Davis Math Lab Project Report, Spring 2026
Stabilizing Mean and Covariance Estimation in Fourier-Bessel Cryo-EM PCA

Ajay Penmetsa
Graduate Mentors: Anna Luo and Caleb Farrow
Faculty Mentor: Yunpeng Shi

June 15, 2026

Contents

1 Introduction

1.1 Background

Single-particle cryo-electron microscopy (cryo-EM) produces large collections of very noisy two-dimensional particle images. A central computational task is to identify meaningful structural variation inside these datasets. Principal component analysis (PCA) is useful for this purpose because the leading eigenvectors of the covariance matrix describe dominant modes of variation, and this covariance information can also be used for denoising through covariance Wiener filtering.

We consider the image formation model

$$g_i = h_i * f_i + \varepsilon_i, \qquad i = 1, \ldots, N, \tag{1}$$

where f_i is the clean image, h_i is the point-spread function or contrast transfer function (CTF), and ε_i is noise. In the Fourier-Bessel basis, radial convolution becomes diagonal, so the model becomes

$$G_i = H_i \odot F_i + E_i, \tag{2}$$

where G_i, F_i, E_i are coefficient vectors, H_i stores the CTF weights, and $\odot$ denotes entrywise multiplication. This representation is useful because the clean covariance is block diagonal by angular frequency, allowing the covariance blocks to be estimated separately.

The fast covariance estimator of Marshall, Mickelin, Shi, and Singer [1] uses this Fourier-Bessel structure to avoid expensive conjugate-gradient covariance estimation and obtains complexity essentially independent of the number of defocus groups. The estimator solves Hadamard moment equations for the covariance and then applies positive semidefinite or eigenvalue shrinkage. However, the original approach uses a global noise/SNR scale and does not explicitly estimate how signal strength changes with frequency.

1.2 Current Goal

The goal of this project is to improve the stability and accuracy of mean and covariance estimation in the Fourier-Bessel coefficient model. The first direction was regularizing the mean estimator when CTF weights are small. The second and more recent direction is to build a frequency-adaptive covariance estimator that estimates diagonal signal power, measures frequency-dependent uncertainty, and applies SNR-adaptive shrinkage directly inside the Hadamard covariance equations.

2 Progress

2.1 Regularized Mean Estimation

The unregularized mean estimator solves

$$\widetilde{\mu} = \arg\min_{\mu} \sum_{i=1}^{N} \|G_i - H_i \odot \mu\|^2. \tag{3}$$

This gives the closed-form coefficientwise estimator

$$\widetilde{\mu}_j = \frac{\sum_{i=1}^{N} G_i(j) H_i(j)}{\sum_{i=1}^{N} H_i(j)^2}. \tag{4}$$

This expression can become unstable when $H_i(j)$ is small across many images, since the denominator may be close to zero. To reduce this instability, we introduced a ridge-type regularization term:

$$\widetilde{\mu} = \arg\min_{\mu} \sum_{i=1}^{N} \|G_i - H_i \odot \mu\|^2 + \lambda \|\mu\|^2. \tag{5}$$

The resulting estimator is

$$\widetilde{\mu}_j = \frac{\sum_{i=1}^{N} G_i(j) H_i(j)}{\sum_{i=1}^{N} H_i(j)^2 + \lambda}. \tag{6}$$

An initial experiment selected λ using a block-correlation heuristic. For each covariance block k, the coefficient matrix was split into two disjoint image subsets, giving matrices $a_k^{(1)}$ and $a_k^{(2)}$. The reliability of the block was measured by

$$r_k = \frac{\langle a_k^{(1)}, a_k^{(2)} \rangle}{\|a_k^{(1)}\| \|a_k^{(2)}\|}. \tag{7}$$

The preliminary choice

$$\lambda = C(1 - r_0) \sum_{i=1}^{N} H_i(j) \tag{8}$$

uses the lowest-frequency block $k = 0$, where $1 - r_0$ increases when the signal appears less reliable. For SNR = 1, grid search found a value of C that slightly improved mean estimation error over the unregularized estimator. This suggested that regularization can help, but also showed that a more principled frequency-dependent choice of regularization is needed.

2.2 SNR-Adaptive Weighted Hadamard Shrinkage Covariance Estimator

After estimating the mean, define centered coefficients

$$R_i = G_i - H_i \odot \widetilde{\mu}. \tag{9}$$

The classical Hadamard covariance estimator is based on the moment equation

$$\mathbb{E}\big[R_{ia}R_{ib} - \sigma^2 \mathbf{1}_{a=b} \mid H_i\big] = H_{ia}H_{ib}C_{ab}. \tag{10}$$

Thus the old unweighted estimator is

$$\widehat{C}^{\mathrm{old}}_{ab} = \frac{\sum_i H_{ia}H_{ib}\left(R_{ia}R_{ib} - \sigma^2 \mathbf{1}_{a=b}\right)}{\sum_i (H_{ia}H_{ib})^2}. \tag{11}$$

This estimator uses CTF observability through the denominator, but it treats moment residuals as equally reliable and does not adapt to frequency-dependent signal power.

The main statistical correction is that white coefficient noise does not imply equal SNR. If the clean covariance is C, then the deconvolved diagonal SNR at coefficient a is

$$\mathrm{SNR}_a = \frac{C_{aa}}{\sigma^2}, \tag{12}$$

while the observed SNR for image i is

$$\mathrm{SNR}^{\mathrm{obs}}_{ia} = \frac{H_{ia}^2 C_{aa}}{\sigma^2}. \tag{13}$$

Because clean cryo-EM signal power typically decays with frequency, high-frequency covariance entries can have very small true signal but very large deconvolution variance.

We estimate diagonal signal power using

$$\mathbb{E}\big[R_{ia}^2 - \sigma^2 \mid H_i\big] = H_{ia}^2 C_{aa}. \tag{14}$$

Let

$$y_{ia} = R_{ia}^2 - \sigma^2, \qquad x_{ia} = H_{ia}^2. \tag{15}$$

Then a weighted diagonal estimate is

$$\widehat{C}_{aa} = \frac{\sum_i w_{ia}x_{ia}y_{ia}}{\sum_i w_{ia}x_{ia}^2}. \tag{16}$$

Under a Gaussian working model,

$$\mathrm{Var}(y_{ia} \mid H_i) = 2(H_{ia}^2 C_{aa} + \sigma^2)^2, \tag{17}$$

which gives natural inverse-variance weights. The raw diagonal estimates are then smoothed into a conservative decreasing frequency envelope s_a.

For full covariance entries, write

$$z_{iab} = R_{ia}R_{ib} - \sigma^2 \mathbf{1}_{a=b}, \qquad x_{iab} = H_{ia}H_{ib}. \tag{18}$$

If

$$\Sigma_i = \mathrm{diag}(H_i) C \,\mathrm{diag}(H_i) + \sigma^2 I, \tag{19}$$

then the real Gaussian fourth-moment identity gives

$$\mathrm{Var}(R_{ia}R_{ib} \mid H_i) = \Sigma_{i,aa}\Sigma_{i,bb} + \Sigma_{i,ab}^2. \tag{20}$$

This motivates inverse-variance weights for the Hadamard covariance equations.

The recommended estimator solves the regularized MAP/GMM problem

$$\widehat{C} = \arg\min_{C \succeq 0} \frac{1}{2} \sum_{i, a \le b} w_{iab}\left(z_{iab} - x_{iab}C_{ab}\right)^2 + \frac{1}{2}\sum_{a \le b} \frac{(C_{ab} - T_{ab})^2}{\tau_{ab}^2}. \tag{21}$$

The default target is

$$T_{aa} = s_a, \qquad T_{ab} = 0 \quad (a \neq b). \tag{22}$$

The prior scale τ_{ab}^2 is chosen from the estimated frequency-power envelope. Since positive semidefiniteness implies

$$|C_{ab}|^2 \leq C_{aa} C_{bb}, \tag{23}$$

small diagonal power at either frequency implies that the off-diagonal covariance should also be small. Thus high-frequency, low-SNR entries receive stronger shrinkage.

Without the PSD constraint, the entrywise solution is

$$\widehat{C}_{ab} = \frac{\sum_i w_{iab} x_{iab} z_{iab} + \tau_{ab}^{-2} T_{ab}}{\sum_i w_{iab} x_{iab}^2 + \tau_{ab}^{-2}}. \tag{24}$$

Equivalently,

$$\widehat{C}_{ab} = T_{ab} + \alpha_{ab} \left(\widehat{C}_{ab}^{\mathrm{WLS}} - T_{ab} \right), \qquad \alpha_{ab} = \frac{A_{ab}}{A_{ab} + \tau_{ab}^{-2}}, \tag{25}$$

where $A_{ab} = \sum_i w_{iab} x_{iab}^2$. Therefore shrinkage is not a post-processing step, as it is the closed-form solution of a regularized version of the original Hadamard CTF equations.

Synthetic experiments with decaying frequency SNR confirmed the need for shrinkage. In one comparison, the global covariance relative error of the old unweighted estimator was about 40.27, while the SNR-adaptive shrinkage estimator reduced it to about 0.157, close to the oracle entrywise shrinkage benchmark of about 0.132. The estimated SNR and the learned shrinkage factors both decreased strongly with frequency, matching the intended behavior.

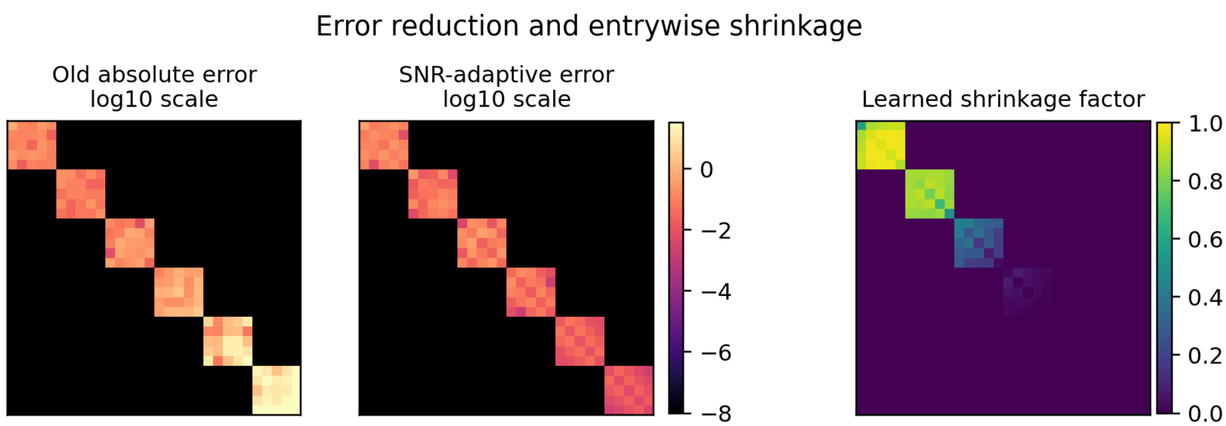

Figure 1: Entrywise error reduction and learned shrinkage in the decaying-SNR synthetic experiment. The first two panels show $\log_{10}(|\widehat{C}_{ab} - C_{ab}| + 10^{-8})$, so values near 0 indicate order-one absolute error, while values near -8 indicate nearly zero error. The right panel shows the learned shrinkage factor $\alpha_{ab} \in [0, 1]$, where $\alpha_{ab} \approx 1$ means the data are trusted and $\alpha_{ab} \approx 0$ means the estimate is strongly shrunk toward the target.

3 Future Plans

We have several potential future directions. First, the old mean regularization parameter should be chosen from a frequency-dependent uncertainty model rather than a scalar grid search. Second, the SNR-adaptive covariance estimator should be tested more extensively on experimental datasets, including its effect on PCA eigenimages and covariance Wiener filtering. Third, split-half and repeated-subset methods can be used to calibrate entrywise uncertainty:

$$\widetilde{\gamma}_{ij} = \widehat{C}_{ij}^{(1)} - \widehat{C}_{ij}^{(2)}. \tag{26}$$

Balanced splits by CTF or defocus group are important so that differences between subsets reflect estimator uncertainty rather than systematic sampling bias.

Finally, evaluation should be frequency-aware. A single global covariance error can hide the fact that high-frequency relative error may be large simply because the true high-frequency covariance norm is tiny. Future tests should observe signal norm, absolute RMSE, relative error, shrinkage factor, uncertainty calibration, held-out likelihood, and downstream denoising or PCA quality by frequency band.

4 References

1. N. F. Marshall, O. Mickelin, Y. Shi, and A. Singer. Fast principal component analysis for cryo-electron microscopy images. *Biological Imaging*, 3:e2, 2023.

2. T. Bhamre, T. Zhang, and A. Singer. Denoising and covariance estimation of single particle cryo-EM images. *Journal of Structural Biology*, 195(1):72–81, 2016.

3. D. L. Donoho, M. Gavish, and I. M. Johnstone. Optimal shrinkage of eigenvalues in the spiked covariance model. *Annals of Statistics*, 46(4):1742–1778, 2018.