hat{M}_F^1, \hat{M}_F^2) \rightarrow z_\rho$ ($\hat{M}_F^1, \hat{M}_F^2$) are estimators of (14) where z_ρ should approximate $(\rho(R))_{R \in Q}$ and ρ is a density on $\text{SO}(3)$.

However, unlike the case of MRA, we now want to have a continuous representation of the Fourier volume where the benefit is explained as follows. Let $\hat{v}_\phi : \mathbb{R}^3 \rightarrow \mathbb{C}$ be an NN that represents a volume on the Fourier domain, and K_2 be n^2 equispaced points on $[-\pi, \pi]^2$. Suppose $\hat{v}_\phi = \hat{v}$ and $z_\rho = \rho(Q)$, one can evaluate $M_F^1[\hat{v}, \rho](K_2, K_2)$, $M_F^2[\hat{v}, \rho](K_2, K_2)$ defined in (14) approximately via the quadrature rule

$$\begin{aligned} M_F^1[\hat{v}_\phi, z_\rho](K_2) &\approx \sum_{j=1}^q z_\rho(j) S \circ Q(j) \circ \hat{v}_\phi(K_2), \\ M_F^2[\hat{v}_\phi, z_\rho](K_2, K_2) &\approx \sum_{j=1}^q z_\rho(j) \left(S \circ Q(j) \circ \hat{v}_\phi(K_2) \right) \otimes \left(S \circ Q(j) \circ \hat{v}_\phi(K_2) \right), \end{aligned} \quad (20)$$

where, by an abuse of notation, we think about $z_\rho = \rho(Q)$, i.e., the density ρ discretized on Q , as ρ itself and $Q(j)$ is an element in the set Q . For simplicity, in this paper, we consider a quadrature rule with uniform quadrature weights, as seen in (20). It is clear that having a continuous $\hat{v}_\phi$ allows us to obtain $\hat{v}_\phi(Q(j)^T(k_x, k_y, 0))$ for any $(k_x, k_y) \in K_2$ easily.

Note that we also allow the flexibility to have an encoder ξ_θ^v just as in the case of MRA. In this case, $\xi_\theta^v : (\hat{M}_F^1, \hat{M}_F^2) \rightarrow z_v$ where z_v is some latent variable of the volume. In this case, we simply let $\hat{v}_\phi : \mathbb{R}^{3+|z_v|} \rightarrow \mathbb{C}$ where the extra inputs of $\hat{v}_\phi$ corresponds to the output of ξ_θ^v . The neural network pipeline we devise is shown in Fig. 2, where $\xi_\theta = \begin{pmatrix} \xi_\theta^\rho, \xi_\theta^v \end{pmatrix}$. As for the architecture of $\xi_\theta^\rho, \xi_\theta^v$, we adopt the type of architecture we use in Section 4.1, though one should be able to improve it according to the structure of the cryo-EM problem. The details of ξ_θ and $\hat{v}_\phi$ are given in Appendix.

5. Numerical examples

This section presents the results of numerical experiments with our Method of Moments algorithm with NN prior done using PyTorch [24].




Fig. 2. Overview of our cryo-EM pipeline: The encoder ξ_θ takes moments $(\hat{M}_F^1, \hat{M}_F^2)$ as input, and outputs $z_\rho \in \mathbb{R}^{|Q|}$, approximating a discretized probability density $(\rho(R))_{R \in Q}$ for some fixed set of quadrature points $Q \subset \text{SO}(3)$. Next, we create copies of the grid K_2 rotated corresponding to the elements of Q and input them to our neural representation $\hat{v}_\phi$, which outputs corresponding slices of a running estimate of $\hat{v}$. These slices $\{S \circ Q(j) \circ \hat{v}_\phi(K_2)\}_j$ along with z_ρ are used to create $(M_F^1[\hat{v}_\phi, z_\rho](K_2), M_F^2[\hat{v}_\phi, z_\rho](K_2, K_2))$ via equation (20), which we then compare with the inputs to the encoder, i.e., $(\hat{M}_F^1, \hat{M}_F^2)$ via the loss function $\mathcal{L}_{\text{recon}}$ in (28). Optionally, ξ_θ can also be used to output an extra z_v , a latent variable of $\hat{v}$ that can be inputted into $\hat{v}_\phi$.

5.1. MRA

We first present results using the method for MRA in Section 4.1. There are two phases when using the NN detailed in Section 4.1: training phase consisting of supervised learning (Section 5.1.1) and reconstruction phase consisting of unsupervised learning (Section 5.1.2).

For evaluation purposes, we define the reconstruction error (also referred to as relative error) of an estimator $u \in \mathbb{R}^n$ of a signal v (or a distribution ρ) discretized at X_1 , to be

$$\inf_{s \in I} \frac{\|s \circ v(X_1) - u\|_F}{\|v(X_1)\|_F}, \quad \inf_{s \in I} \frac{\|s \circ \rho(X_1) - u\|_F}{\|\rho(X_1)\|_F}. \quad (21)$$

Note that in this case, we identify the group with X_1 . In addition, we define the relative errors for any moment estimators A_1, A_2 for the first and second moments, respectively, as

$$\frac{\|M_F^1[\hat{v}, \rho](K_1) - A_1\|_F}{\|A_1\|_F}, \quad \frac{\|M_F^2[\hat{v}, \rho](K_1, K_1) - A_2\|_F}{\|A_2\|_F}. \quad (22)$$

5.1.1. Supervised training phase

The goal of the training phase is to use the encoder ξ_θ to predict $\rho(X_1)$ and $v(X_1)$ given moments. In other words, we demonstrate that the moment inversion map can be learned by neural networks in a supervised way.

To this end, we draw the signal v and the density ρ from a distribution. After forming their corresponding first and second moments, we train our encoder ξ_θ in a supervised way to take inputs of the form $(M_F^1[\hat{v}, \rho](K_1), M_F^2[\hat{v}, \rho](K_1, K_1))$ and output $(\rho(X_1), \hat{v}(K_1))$. In our experiments, we let the distribution of v and ρ be a mixture of Gaussians on the interval I , where we repeat our training procedure separately for a different number of Gaussians. We take 1750k of input-output pairs to do the training. We compute test error on 250k of samples using (21).

We now discuss the hyperparameters for training. We train the encoders ξ_θ^ρ and ξ_θ^v separately; let us consider ξ_θ^ρ . We take the training set and feed the moments pairs $(M_F^1[\hat{v}, \rho](K_1), M_F^2[\hat{v}, \rho](K_1, K_1))$ to ξ_θ^ρ , which outputs corresponding z_ρ for each pair as a prediction for $\rho(X_1)$. We train ξ_θ^ρ over a total of 30,000 epochs with learning rates of 10^{-4} , 10^{-5} and 10^{-6} over 10,000 epochs successively. We then repeat the same process for ξ_θ^v .

Table 1 summarizes the average relative error on the training and test sets, using (21), while evaluating our trained encoders on mixtures of different numbers of Gaussians. The left and right columns of Fig. 3 show some comparisons of the encoder output (z_ρ, z_v) with ground truth $(\rho(X_1), \hat{v}(K_1))$ during prediction time.

5.1.2. Reconstruction phase

In the previous section, we discussed our process of training the encoder ξ_θ in a supervised way such that it learns the moment inversion map. A useful application of this trained encoder is when supplied with new, possibly noisy, moments $(\hat{M}_F^1, \hat{M}_F^2)$, we can use its outputs as a good initialization for further refinement. In this section, we demonstrate that this procedure leads to faster convergence.




Fig. 3. Predictions for the distribution ρ (Left) and volume v (Right), outputted by trained encoders ξ_{θ}^{ρ} and ξ_{θ}^v respectively, for ρ, v being mixture of 2 Gaussians. The solid lines are the ground truth ρ and v , while the dotted lines are the corresponding predictions by an NN.

Table 1

Average reconstruction errors (defined in (21)) of predictions z_{ρ} and z_v on training and test sets for mixtures of Gaussians.

No. of Gaussians	z_{ρ} (Train error)	z_{ρ} (Test error)	z_v (Train error)	z_v (Test error)
1	0.042	0.048	0.048	0.052
2	0.121	0.141	0.156	0.170
3	0.177	0.195	0.180	0.206

We first talk about how we obtain the estimators $\hat{M}_F^1, \hat{M}_F^2$ from observations of the form

$$v_j = s_j \circ v(X_1) + \epsilon_j, \quad j = 1, \dots, N \quad (23)$$

where $\epsilon_j \sim N(0, \sigma^2 I_n)$. Let $F \in \mathbb{C}^{n \times n}$ again be the Fourier matrix, we form unbiased moment estimators of the form

$$\hat{M}_F^1 = \frac{1}{N} \sum_{j=1}^N F v_j, \quad \hat{M}_F^2 = \frac{1}{N} \sum_{j=1}^N (F v_j)(F v_j)^* - \sigma^2 I_n \quad (24)$$

by subtracting a constant term on the diagonal of the empirical second moment. These are used as input to the trained encoder ξ_{θ} for prediction.

Notice that the solution to the MRA problem has a global translation ambiguity. Therefore, it is possible for the encoders $\xi_{\theta}^v, \xi_{\theta}^{\rho}$, to output an approximation to signal v and density ρ up to some arbitrary translations. While this is not a big issue if the predicted signal is all we want, it becomes an issue if we want to refine the predictions further. More precisely, before deploying the encoder for refinement with new incoming moments $\hat{M}_F^1, \hat{M}_F^2$, we conduct an alignment to ensure that the outputs $z_{\rho} = \xi_{\theta}^{\rho}(\hat{M}_F^1, \hat{M}_F^2)$ and




Fig. 4. Plots of logarithms (with base 10) of Sum of relative errors (defined in (22)) for $\hat{M}_F^1$ and $\hat{M}_F^2$ across 3000 iterations (Top), and Reconstruction error (defined in (21)) across 3000 iterations (Bottom); averaged over 20 reconstructions of $(\rho(X_1), \hat{v}(K_1))$ pairs drawn from the family of a mixture of 2 Gaussians. In both plots, the blue curve corresponds to the scenario where the encoder underwent supervised training, while the orange corresponds to the scenario where it did not.

$z_v = \xi_\theta^v(\hat{M}_F^1, \hat{M}_F^2)$, upon forming

$$\begin{aligned} M_F^1[z_v, z_\rho](K_1) &= \sum_{j=1}^n z_\rho(j) \exp(-iK_1 s(j)) \odot z_v, \\ M_F^2[z_v, z_\rho](K_1, K_1) &= \sum_{j=1}^n z_\rho(j) (\exp(-iK_1 s(j)) \odot z_v) (\exp(-iK_1 s(j)) \odot z_v)^*, \end{aligned} \quad (25)$$

matches the inputs $(\hat{M}_F^1, \hat{M}_F^2)$ of the encoder. Here by abuse of notation, we treat z_v, z_ρ as a continuous object and apply the functionals M_F^1, M_F^2 to them. With an alignment, we can make sure the initial loss

$$\mathcal{L}_{\text{recon}} = \left\| \hat{M}_F^1 - M_F^1[z_v, z_\rho](K_1) \right\|_F + \lambda \left\| \hat{M}_F^2 - M_F^2[z_v, z_\rho](K_1, K_1) \right\|_F, \quad (26)$$

is small. Recall that $z_v = \xi_\theta^v(\hat{M}_F^1, \hat{M}_F^2)$, $z_\rho = \xi_\theta^\rho(\hat{M}_F^1, \hat{M}_F^2)$, we further optimize the NN parameters θ to refine z_v, z_ρ with the loss in (26).

We now show the results of the deployment of our architecture ξ_θ when working with noisy moments. We take 20 different $\hat{M}_F^1, \hat{M}_F^2$, and determine z_v, z_ρ by minimizing (26) over the parameters of ξ_θ . The relative errors (defined in (21) and (22)) of the reconstructed $(\rho(X_1), \hat{v}(K_1))$ and the moments are plotted in Fig. 4. The errors are averaged over 20 different instances of (ρ, v) combinations from mixtures of 2 Gaussians, and the empirical moments are formed from 1000k observations for each pair of (ρ, v) as in (23), with Gaussian noise $\sigma = 1.0$. Depending on whether the encoder underwent supervised training, we observe the trajectory of this “average” reconstruction error to be different. Fig. 4 illustrates that the average reconstruction error indeed converges faster when the encoder is trained in a supervised phase.

5.2. Cryo-EM

We now present the results using our method for cryo-EM as illustrated in Section 4.2. Again for evaluation purposes, the relative error for an estimate $u \in \mathbb{R}^{n^3}$ of a signal v discretized at n^3 equispaced points X_3 on T^3 , is defined as

$$\inf_{R \in \text{SQ}(3)} \frac{\left\| R \circ v(X_3) - u \right\|_F}{\left\| v(X_3) \right\|_F}. \quad (27)$$

The relative errors for moment estimators of the first and second moments are defined analogously to (22).

While we do not describe any supervised training phase like in the MRA case, our architecture keeps this option open. We believe that even for cryo-EM, it would be possible to train our encoder ξ_θ^ρ in a supervised way to learn the moment inversion map, i.e., to take inputs of the form $(M_F^1[\hat{v}, \rho](K_2), M_F^2[\hat{v}, \rho](K_2, K_2))$ and predict $(\rho(R))_{R \in Q}$ for training and reconstruction, where Q is the set of quadrature points on $\text{SO}(3)$ defined in 4.2. It would also be possible to train ξ_θ^v such that it outputs a discretized approximation of the volume from the moments, or at least some vector containing important feature information about it.

The reconstruction is carried out by optimizing the NN parameters θ and ϕ of our encoder $z_\rho = \xi_\theta^\rho(\hat{M}_F^1, \hat{M}_F^2)$ and neural representation $\hat{v}_\phi$, respectively, to minimize the loss function

$$\mathcal{L}_{\text{recon}} = \left\| \hat{M}_F^1 - M_F^1[\hat{v}_\phi, z_\rho](K_2) \right\|_F + \lambda \left\| \hat{M}_F^2 - M_F^2[\hat{v}_\phi, z_\rho](K_2, K_2) \right\|_F. \quad (28)$$

During reconstruction, one of the challenges we face is fixing a good set $Q \subset \text{SO}(3)$ on which we shall use a quadrature rule with uniform weights to evaluate the functionals M_F^1, M_F^2 , as described in (20). In our experiments, we do so in two steps. First, we choose a q_1 -point spherical design on S^2 , see, e.g., [25]. A q_1 -point spherical t -design is a finite set of points with cardinality q_1 on




Fig. 5. (Left) 1000 points sampled from a mixture of eight von Mises–Fisher random variables shown in different colors, and (Right) 100-point 13-design plotted on a 3D unit sphere.




Fig. 6. (Left) A clean projection, and (Right) its noisy counterpart with noise level $\sigma = 0.5$ as defined in (29), for EMD-0409.

S^2 , such that their quadrature over S^2 with uniform unit weights is exact for any polynomial (spherical harmonics) with degree $\leq \iota$. Then, for each point of the design, treating the axis connecting that point to the center as viewing direction, we consider in-plane rotations with q_2 equally spaced angles in $[0, 2\pi)$ radians. This gives us a set Q with $|Q| = q_1 q_2$ quadrature points on $SO(3)$. In experiments, we take $q_1 = 100$ and $q_2 = 12$ for a total of $|Q| = 1200$ quadrature points. To illustrate these quadrature points, we use a 100-point 13-design on S^2 as the set of viewing directions, as seen in the right side of Fig. 5.

We now discuss our data generation process for cryo-EM and the moment estimators to be inputted into the encoders. In practice, given real observations of the form

$$v_j = \mathcal{P} \circ R_j \circ v(X_2) + \epsilon_j, \quad j = 1, \dots, N \quad (29)$$

where $\epsilon_j \sim N(0, \sigma^2 I_{n^2})$ and X_2 is n^2 equispaced points on $\mathcal{I}^2$, we could form unbiased moment estimators

$$\hat{M}_F^1 = \frac{1}{N} \sum_{j=1}^N F_2 v_j, \quad \hat{M}_F^2 = \frac{1}{N} \sum_{j=1}^N (F_2 v_j) \otimes (F_2 v_j) - \sigma^2 I_{n^2}, \quad (30)$$

letting $F_2 \in \mathbb{C}^{n^2 \times n^2}$ be the two-dimension Fourier transform matrix. Clean observations v_j are depicted alongside their noisy counterparts in Figs. 6 and 7.

We next discuss our choices of ground truth volumes v and rotational distributions ρ . For our experiments, we use three volumes: EMD-0409 and EMD-25892 taken from the Electron Microscopy Data Bank (EMDB); and a mixture of four Gaussians not lying on the same plane in three dimensions. The dimensions of EMD-0409 are $128 \times 128 \times 128$ with voxel size $1.117 \text{ \AA}$, while the dimensions of EMD-25892 are $320 \times 320 \times 320$ with voxel size $1.68 \text{ \AA}$. Both volumes were downsampled to $63 \times 63 \times 63$ and scaled to have norm 1. The mixture of Gaussians has dimensions $25 \times 25 \times 25$, whose voxel size is taken to be $1 \text{ \AA}$ since it is a simulated volume. We represent the ground truth using $\hat{v}_\phi$, and report the approximation error as defined in (27), between the original and this neural ground truth as 0.043 for EMD-0409, 0.076 for EMD-25892, and 0.004 for the mixture of Gaussians. These NN-approximated




Fig. 7. (Left) A clean projection, and (Right) its noisy counterpart with noise level $\sigma = 0.5$ as defined in (29), for EMD-25892.

Table 2

Final relative errors of moment estimates $\hat{M}_F^1$ and $\hat{M}_F^2$ after reconstruction phase.

Volume	Relative error in $\hat{M}_F^1$	Relative error in $\hat{M}_F^2$
EMD-0409	0.003	0.013
EMD-25892	0.007	0.035
Mixture of Gaussians	0.007	0.016

Table 3

Optimal resolutions between ground truth volumes and their reconstructions.

Volume	Resolution (in Å)
EMD-0409	16.86
EMD-25892	21.52
Mixture of Gaussians	4.45

volumes are then used as the ground truths for the rest of the simulations. The ground truth distribution of rotations ρ is chosen in the following way. The viewing directions are distributed as a mixture of 8 von Mises–Fisher distributions with different mean directions μ and concentration parameters κ , respectively, to ensure a sufficiently non-uniform distribution on S^2 . 1000 points from this distribution are shown on the left side of Fig. 5. The in-plane rotations are uniform on $[0, 2\pi)$ and independent of the viewing directions. We then create moment estimators from $N = 5,000,000$ noisy observations with noise level $\sigma = 0.5$ using (30), where a neural slice approximates $F_2 v_j$.

We run our algorithm with learning rates 10^{-5} and 10^{-6} successively for 10,000 epochs each, to minimize the loss function in (28). The reconstructed volumes are visualized in Figs. 8, 9, and 10, alongside their corresponding neural ground truth volumes for EMD-0409, EMD-25892, and mixture of Gaussian volumes, respectively. Table 2 shows the relative errors of our moments from the reconstructed volumes, defined analogously to (22), at the end of our reconstruction.

Finally to evaluate the quality of reconstruction, we first align the reconstructed volumes with the ground truth. For that purpose, we run the algorithm for aligning three-dimensional density maps in [26] multiple times and pick the best alignment. We then calculate the Fourier Shell Correlation (FSC) between the ground truth volumes and their corresponding aligned reconstructions. We denote the resolution of the reconstructed volume as the point where the FSC curve goes below 0.5. The final resolutions between the ground truths and reconstructed volumes are provided in Table 3.

6. Conclusion and outlook

Single-particle cryo-EM is a prominent method for determining the atomic-resolution 3D structure of biological macromolecules. This technique underwent a “resolution revolution” a decade ago [28], and three of its pioneers were awarded the 2017 Nobel Prize in Chemistry [29]. These days, cryo-EM provides researchers access to some of the molecules’ tiniest and most essential building blocks. In this paper, we addressed the reconstruction problem in cryo-EM as well as one of its simpler versions, namely, multireference alignment. Both cryo-EM and MRA fall under the class of orbit recovery problems.

Although deep NN-based methods have been successfully used in maximum likelihood estimation for orbit recovery problems, they have not historically exploited the benefits offered by the MoM, like noise resilience, due to the central limit theorem when averaging data. In this paper, we take a first step towards using NNs for solving moment systems in orbit recovery problems. In the case of MRA, we demonstrate theoretically and numerically that a map can be learned to take moments as input and output the




Fig. 8. Ground truth volume (in gray) and reconstructed volume (in yellow) for the EMD-0409 volume, visualized using UCSF Chimera [27].




Fig. 9. Ground truth volume (in gray) and reconstructed volume (in yellow) for the EMD-25892 volume, visualized using UCSF Chimera [27].




Fig. 10. Two views of recovery of a mixture of Gaussians. Ground truth volume (in gray) and reconstructed volume (in yellow) for a mixture of 4 Gaussians in three dimensions, visualized using UCSF Chimera [27].

signal and density of translations, and develop novel neural network architectures for the same. This map can then be used as a deep NN prior to accelerating convergence in unsupervised reconstruction from new incoming moments.

We also apply this approach to cryo-EM with encouraging results, but further work is needed to demonstrate the superiority of supervised learning and tackle more general cryo-EM models, like those dealing with small translations in addition to the rotations, and further image contamination due to aberrations (which would involve accounting for contrast transfer functions). Supervised learning would effectively enable low-dimension reconstruction of volumes near-instantly and would serve as an inexpensive and time-efficient method of generating *ab-initio* models for iterative refinement algorithms. Other future work includes investigating the use of higher-order moments to improve reconstruction accuracy and parallelizing the model on multiple GPUs to enable reconstruction with larger images and improve speed and accuracy. Additionally, tackling more general cryo-EM models will bring us closer to operating on real-world datasets.

Data availability

Code hosted on Github.

Acknowledgments

YK is thankful to DOE, USA for funding DE-SC0022232. NS is partially supported by the NSF-BSF award 2019752 and DFG, USA award 514588180.




Fig. A.11. Architecture of ξ_θ^ρ in case of MRA.

Appendix. Architecture of neural networks

In this appendix, we describe the details of the NN for both MRA and cryo-EM. To facilitate the discussion, we first define $\text{conv1D}_{w,c}$ to be a 1D convolutional layer with periodic padding, kernel window size w and channel number c . In a similar way, we also denote a 2D convolutional layer with window size $w \times w$ and channel number c as $\text{conv2D}_{w,c}$. Furthermore we define $\text{input1D}_{\ell,c}$ to be an input layer that prepares the input as a length ℓ 1D vector field with channel number c . We then define a fully connected layer full_w that takes an input vector field and output a vector with size w . The nonlinearities we use in this paper are leaky ReLU (LReLU) nonlinear activation with parameter 0.02, $\tanh(\cdot)$ function, and just linear activation (without nonlinearities). We make no distinction between real or complex input, since changing real to complex input only requires doubling the input or output channel number.

A.1. MRA

In the MRA case, we present the proposed architecture for the encoder ξ_θ . An illustration of ξ_θ^ρ is presented in Fig. A.11, and the same architecture is used for ξ_θ^ν . The input layers $\text{input1D}_{n,1}$ and $\text{input1D}_{n,n}$ take the moments as inputs. After a few layers of conv1D , we stack the output of the upper branch and lower branch in Fig. A.11 together into a 1D vector field of length n and 6 channels. Then after a few more layers of CNN conv1D and fully connected layers full , we output z_ρ .

A.2. Cryo-EM

The encoders ξ_θ^ρ and ξ_θ^ν are very similar to the one presented in Fig. A.11 for MRA, except we replace all conv1D with conv2D with the same window sizes and channel numbers. As for $\hat{v}_\phi$, currently, it is chosen to be the FourierNet of [16]. FourierNet finds success in representing the Fourier transforms of three-dimensional volumes of molecules and other volumes arising in nature, with values that often span multiple orders of magnitude. The main point of such a representation is that, instead of approximating $v(x)$ directly by an NN, it is often easier to approximate its Fourier coefficients $\hat{v}(x)$ by an NN on k -space when $v(x)$ exhibits oscillatory patterns. This is also similar to the approach taken in [30] for solving high-frequency wave equations. More precisely, it lets

$$\hat{v}(k) \approx \hat{v}_\phi(k) = a_{\phi_1}(k) \exp(ib_{\phi_2}(k)) \quad (\text{A.1})$$

with two NNs $a_{\phi_1}(k) \in \mathbb{C}$ and $b_{\phi_2}(k) \in \mathbb{C}$ where a_{ϕ_1} gives the amplitude of the Fourier coefficients and b_{ϕ_2} gives the phase variations. By representing v in Fourier domain instead of real domain, one can bypass the oscillatory pattern caused by the Fourier series $\exp(ikx)$ in $v(x) = \sum_k \hat{v}(k) \exp(ikx)$. More details regarding the architecture, its effectiveness, and its memory requirements are provided in [16].

References

- [1] Joris Portegies Zwart, René van der Heiden, Sjoerd Gelsema, Frans Groen, Fast translation invariant classification of HRR range profiles in a zero phase representation, *IEE Proc.-Radar, Sonar Navig.* 150 (6) (2003) 411–418.
- [2] Sjors HW Scheres, Mikel Valle, Rafael Nuñez, Carlos OS Sorzano, Roberto Marabini, Gabor T Herman, Jose-Maria Carazo, Maximum-likelihood multi-reference refinement for electron microscopy images, *J. Mol. Biol.* 348 (1) (2005) 139–149.
- [3] Douglas L. Theobald, Phillip A. Steindel, Optimal simultaneous superpositioning of multiple structures with missing data, *Bioinformatics* 28 (15) (2012) 1972–1979.
- [4] Alexander S. Wein, Statistical Estimation in the Presence of Group Actions (Ph.D. thesis), Massachusetts Institute of Technology, 2018.
- [5] Emmanuel Abbe, Tamir Bendory, William Leeb, João M Pereira, Nir Sharon, Amit Singer, Multireference alignment is easier with an aperiodic translation distribution, *IEEE Trans. Inform. Theory* 65 (6) (2018) 3565–3584.
- [6] Ankur Moitra, Alexander S. Wein, Spectral methods from tensor networks, in: *Proceedings of the 51st Annual ACM Symposium on Theory of Computing*, ACM, 2019, pp. 926–937.
- [7] Nir Sharon, Joe Kileel, Yuehaw Khoo, Boris Landa, Amit Singer, Method of moments for 3D single particle ab initio modeling with non-uniform distribution of viewing angles, *Inverse Problems* 36 (4) (2020) 044003.
- [8] Amit Singer, Mathematics for cryo-electron microscopy, in: *Proceedings of the International Congress of Mathematicians*, Vol. 4, 2018, pp. 4013–4032.

- [9] Joachim Frank, Three-Dimensional Electron Microscopy of Macromolecular Assemblies: Visualization of Biological Molecules in their Native State, Oxford University Press, 2006.
- [10] Ali Punjani, David Fleet, Advances in modelling continuous heterogeneity from single particle cryo-EM data, *Found. Crystallogr.* 77 (2021) A235.
- [11] Ellen D Zhong, Tristan Bepler, Bonnie Berger, Joseph H Davis, CryoDRGN: reconstruction of heterogeneous cryo-EM structures using neural networks, *Nature Methods* 18 (2) (2021) 176–185.
- [12] Ellen D Zhong, Adam Lerer, Joseph H Davis, Bonnie Berger, CryoDRGN2: Ab initio neural reconstruction of 3D protein structures from real cryo-EM images, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 4066–4075.
- [13] A. Jiménez-Moreno, D. Střelák, J Filipovič, JM Carazo, CÓS Sorzano, DeepAlign, a 3D alignment method based on regionalized deep learning for cryo-EM, *J. Struct. Biol.* 213 (2) (2021) 107712.
- [14] Tristan Bepler, Kotaro Kelley, Alex J Noble, Bonnie Berger, Topaz-denoise: general deep denoising models for cryoEM and cryoET, *Nature Commun.* 11 (1) (2020) 1–12.
- [15] Dari Kimanius, Liyi Dong, Grigory Sharov, Takanori Nakane, Sjors HW Scheres, New tools for automated cryo-EM single-particle analysis in RELION-4.0, *Biochem. J.* 478 (24) (2021) 4169–4185.
- [16] Axel Levy, Frédéric Poitevin, Julien Martel, Youssef Nashed, Ariana Peck, Nina Miolane, Daniel Ratner, Mike Dunne, Gordon Wetzstein, CryoAI: Amortized inference of poses for ab initio reconstruction of 3D molecular volumes from real cryo-EM images, 2022, arXiv preprint [arXiv:2203.08138](https://arxiv.org/abs/2203.08138).
- [17] Joakim Andén, Amit Singer, Structural variability from noisy tomographic projections, *SIAM J. Imaging Sci.* 11 (2) (2018) 1441–1492.
- [18] Yoel Shkolnisky, Amit Singer, Viewing direction estimation in cryo-EM using synchronization, *SIAM J. Imaging Sci.* 5 (3) (2012) 1088–1110.
- [19] James O. Berger, Brunero Liseo, Robert L. Wolpert, Integrated likelihood methods for eliminating nuisance parameters, *Statist. Sci.* 14 (1) (1999) 1–28.
- [20] Tamir Bendory, Nicolas Boumal, Chao Ma, Zhizhen Zhao, Amit Singer, Bispectrum inversion with application to multireference alignment, *IEEE Trans. Signal Process.* 66 (4) (2017) 1037–1050.
- [21] Emmanuel Abbe, João M. Pereira, Amit Singer, Estimation in the group action channel, in: *2018 IEEE International Symposium on Information Theory, IEEE*, 2018, pp. 561–565.
- [22] Amelia Perry, Jonathan Weed, Afonso S Bandeira, Philippe Rigollet, Amit Singer, The sample complexity of multireference alignment, *SIAM J. Math. Data Sci.* 1 (3) (2019) 497–517.
- [23] Ayelet Heimowitz, Nir Sharon, Amit Singer, Centering noisy images with application to cryo-EM, *SIAM J. Imaging Sci.* 14 (2) (2021) 689–716.
- [24] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimeshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, Soumith Chintala, Pytorch: An imperative style, high-performance deep learning library, in: *Advances in Neural Information Processing Systems*, vol. 32, Curran Associates, Inc., 2019, pp. 8024–8035, URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [25] Robert S. Womersley, Efficient spherical designs with good geometric properties, in: *Contemporary Computational Mathematics-A Celebration of the 80th Birthday of Ian Sloan*, Springer, 2018, pp. 1243–1285.
- [26] Yael Harpaz, Yoel Shkolnisky, Three-dimensional alignment of density maps in cryo-electron microscopy, *Biol. Imaging* 3 (2023) e8.
- [27] Eric F Pettersen, Thomas D Goddard, Conrad C Huang, Gregory S Couch, Daniel M Greenblatt, Elaine C Meng, Thomas E Ferrin, UCSF chimera—A visualization system for exploratory research and analysis, *J. Comput. Chem.* 25 (13) (2004) 1605–1612.
- [28] Werner Kühlbrandt, The resolution revolution, *Science* 343 (6178) (2014) 1443–1444.
- [29] Jacques Dubochet, Joachim Frank, Richard Henderson, The nobel prize in chemistry 2017, *Nobel Media AB* (2017).
- [30] Wei Cai, Xiaoguang Li, Lizuo Liu, A phase shift deep neural network for high frequency approximation and wave problems, *SIAM J. Sci. Comput.* 42 (5) (2020) A3285–A3312.