Reconstructing without Angles in Unknown-View Tomography

Zhizhen Jane Zhao

Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign

ICERM Mathematics of 3D Reconstruction

July 15, 2026

Introduction

- Scientific imaging is used to characterize physical, material, chemical, and biological processes.

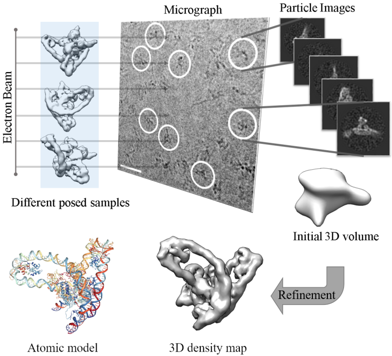

Cryo-electron microscopy single particle analysis

Introduction

- Scientific imaging is used to characterize physical, material, chemical, and biological processes.

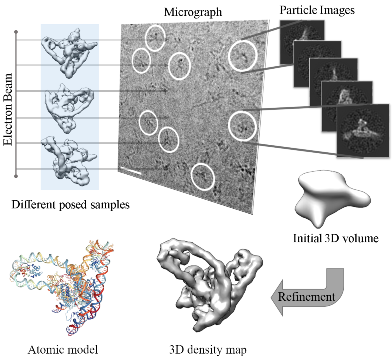

Cryo-electron microscopy single particle analysis

Imaging data challenges include

- Incomplete data
- measurement uncertainty (e.g., noise, unknown latent parameters)

Introduction

- Scientific imaging is used to characterize physical, material, chemical, and biological processes.

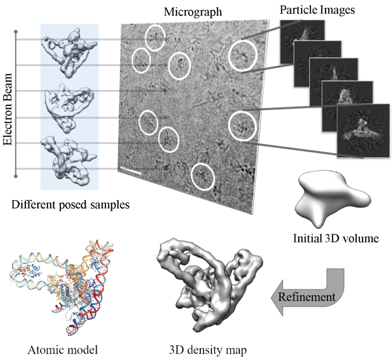

Cryo-electron microscopy single particle analysis

Imaging data challenges include

- Incomplete data
- measurement uncertainty (e.g., noise, unknown latent parameters)

Here we focus on **tomography with unknown view angles.**

Introduction

- Scientific imaging is used to characterize physical, material, chemical, and biological processes.

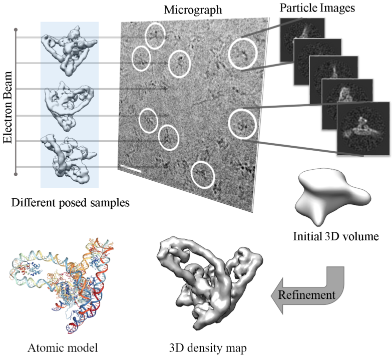

Cryo-electron microscopy single particle analysis

Imaging data challenges include

- Incomplete data
- measurement uncertainty (e.g., noise, unknown latent parameters)

Here we focus on **tomography with unknown view angles.**

Part I. Orthogonal matrix retrieval with spatial consensus

Joint work with Shuai Huang, Mona Zehni, and Ivan Dokmanic

Image formation model

- A simplified image formation model for a 3D electron density map ρ and $\omega_i \in \mathrm{SO}(3)$:

$$\omega_i \circ \rho$$

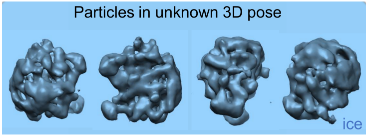

Image formation model

- A simplified image formation model for a 3D electron density map ρ and $\omega_i \in \mathrm{SO}(3)$:

$$I_i = \quad P(\omega_i \circ \rho)$$

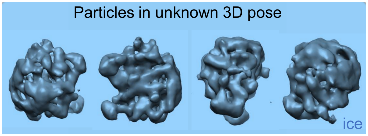

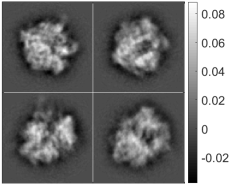

projections

Image formation model

- A simplified image formation model for a 3D electron density map ρ and $\omega_i \in \mathrm{SO}(3)$:

$$I_i = H_i \star P(\omega_i \circ \rho)$$

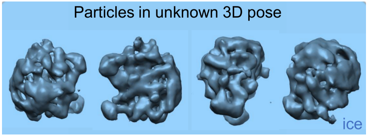

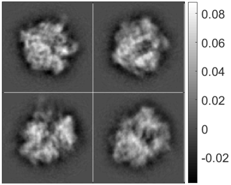

projections

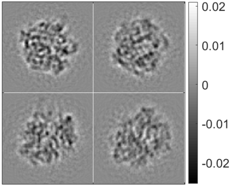

w. CTFs

Image formation model

- A simplified image formation model for a 3D electron density map ρ and $\omega_i \in \mathrm{SO}(3)$:

$$I_i = {\color{green}H_i} \star {\color{red}P(}{\color{blue}\omega_i \circ \rho}{\color{red})} + \text{noise}$$

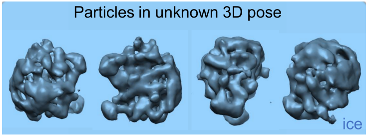

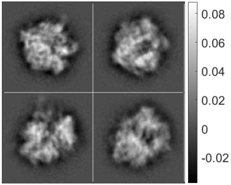

projections

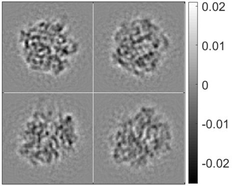

w. CTFs

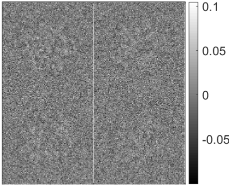

SNR $= 0.05$

Noisy images

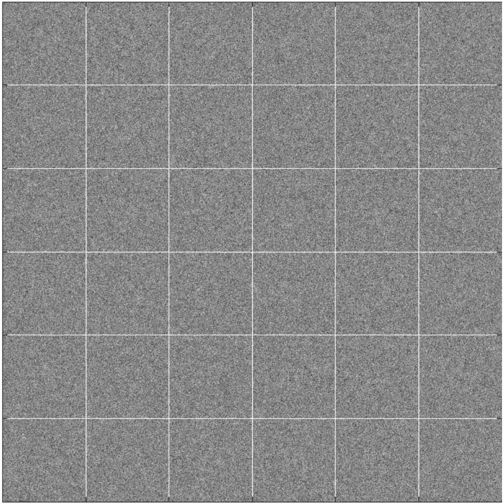

Typically contains over a million raw particle images

https://www.ebi.ac.uk/pdbe/emdb/empiar/entry/10005/

Traditional estimation approach

Alternate between estimating the particle orientation and reconstructing the 3D density

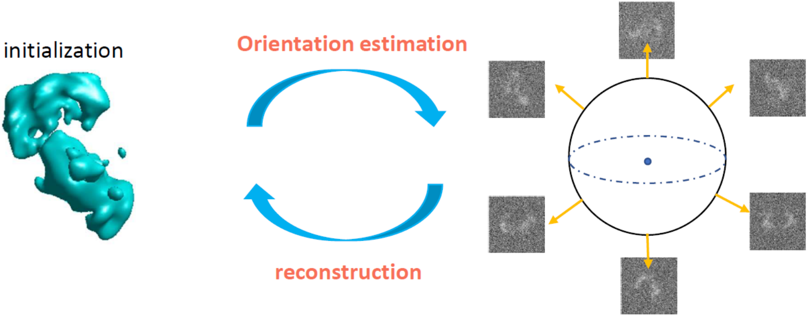

- Multiple passes over the whole data set

Method of moments

Reconstruct the 3D density via the method of moments

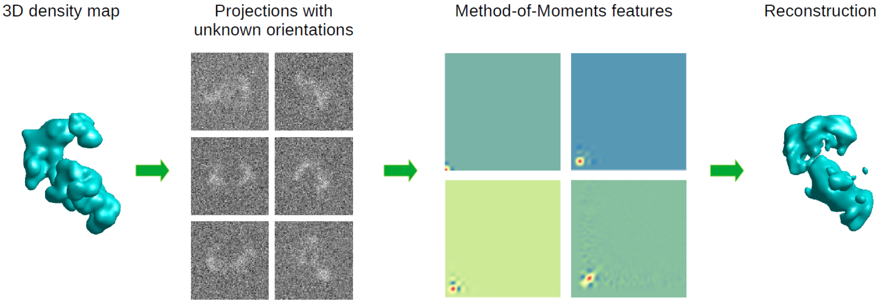

- Denoised moments can be obtained

Method of moments

Reconstruct the 3D density via the method of moments

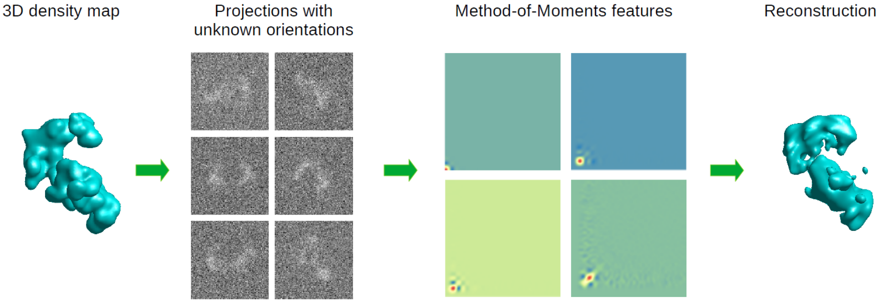

- Denoised moments can be obtained
- Only go through the images once to extract features
- Bypass the orientation estimation

Method of moments

Reconstruct the 3D density via the method of moments

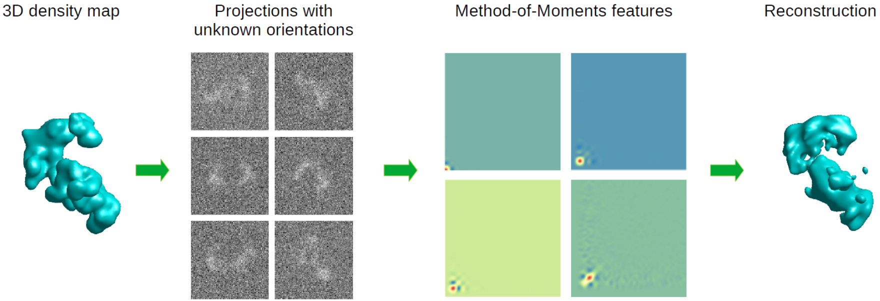

- Denoised moments can be obtained
- Only go through the images once to extract features
- Bypass the orientation estimation
- The sample complexity scales as σ^{2m} for the estimation of m^{th} order moments

Method of moments

Reconstruct the 3D density via the method of moments

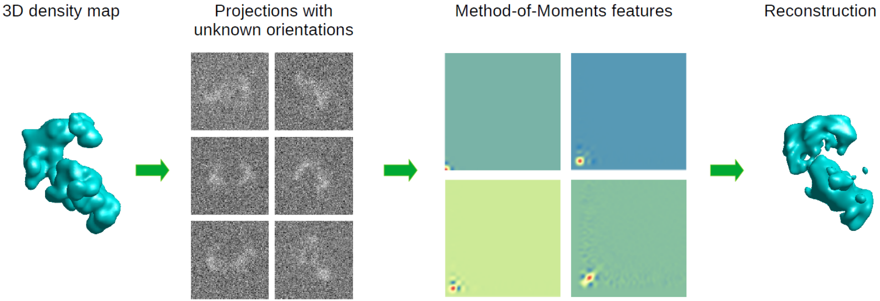

- Denoised moments can be obtained
- Only go through the images once to extract features
- Bypass the orientation estimation
- The sample complexity scales as σ^{2m} for the estimation of m^{th} order moments
- Some assumptions: centered images, distribution of the orientations is uniform over SO(3)

Projection-slice theorem

- Let $\boldsymbol{R}_i$ be the rotation matrix corresponding to ω_i, for clean measurement data, we have

$$\widehat{I}_i(k_x, k_y) = \widehat{\rho}\left(\boldsymbol{R}_i^T [k_x \ k_y \ 0]^T\right)$$

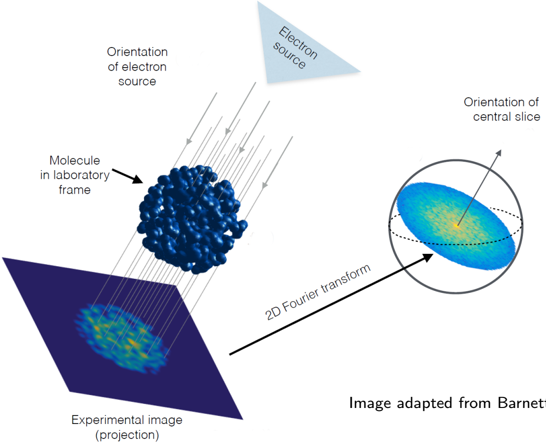

Image adapted from Barnett et al. SIIMS, 2017

Projection-slice theorem

- Let $\boldsymbol{R}_i$ be the rotation matrix corresponding to ω_i, for clean measurement data, we have

$$\widehat{I}_i(k_x, k_y) = \widehat{\rho}\left(\boldsymbol{R}_i^T [k_x \ k_y \ 0]^T\right)$$

- Spherical harmonic expansion:

$$\widehat{\rho}(\boldsymbol{k}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} A_{lm}(k) \cdot Y_{lm}(\theta_{\boldsymbol{k}}, \varphi_{\boldsymbol{k}}),$$

$A_{lm}(k)$ is the real spherical harmonic expansion coefficients.

Projection-slice theorem

- Let $\boldsymbol{R}_i$ be the rotation matrix corresponding to ω_i, for clean measurement data, we have

$$\widehat{I}_i(k_x, k_y) = \widehat{\rho}\left(\boldsymbol{R}_i^T [k_x \ k_y \ 0]^T\right)$$

- Spherical harmonic expansion:

$$\widehat{\rho}(\boldsymbol{k}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} A_{lm}(k) \cdot Y_{lm}(\theta_{\boldsymbol{k}}, \varphi_{\boldsymbol{k}}),$$

  $A_{lm}(k)$ is the real spherical harmonic expansion coefficients.

- In terms of the spherical expansion coefficients,

$$\widehat{I}_i(k_x, k_y) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} \sum_{m'=-l}^{l} A_{lm}(k) \cdot Y_{lm'}\left(\frac{\pi}{2}, \varphi_{\boldsymbol{k}}\right) \cdot D^l_{mm'}(\omega_i),$$

  where D^l is the Wigner D-matrix.

Kam's autocorrelation analysis

- Zvi Kam showed in 1980 that the autocorrelation function of the 3D molecule over the rotation group SO(3) can be estimated from 2D projection images whose viewing directions are uniformly distributed over the sphere.

- Zvi Kam showed in 1980 that the autocorrelation function of the 3D molecule over the rotation group SO(3) can be estimated from 2D projection images whose viewing directions are uniformly distributed over the sphere.

$$C_N(k_1, k_2, \psi) = \frac{1}{N}\sum_{n=1}^{N} \frac{1}{2\pi}\int_0^{2\pi} \widehat{I}_n(k_1, \varphi_1) \cdot \widehat{I}_n^*(k_2, \varphi_1 + \psi)\ d\varphi_1$$

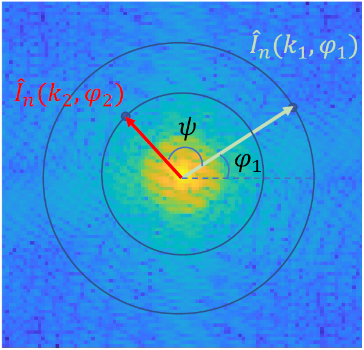

Kam's autocorrelation analysis

- Zvi Kam showed in 1980 that the autocorrelation function of the 3D molecule over the rotation group SO(3) can be estimated from 2D projection images whose viewing directions are uniformly distributed over the sphere.

$$C_N\left(k_1, k_2, \psi\right) = \frac{1}{N}\sum_{n=1}^{N}\frac{1}{2\pi}\int_0^{2\pi}\widehat{I}_n\left(k_1, \varphi_1\right)\cdot\widehat{I}_n^*\left(k_2, \varphi_1+\psi\right)\, d\varphi_1$$

- For clean images,

$$C_N\left(k_1, k_2, \psi\right) \xrightarrow{N\to\infty} \underbrace{\frac{1}{4\pi}\sum_{l=0}^{\infty}P_l(\cos\psi)\sum_{m=-l}^{l}A_{lm}(k_1)\cdot A_{lm}^*(k_2)}_{=:\, C(k_1, k_2, \psi)},$$

where $P_l(\cdot)$ is the Legendre polynomial of degree l, and $C(k_1, k_2, \psi)$ is the asymptotic noiseless autocorrelation function of the Fourier transform $\widehat{\rho}(\boldsymbol{k})$ in the frequency domain over the rotation group $\mathrm{SO}(3)$.

Kam's autocorrelation analysis

- Kam further computed the contribution of the subspace of all spherical harmonics with degree l, using orthogonality of Legendre polynomials,

$$C_l(k_1, k_2) = 2\pi(2l+1)\int_0^\pi C(k_1, k_2, \psi) \cdot P_l(\cos\psi) \cdot \sin\psi \, d\psi$$

$$= \sum_{m=-l}^{l} A_{lm}(k_1) \cdot A_{lm}^*(k_2), \quad l = 0, 1, \ldots, L$$

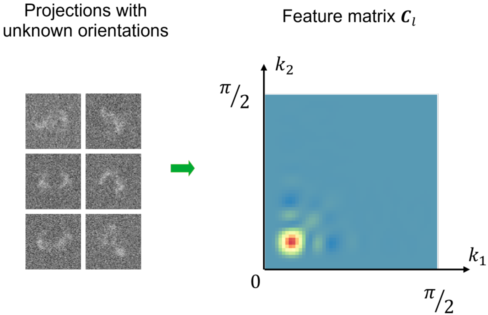

Kam's autocorrelation analysis

- Kam further computed the contribution of the subspace of all spherical harmonics with degree l, using orthogonality of Legendre polynomials,

$$C_l(k_1, k_2) = 2\pi(2l+1)\int_0^\pi C(k_1, k_2, \psi) \cdot P_l(\cos\psi) \cdot \sin\psi \, d\psi$$
$$= \sum_{m=-l}^{l} A_{lm}(k_1) \cdot A^*_{lm}(k_2)\ , \quad l = 0, 1, \ldots, L$$

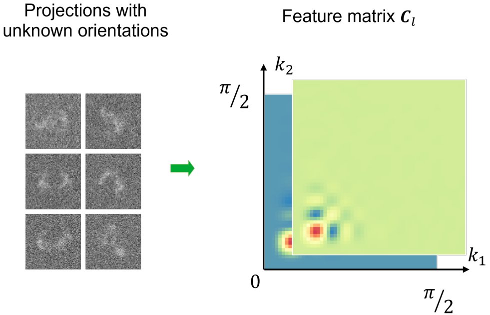

Kam's autocorrelation analysis

- Kam further computed the contribution of the subspace of all spherical harmonics with degree l, using orthogonality of Legendre polynomials,

$$C_l\,(k_1, k_2) = 2\pi(2l+1)\int_0^{\pi} C\,(k_1, k_2, \psi)\cdot P_l(\cos\psi)\cdot \sin\psi\; d\psi$$

$$= \sum_{m=-l}^{l} A_{lm}\,(k_1)\cdot A^*_{lm}\,(k_2)\ , \quad l = 0, 1, \ldots, L$$

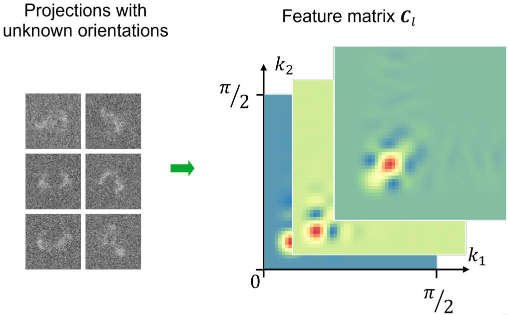

Estimated autorcorrelation features

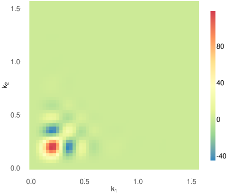

Noiseless C_1

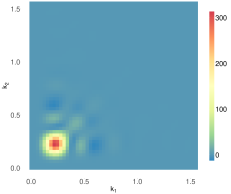

Noiseless C_5

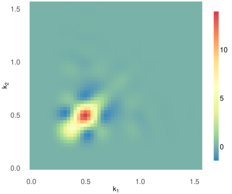

Noiseless C_{10}

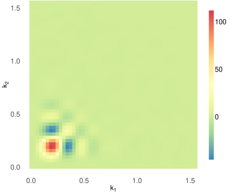

RRMSE of $\widetilde{C}_1$: 0.0821

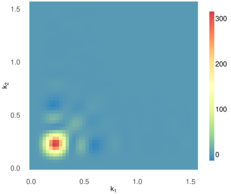

RRMSE of $\widetilde{C}_5$: 0.0197

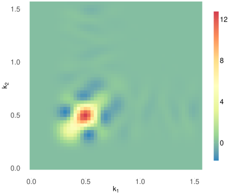

RRMSE of $\widetilde{C}_{10}$: 0.1034

Fourier autocorrelation feature $\widetilde{C}_l$ is extracted from $N = 10^4$ projection images with SNR=0.1.

Zhao et al., Fast steerable principal component analysis, IEEE TCI, 2016.
Bhamre et al., Denoising and covariance estimation of single particle cryo-EM Images, JSB, 2016

Kam's autocorrelation analysis

- If we can determine $A_{lm}(k)$, then

$$\widehat{\rho}(\boldsymbol{k}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} A_{lm}(k) \cdot Y_{lm}(\theta_{\boldsymbol{k}}, \varphi_{\boldsymbol{k}}) \,,$$

Kam's autocorrelation analysis

- If we can determine $A_{lm}(k)$, then

$$\widehat{\rho}(\boldsymbol{k}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} A_{lm}(k) \cdot Y_{lm}(\theta_{\boldsymbol{k}}, \varphi_{\boldsymbol{k}}),$$

- Truncated reconstruction:

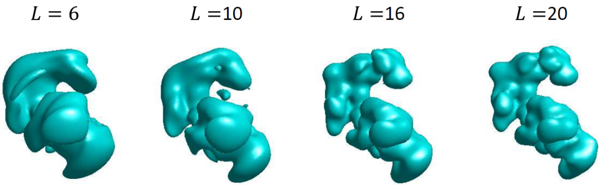

Orthogonal matrix retrieval

- We have for any orthogonal $\boldsymbol{O}_l$ of size $(2l+1) \times (2l+1)$ that

$$\boldsymbol{C}_l = \boldsymbol{A}_l \boldsymbol{O}_l^T (\boldsymbol{A}_l \boldsymbol{O}_l^T)^* = \boldsymbol{F}_l \boldsymbol{F}_l^* \, .$$

[*] T. Bhamre, T. Zhang, and A. Singer, Orthogonal matrix retrieval in cryo-electron microscopy, ISBI, 2015

[+] E. Levin, T. Bendory, N. Boumal, J. Kileel, and A. Singer, 3D ab initio modeling in cryo-EM by autocorrelation analysis, ISBI, 2018

Orthogonal matrix retrieval

- We have for any orthogonal $\boldsymbol{O}_l$ of size $(2l+1) \times (2l+1)$ that

$$\boldsymbol{C}_l = \boldsymbol{A}_l \boldsymbol{O}_l^T (\boldsymbol{A}_l \boldsymbol{O}_l^T)^* = \boldsymbol{F}_l \boldsymbol{F}_l^* \,.$$

- The Cholesky decomposition of $\boldsymbol{C}_l$ yields $\boldsymbol{F}_l$. So to recover the volume, we need to recover orthogonal matrix $\boldsymbol{O}_l$ using additional information.

---

* T. Bhamre, T. Zhang, and A. Singer, Orthogonal matrix retrieval in cryo-electron microscopy, ISBI, 2015

+E. Levin, T. Bendory, N. Boumal, J. Kileel, and A. Singer, 3D ab initio modeling in cryo-EM by autocorrelation analysis, ISBI, 2018

Orthogonal matrix retrieval

- We have for any orthogonal $\boldsymbol{O}_l$ of size $(2l+1) \times (2l+1)$ that

$$\boldsymbol{C}_l = \boldsymbol{A}_l \boldsymbol{O}_l^T (\boldsymbol{A}_l \boldsymbol{O}_l^T)^* = \boldsymbol{F}_l \boldsymbol{F}_l^* \,.$$

- The Cholesky decomposition of $\boldsymbol{C}_l$ yields $\boldsymbol{F}_l$. So to recover the volume, we need to recover orthogonal matrix $\boldsymbol{O}_l$ using additional information.
- Bhamre et al* proposed to use a known homologous volume to help determine the $\boldsymbol{O}_l$.
- Levin et al+ proposed to use two denoised projection images taken from different views to determine $\boldsymbol{O}_l$. (OMR-PM)

---

* T. Bhamre, T. Zhang, and A. Singer, Orthogonal matrix retrieval in cryo-electron microscopy, ISBI, 2015

+ E. Levin, T. Bendory, N. Boumal, J. Kileel, and A. Singer, 3D ab initio modeling in cryo-EM by autocorrelation analysis, ISBI, 2018

Orthogonal matrix retrieval

- We have for any orthogonal $\boldsymbol{O}_l$ of size $(2l+1) \times (2l+1)$ that

$$\boldsymbol{C}_l = \boldsymbol{A}_l \boldsymbol{O}_l^T (\boldsymbol{A}_l \boldsymbol{O}_l^T)^* = \boldsymbol{F}_l \boldsymbol{F}_l^* \,.$$

- The Cholesky decomposition of $\boldsymbol{C}_l$ yields $\boldsymbol{F}_l$. So to recover the volume, we need to recover orthogonal matrix $\boldsymbol{O}_l$ using additional information.
- Bhamre et al* proposed to use a known homologous volume to help determine the $\boldsymbol{O}_l$.
- Levin et al+ proposed to use two denoised projection images taken from different views to determine $\boldsymbol{O}_l$. (OMR-PM)
- For a pair of images I_1 and I_2, we have $\rho_2 = \omega \circ \rho_1$, with $\omega \in \mathrm{SO}(3)$. Optimize over $\{O_l\}_{l=1}^L$ and ω.

---

* T. Bhamre, T. Zhang, and A. Singer, Orthogonal matrix retrieval in cryo-electron microscopy, ISBI, 2015

+ E. Levin, T. Bendory, N. Boumal, J. Kileel, and A. Singer, 3D ab initio modeling in cryo-EM by autocorrelation analysis, ISBI, 2018

Challenges in OMR

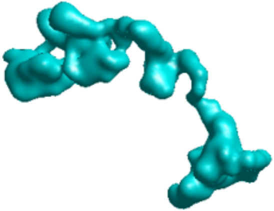

Ground truth

Challenges in OMR

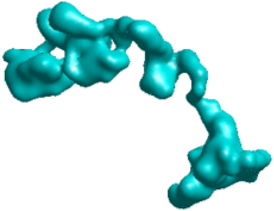

Ground truth

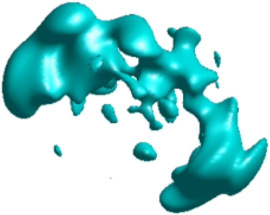

Clean reconstruction

Challenges in OMR

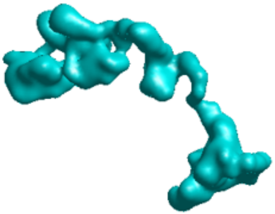

Ground truth

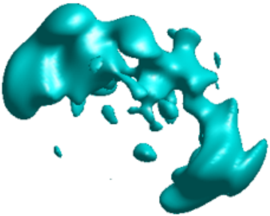

Clean reconstruction

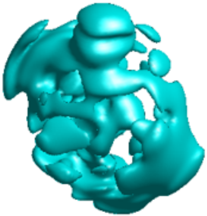

Noisy reconstruction

Challenges in OMR

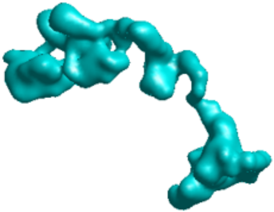

Ground truth

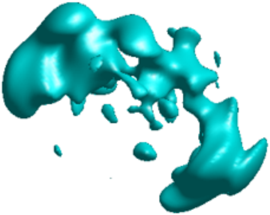

Clean reconstruction

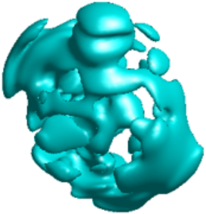

Noisy reconstruction

- Previous approach focuses mainly on the recovery of O_l

Challenges in OMR

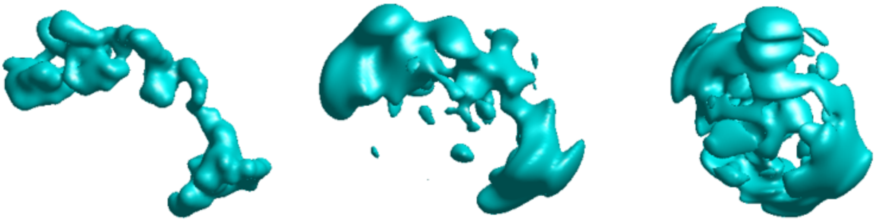

Ground truth   Clean reconstruction   Noisy reconstruction

- Previous approach focuses mainly on the recovery of O_l
- Not robust enough in the noisy case (SNR$= 0.1$)

Challenges with OMR

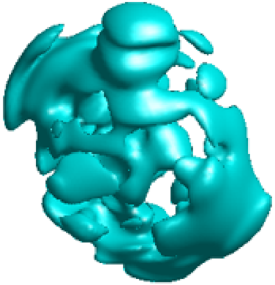

Noisy reconstruction

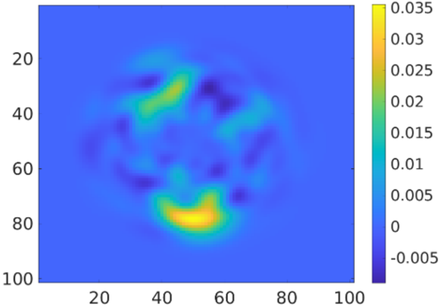

A 2D slice

Challenges with OMR

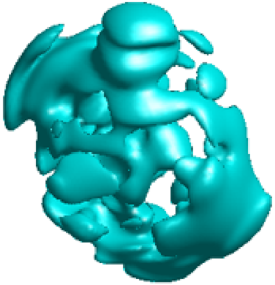

Noisy reconstruction

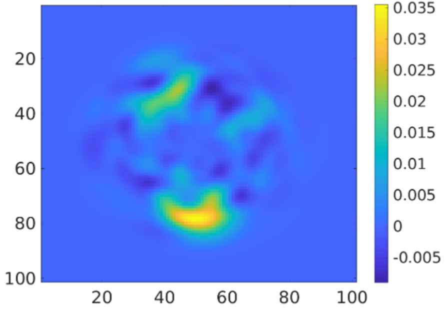

A 2D slice

- How to enforce that the reconstructed density map is non-negative

Challenges with OMR

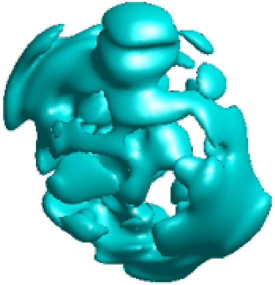

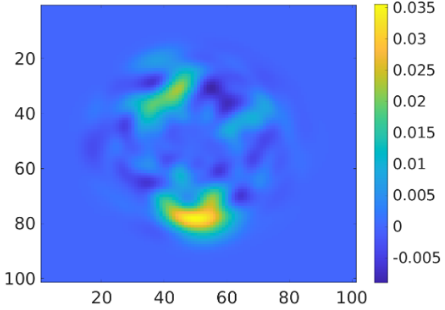

Noisy reconstruction

A 2D slice

- How to enforce that the reconstructed density map is non-negative
- How to improve the reconstruction quality from the noisy measurements

Spatial constraints on the density map

- Non-negativity and total mass constraint

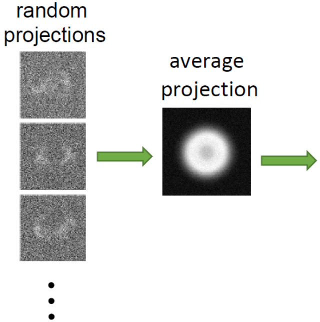

total mass

$$W_\rho = \|\boldsymbol{\rho}\|_1$$

$$0 \leq \rho_d < W_\rho, \qquad d = 1, \cdots, D$$

$$\|\boldsymbol{\rho}\|_1 = \sum_{d=1}^{D} |\rho_d| = W_\rho$$

Spatial constraints on the density map

- Non-negativity and total mass constraint

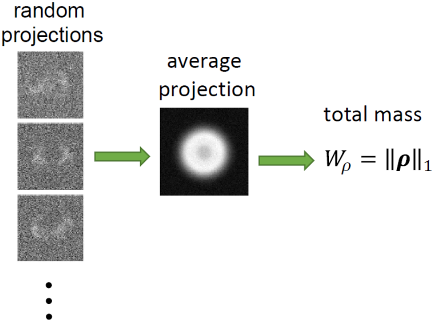

$$0 \le \rho_d < W_\rho, \qquad d = 1, \cdots, D$$

$$\|\boldsymbol{\rho}\|_1 = \sum_{d=1}^{D} |\rho_d| = W_\rho$$

- Reference projection (randomly chosen): $P_\rho(x, y) = I(x, y)$

Spatial autocorrelation features

- The density map has the following real spherical harmonic expansion,

$$\rho(\boldsymbol{r}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} B_{lm}(r) \cdot Y_{lm}(\theta_{\boldsymbol{r}}, \varphi_{\boldsymbol{r}}) \, .$$

Spatial autocorrelation features

- The density map has the following real spherical harmonic expansion,

$$\rho(\boldsymbol{r}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} B_{lm}(r) \cdot Y_{lm}(\theta_{\boldsymbol{r}}, \varphi_{\boldsymbol{r}}) \,.$$

- Spatial autocorrelation features

$$E_l(r_1, r_2) = \sum_{m=-l}^{l} B_{lm}(r_1) \cdot B_{lm}(r_2) \,.$$

are related to the auto-correlation function of the 3D density map over SO(3).

Spatial autocorrelation features

- The density map has the following real spherical harmonic expansion,

$$\rho(\boldsymbol{r}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} B_{lm}(r) \cdot Y_{lm}(\theta_{\boldsymbol{r}}, \varphi_{\boldsymbol{r}}) \,.$$

- Spatial autocorrelation features

$$E_l(r_1, r_2) = \sum_{m=-l}^{l} B_{lm}(r_1) \cdot B_{lm}(r_2) \,.$$

  are related to the auto-correlation function of the 3D density map over SO(3).
- We further show that C_l is connected with E_l as follows

$$C_l(k_1, k_2) = 8\pi^2 \int_0^{\infty} \left(\int_0^{\infty} E_l(r_1, r_2) \cdot j_l(k_1 r_1) r_1^2 \; dr_1 \right) \cdot j_l(k_2 r_2) r_2^2 \; dr_2 \,.$$

Spatial autocorrelation features

- The density map has the following real spherical harmonic expansion,

$$\rho(\boldsymbol{r}) = \sum_{l=0}^{\infty} \sum_{m=-l}^{l} B_{lm}(r) \cdot Y_{lm}(\theta_{\boldsymbol{r}}, \varphi_{\boldsymbol{r}})\,.$$

- Spatial autocorrelation features

$$E_l(r_1, r_2) = \sum_{m=-l}^{l} B_{lm}(r_1) \cdot B_{lm}(r_2)\,.$$

  are related to the auto-correlation function of the 3D density map over SO(3).
- We further show that C_l is connected with E_l as follows

$$C_l(k_1, k_2) = 8\pi^2 \int_0^{\infty} \left(\int_0^{\infty} E_l(r_1, r_2) \cdot j_l(k_1 r_1) r_1^2 \, dr_1 \right) \cdot j_l(k_2 r_2) r_2^2 \, dr_2\,.$$

- Use the quadrature points and weights according to Gauss-Legendre Quadrature rule, we express the relations in the matrix form,

$$\boldsymbol{A}_l = \boldsymbol{Q}_l^* \boldsymbol{B}_l, \quad \boldsymbol{F}_l \boldsymbol{O}_l = \boldsymbol{Q}_l^* \boldsymbol{B}_l(\rho), \quad \boldsymbol{B}_l(\rho) \text{ is linear in } \rho$$

Spatial radial features

- We average $\widehat{I}_n(\boldsymbol{k})$ over all directions of $\boldsymbol{k}$ with the same norm k:

$$M_N(k) = \frac{1}{N}\sum_{n=1}^{N}\frac{1}{2\pi}\int_0^{2\pi}\widehat{I}_n(\boldsymbol{k})\,d\varphi$$

Spatial radial features

- We average $\widehat{I}_n(\boldsymbol{k})$ over all directions of $\boldsymbol{k}$ with the same norm k:

$$M_N(k) = \frac{1}{N}\sum_{n=1}^{N}\frac{1}{2\pi}\int_0^{2\pi}\widehat{I}_n(\boldsymbol{k})\,d\varphi \xrightarrow{N\to\infty} \underbrace{\iiint \frac{\sin(kr)}{kr}\cdot\rho(\boldsymbol{r})\,d\boldsymbol{r}}_{=:\,M(k)}$$

Spatial radial features

- We average $\widehat{I}_n(\boldsymbol{k})$ over all directions of $\boldsymbol{k}$ with the same norm k:

$$M_N(k) = \frac{1}{N}\sum_{n=1}^{N}\frac{1}{2\pi}\int_0^{2\pi}\widehat{I}_n(\boldsymbol{k})\ d\varphi \xrightarrow{N\to\infty} \underbrace{\iiint \frac{\sin(kr)}{kr}\cdot\rho(\boldsymbol{r})\ d\boldsymbol{r}}_{=:\,M(k)}$$

- The integration of the 3D density map $\rho(\boldsymbol{r})$ on the sphere with radius r is given by

$$W(r) = \frac{2r}{\pi}\int_0^{\infty} k\cdot M(k)\cdot\sin(kr)\ dk = \iiint \rho(\tilde{\boldsymbol{r}})\cdot\delta(\tilde{r}-r)\ d\tilde{\boldsymbol{r}}\,.$$

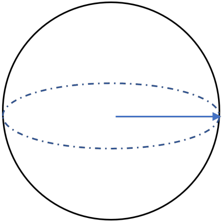

Integration of ρ on a sphere

Spatial radial features

- We average $\widehat{I}_n(\boldsymbol{k})$ over all directions of $\boldsymbol{k}$ with the same norm k:

$$M_N(k) = \frac{1}{N}\sum_{n=1}^{N}\frac{1}{2\pi}\int_0^{2\pi}\widehat{I}_n(\boldsymbol{k})\;d\varphi \xrightarrow{N\to\infty} \underbrace{\iiint \frac{\sin(kr)}{kr}\cdot\rho(\boldsymbol{r})\;d\boldsymbol{r}}_{=:\,M(k)}$$

- The integration of the 3D density map $\rho(\boldsymbol{r})$ on the sphere with radius r is given by

$$W(r) = \frac{2r}{\pi}\int_0^{\infty} k\cdot M(k)\cdot\sin(kr)\;dk = \iiint \rho(\tilde{\boldsymbol{r}})\cdot\delta(\tilde{r}-r)\;d\tilde{\boldsymbol{r}}\,.$$

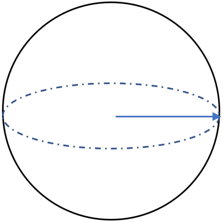

Integration of ρ on a sphere

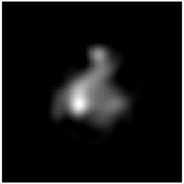

Denoised reference image

Orthogonal matrix retrieval with spatial consensus

We formulate the following nonconvex problem to recover the density in the spatial domain:

$$\underset{\rho,\{\boldsymbol{O}_l\}_{l=0}^{L}}{\text{minimize}} \quad f(\rho, \boldsymbol{O}_l) := \sum_{l=0}^{L} \|\boldsymbol{F}_l \boldsymbol{O}_l - \boldsymbol{Q}_l^* \boldsymbol{B}_l(\rho)\|_2^2$$

autocorrelation features

Orthogonal matrix retrieval with spatial consensus

We formulate the following nonconvex problem to recover the density in the spatial domain:

$$\begin{aligned}\underset{\rho,\{\boldsymbol{O}_l\}_{l=0}^{L}}{\text{minimize}} \quad f(\rho, \boldsymbol{O}_l) := \sum_{l=0}^{L} \|\boldsymbol{F}_l \boldsymbol{O}_l - \boldsymbol{Q}_l^* \boldsymbol{B}_l(\rho)\|_2^2 & \qquad \text{autocorrelation features} \\ + \lambda \cdot \sum_{v=1}^{V} \left(\boldsymbol{g}_v^T \rho - W_v\right)^2 & \qquad \text{radial features}\end{aligned}$$

Orthogonal matrix retrieval with spatial consensus

We formulate the following nonconvex problem to recover the density in the spatial domain:

$$
\begin{aligned}
\underset{\rho,\{\boldsymbol{O}_l\}_{l=0}^{L}}{\text{minimize}} \quad f(\rho,\boldsymbol{O}_l) := & \sum_{l=0}^{L} \|\boldsymbol{F}_l\boldsymbol{O}_l - \boldsymbol{Q}_l^* \boldsymbol{B}_l(\rho)\|_2^2 && \text{autocorrelation features} \\
& + \lambda \cdot \sum_{v=1}^{V} \left(\boldsymbol{g}_v^T \rho - W_v\right)^2 && \text{radial features} \\
& + \xi \cdot \sum_{x,y} \left(P_\rho(x,y) - \overline{I}(x,y)\right)^2, && \text{a denoised reference image}
\end{aligned}
$$

Orthogonal matrix retrieval with spatial consensus

We formulate the following nonconvex problem to recover the density in the spatial domain:

$$\underset{\rho, \{\boldsymbol{O}_l\}_{l=0}^{L}}{\text{minimize}} \quad f(\rho, \boldsymbol{O}_l) := \sum_{l=0}^{L} \|\boldsymbol{F}_l \boldsymbol{O}_l - \boldsymbol{Q}_l^* \boldsymbol{B}_l(\rho)\|_2^2 \qquad \text{autocorrelation features}$$

$$+ \lambda \cdot \sum_{v=1}^{V} \left(\boldsymbol{g}_v^T \rho - W_v \right)^2 \qquad \text{radial features}$$

$$+ \xi \cdot \sum_{x,y} \left(P_\rho(x,y) - \overline{I}(x,y) \right)^2, \qquad \text{a denoised reference image}$$

$$\text{subject to} \quad 0 \le \rho_d \le W_\rho,$$

$$\sum_{d=1}^{D} \rho_d = W_\rho \,,$$

Orthogonal matrix retrieval with spatial consensus

We formulate the following nonconvex problem to recover the density in the spatial domain:

$$\underset{\rho,\{\boldsymbol{O}_l\}_{l=0}^{L}}{\text{minimize}} \quad f(\rho, \boldsymbol{O}_l) := \sum_{l=0}^{L} \|\boldsymbol{F}_l \boldsymbol{O}_l - \boldsymbol{Q}_l^* \boldsymbol{B}_l(\rho)\|_2^2 \qquad \text{autocorrelation features}$$

$$+ \lambda \cdot \sum_{v=1}^{V} \left(\boldsymbol{g}_v^T \rho - W_v \right)^2 \qquad \text{radial features}$$

$$+ \xi \cdot \sum_{x,y} \left(P_\rho(x,y) - \overline{I}(x,y) \right)^2, \qquad \text{a denoised reference image}$$

$$\text{subject to} \quad 0 \leq \rho_d \leq W_\rho,$$

$$\sum_{d=1}^{D} \rho_d = W_\rho\,,$$

$$\boldsymbol{O}_l^T \boldsymbol{O}_l = \boldsymbol{O}_l \boldsymbol{O}_l^T = \boldsymbol{I}, \quad l \in \{0, \ldots, L\}\,.$$

Alternating minimization

Given an initialization ρ_0, we use the following alternating minimization approach:

- **($\boldsymbol{O}$-update)**. Fix ρ, and update $\{\boldsymbol{O}_l\}_l$ with respect to ρ:

$$\begin{aligned} \underset{\{\boldsymbol{O}_l\}_{l=0}^{L}}{\text{minimize}} \quad & f_1(\boldsymbol{O}_l) := \sum_{l=0}^{L} \|\boldsymbol{F}_l \boldsymbol{O}_l - \boldsymbol{Q}_l^* \boldsymbol{B}_l(\rho)\|_2^2 \\ \text{subject to} \quad & \boldsymbol{O}_l^T \boldsymbol{O}_l = \boldsymbol{O}_l \boldsymbol{O}_l^T = \boldsymbol{I}, \quad l \in \{0, \ldots, L\}. \end{aligned}$$

Alternating minimization

Given an initialization ρ_0, we use the following alternating minimization approach:

- **($\boldsymbol{O}$-update)**. Fix ρ, and update $\{\boldsymbol{O}_l\}_l$ with respect to ρ:

$$\begin{aligned} \underset{\{\boldsymbol{O}_l\}_{l=0}^{L}}{\text{minimize}} \quad & f_1(\boldsymbol{O}_l) := \sum_{l=0}^{L} \|\boldsymbol{F}_l \boldsymbol{O}_l - \boldsymbol{Q}_l^* \boldsymbol{B}_l(\rho)\|_2^2 \\ \text{subject to} \quad & \boldsymbol{O}_l^T \boldsymbol{O}_l = \boldsymbol{O}_l \boldsymbol{O}_l^T = \boldsymbol{I}, \quad l \in \{0, \ldots, L\}\,. \end{aligned}$$

Closed-form solutions for the orthogonal matrices $\{\boldsymbol{O}_l\}_l$ are given by

$$\boldsymbol{O}_l = \boldsymbol{V}_l \boldsymbol{U}_l^T\,,$$

where $\boldsymbol{V}_l$ and $\boldsymbol{U}_l$ are obtained from the singular value decomposition (SVD) of $\boldsymbol{B}_l^T(\rho)\boldsymbol{Q}_l\boldsymbol{F}_l$

Alternating minimization

- **(ρ-update)**. Fix $\{\boldsymbol{O}_l\}_l$, and estimate ρ with respect to $\{\boldsymbol{O}_l\}_l$:

$$
\begin{aligned}
\underset{\rho}{\text{minimize}} \quad & f_2(\rho) := \sum_{l=0}^{L} \|\boldsymbol{F}_l \boldsymbol{O}_l - \boldsymbol{Q}_l^* \boldsymbol{B}_l(\rho)\|_2^2 \\
& \qquad + \lambda \cdot \sum_{v=1}^{V} \left(\boldsymbol{g}_v^T \rho - W_v\right)^2 \\
& \qquad + \xi \cdot \sum_{x,y} \left(P_\rho(x,y) - \overline{I}(x,y)\right)^2, \\
\text{subject to} \quad & 0 \le \rho_d \le W_\rho, \\
& \sum_{d=1}^{D} \rho_d = W_\rho \,.
\end{aligned}
$$

- A convex problem which can be solved using projected gradient descent (PGD).

An Initialization

- Compute an initialization ρ_0 by solving the following problem:

$$
\begin{aligned}
\underset{\rho_0}{\text{minimize}} \quad & \sum_{v=1}^{V} \left(\boldsymbol{g}_v^T \rho_0 - W_v \right)^2 + \sum_{x,y} \left(P_{\rho_0}(x,y) - \overline{I}(x,y) \right)^2, \\
\text{subject to} \quad & 0 \le {\rho_0}_d \le W_\rho, \\
& \sum_{d=1}^{D} {\rho_0}_d = W_\rho \,.
\end{aligned}
$$

An Initialization

- Compute an initialization ρ_0 by solving the following problem:

$$\begin{aligned} \underset{\rho_0}{\text{minimize}} \quad & \sum_{v=1}^{V} \left(\boldsymbol{g}_v^T \rho_0 - W_v \right)^2 + \sum_{x,y} \left(P_{\rho_0}(x,y) - \overline{I}(x,y) \right)^2, \\ \text{subject to} \quad & 0 \leq {\rho_0}_d \leq W_\rho, \\ & \sum_{d=1}^{D} {\rho_0}_d = W_\rho \,. \end{aligned}$$

- We use projected gradient descent initialized at $\mathbf{0}$ that favors the minimum l_2-norm solution.

Reconstructions

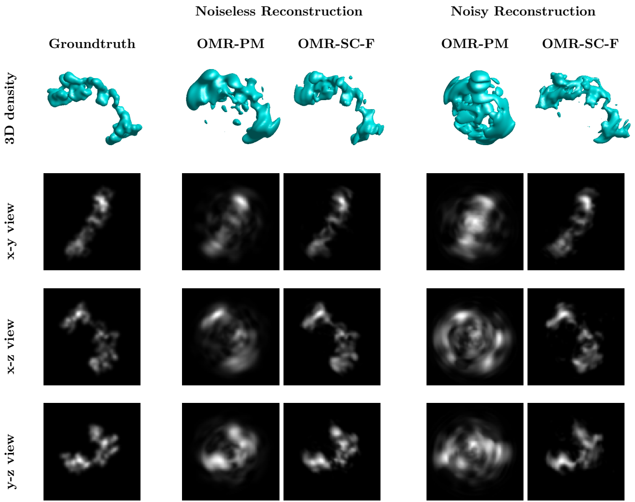

Reconstructions

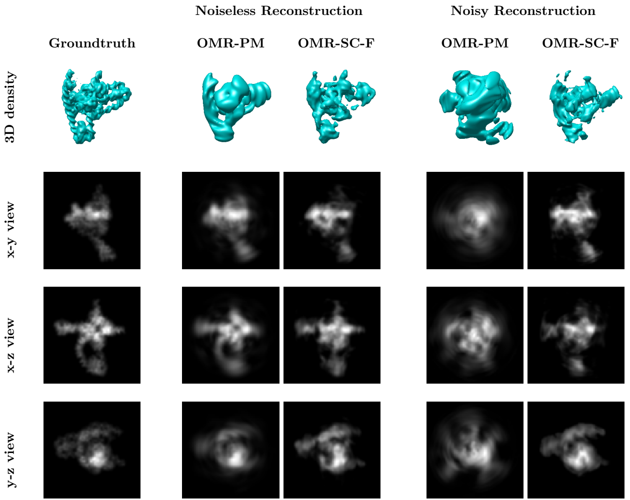

Holliday junction complex (HJC): reconstructions using the OMR-PM and OMR-SC-F approaches in the noiseless case and the noisy case (SNR=0.1).

Sparsity

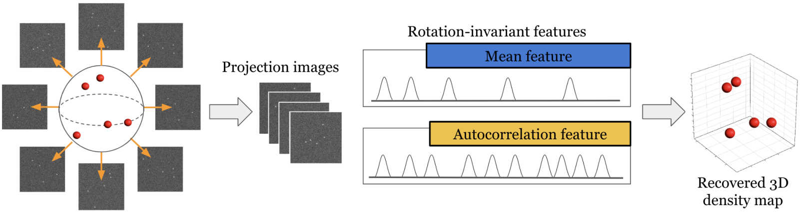

- A subset of the autocorrelation features can recover the density modeled as a mixture of Gaussians *.

---

* M. Zehni et al. "3D unknown view tomography via rotation invariants" *ICASSP*, 2020.

† T. Bendory et al. "Autocorrelation analysis for cryo-EM with sparsity constraints: Improved sample complexity and projection-based algorithms" *PNAS*, 2023.

Sparsity

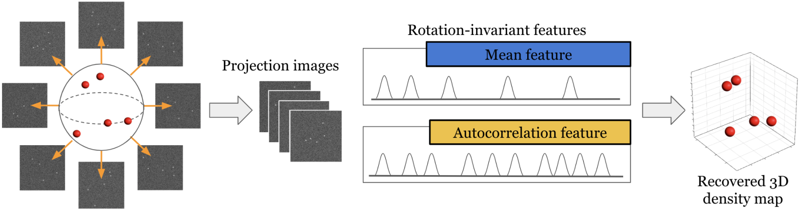

- A subset of the autocorrelation features can recover the density modeled as a mixture of Gaussians *.
- Bendory et al.† proved that 3D density maps modeled as sums of Gaussians (with some assumptions) are uniquely determined by the second-order autocorrelation of their projection images.

---

* M. Zehni et al. "3D unknown view tomography via rotation invariants" *ICASSP*, 2020.

† T. Bendory et al. "Autocorrelation analysis for cryo-EM with sparsity constraints: Improved sample complexity and projection-based algorithms" *PNAS*, 2023.

Sparsity

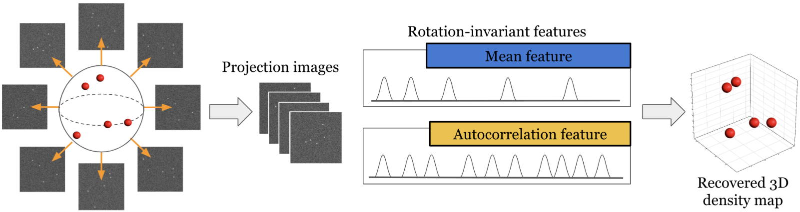

- A subset of the autocorrelation features can recover the density modeled as a mixture of Gaussians *.
- Bendory et al.† proved that 3D density maps modeled as sums of Gaussians (with some assumptions) are uniquely determined by the second-order autocorrelation of their projection images.
- The sample complexity is proportional to σ^4.

* M. Zehni et al. "3D unknown view tomography via rotation invariants" *ICASSP*, 2020.
† T. Bendory et al. "Autocorrelation analysis for cryo-EM with sparsity constraints: Improved sample complexity and projection-based algorithms" *PNAS*, 2023.

Summary and discussions

- Using Kam's autocorrelation analysis with non-negative constraint on the voxel values, we are able to recover the 3D volumes without estimating the individual orientation

Summary and discussions

- Using Kam's autocorrelation analysis with non-negative constraint on the voxel values, we are able to recover the 3D volumes without estimating the individual orientation
- How to incorporate more prior information about the density maps

Summary and discussions

- Using Kam's autocorrelation analysis with non-negative constraint on the voxel values, we are able to recover the 3D volumes without estimating the individual orientation
- How to incorporate more prior information about the density maps
- Require more thorough theoretical analysis on the uniqueness guarantee given certain types of spatial constraints or prior information

Summary and discussions

- Using Kam's autocorrelation analysis with non-negative constraint on the voxel values, we are able to recover the 3D volumes without estimating the individual orientation
- How to incorporate more prior information about the density maps
- Require more thorough theoretical analysis on the uniqueness guarantee given certain types of spatial constraints or prior information
- Need to understand the landscape of the associated optimization problems

Introduction

- Scientific imaging is used to characterize physical, material, chemical, and biological processes.

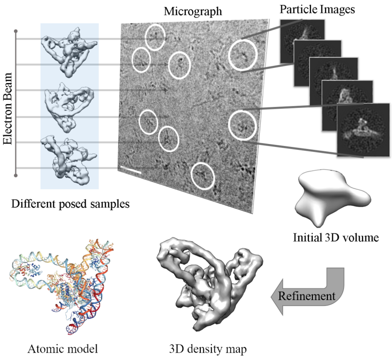

Cryo-electron microscopy single particle analysis

Imaging data challenges include

- Incomplete data
- measurement uncertainty (e.g., noise, unknown latent parameters)

Here we focus on **tomography with unknown view angles.**

Part II. Combine distribution matching and physics-based imaging process

Joint work with Mona Zehni

2D Tomography with unknown view angles

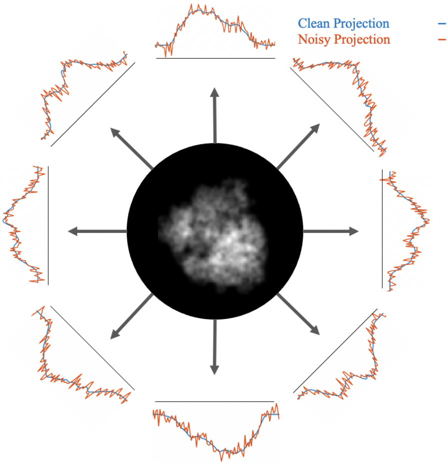

Noisy measurements:

$$\zeta_\ell = \mathcal{P}_{\theta_\ell} I + \varepsilon_\ell$$

$$\theta_\ell \sim p(\theta)$$

Goal

Recover I and p given unordered $\{\zeta_\ell\}_{\ell=1}^{L}$.

2D tomography with unknown view angles

- Viewing angle estimation (angle recovery problem[1,2], graph Laplacian tomography[3,4], Moment-based angular difference estimates (MADE)[5]) and tomographic reconstruction
- Maximum likelihood estimation or maximum a posteriori estimation using expectation-maximization algorithm.
- Generative adversarial learning based approach[6,7]

---

[1] Basu and Bresler, Uniqueness of tomography with unknown view angles, *IEEE TIP*, 2000

[2] Basu and Bresler, Feasibility of tomography with unknown view angles, *IEEE TIP*, 2000

[3] Coifman et al., Graph Laplacian Tomography From Unknown Random Projections, *IEEE TIP*, 2008

[4] Singer and Wu, Two-dimensional tomography from noisy projections taken at unknown random directions, *SIIMS*, 2013

[5] Phan et al., Moment-Based Angular Difference Estimation Between Two Tomographic Projections in 2D and 3D, *J Math Imaging Vis* 2016.

[6] Gupta et al. CryoGAN: A New Reconstruction Paradigm for Single-particle Cryo-EM Via Deep Adversarial Learning, *IEEE TCI*, 2021.

[7] Zehni and Zhao, UVTomo-GAN: An adversarial learning based approach for unknown view X-ray tomographic reconstruction, *ISBI*, 2021

Viewing angle estimation

Graph Laplacian tomography:

- Apply PCA and Wiener filtering to denoise the projections.
- Compute pairwise distance between the denoised projections $d_{ij} = \|\zeta_i - \zeta_j\|_2$.
- Construct k nearest neighbor graph followed by graph denoising.
- Use the top two non-trivial eigenvectors of the normalized similarity matrix to determine the order of the projections.

Viewing angle estimation

Graph Laplacian tomography:

- Apply PCA and Wiener filtering to denoise the projections.
- Compute pairwise distance between the denoised projections $d_{ij} = \|\zeta_i - \zeta_j\|_2$.
- Construct k nearest neighbor graph followed by graph denoising.
- Use the top two non-trivial eigenvectors of the normalized similarity matrix to determine the order of the projections.

Main Challenge: When the viewing angle distribution is not uniform, it is hard to get the precise angle from the ordering.

Viewing angle estimation

Graph Laplacian tomography:

- Apply PCA and Wiener filtering to denoise the projections.
- Compute pairwise distance between the denoised projections $d_{ij} = \|\zeta_i - \zeta_j\|_2$.
- Construct k nearest neighbor graph followed by graph denoising.
- Use the top two non-trivial eigenvectors of the normalized similarity matrix to determine the order of the projections.

Main Challenge: When the viewing angle distribution is not uniform, it is hard to get the precise angle from the ordering.

MADE: Use the the moment-based angular difference estimates instead of the ℓ_2 distance between projections.

Viewing angle estimation

Graph Laplacian tomography:

- Apply PCA and Wiener filtering to denoise the projections.
- Compute pairwise distance between the denoised projections $d_{ij} = \|\zeta_i - \zeta_j\|_2$.
- Construct k nearest neighbor graph followed by graph denoising.
- Use the top two non-trivial eigenvectors of the normalized similarity matrix to determine the order of the projections.

Main Challenge: When the viewing angle distribution is not uniform, it is hard to get the precise angle from the ordering.

MADE: Use the the moment-based angular difference estimates instead of the ℓ_2 distance between projections.
Main Challenge: The projection moment-based features used in the MADE are sensitive to noise.

Bypassing individual view angle estimation?

- **Question:** Can we achieve good image reconstruction without the estimation of individual view angles?

[1] Gupta et al. CryoGAN: A New Reconstruction Paradigm for Single-particle Cryo-EM Via Deep Adversarial Learning, *IEEE TCI*, 2021.

[2] Zehni and Zhao, UVTomo-GAN: An adversarial learning based approach for unknown view X-ray tomographic reconstruction, *ISBI*, 2021.

Bypassing individual view angle estimation?

- **Question:** Can we achieve good image reconstruction without the estimation of individual view angles?
- **Yes!** One approach is to use the generative adversarial learning[1,2].

[1] Gupta et al. CryoGAN: A New Reconstruction Paradigm for Single-particle Cryo-EM Via Deep Adversarial Learning, *IEEE TCI*, 2021.

[2] Zehni and Zhao, UVTomo-GAN: An adversarial learning based approach for unknown view X-ray tomographic reconstruction, *ISBI*, 2021.

Bypassing individual view angle estimation?

- **Question:** Can we achieve good image reconstruction without the estimation of individual view angles?
- **Yes!** One approach is to use the generative adversarial learning[1][2].
- **Challenge:** recover both the image and the viewing angle distribution simultaneously.

[1] Gupta et al. CryoGAN: A New Reconstruction Paradigm for Single-particle Cryo-EM Via Deep Adversarial Learning, *IEEE TCI*, 2021.

[2] Zehni and Zhao, UVTomo-GAN: An adversarial learning based approach for unknown view X-ray tomographic reconstruction, *ISBI*, 2021.

Reconstruction based on distribution matching

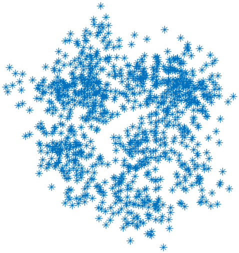

Reconstruction based on distribution matching

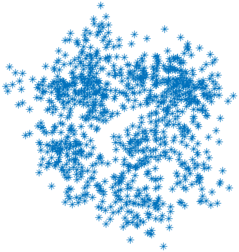

Reconstruction based on distribution matching

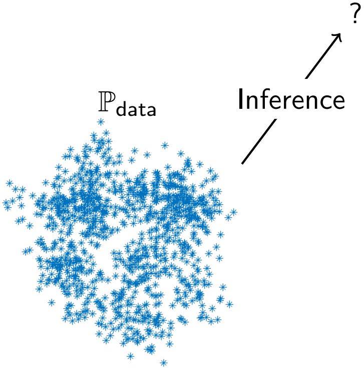

Reconstruction based on distribution matching

$$I,\ p(\theta)$$

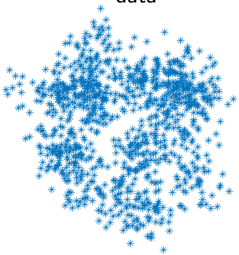

Reconstruction based on distribution matching

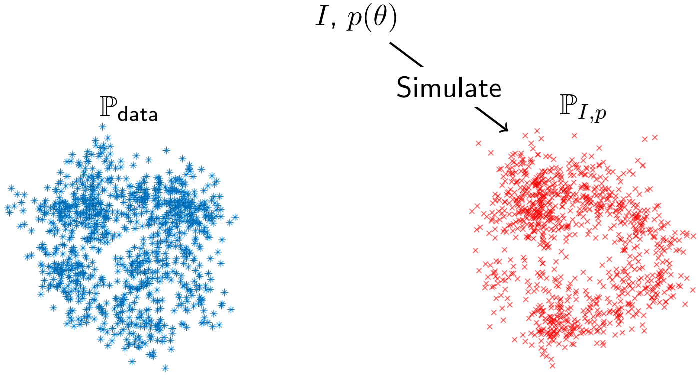

Reconstruction based on distribution matching

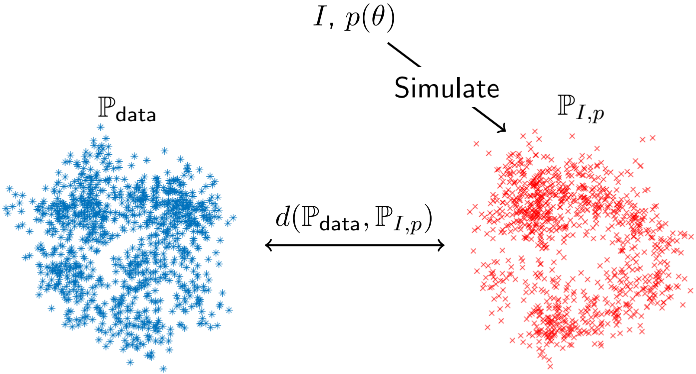

Reconstruction based on distribution matching

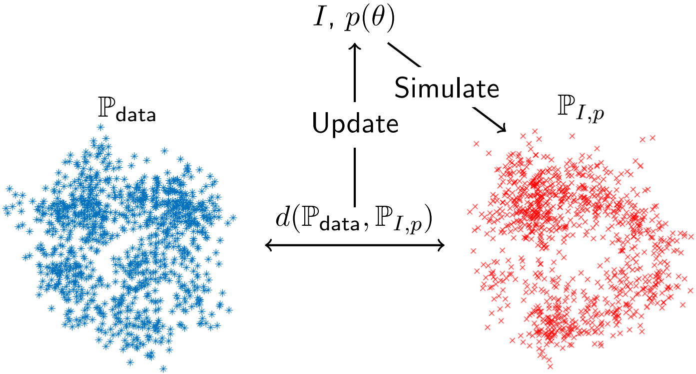

Statistical inference is performed by minimizing, over the parameter space, the distance / divergence between model distribution and the empirical distribution of the data.

Comparing probability distributions

- KL divergence, Jensen-Shannon divergence, total variation, MMD...

Comparing probability distributions

- KL divergence, Jensen-Shannon divergence, total variation, MMD...
- The Earth-Mover's distance or Wasserstein-1 distance

$$W(\mathbb{P}_{\mathsf{data}}, \mathbb{P}_{I,p}) = \inf_{\gamma \in \prod(\mathbb{P}_{\mathsf{data}}, \mathbb{P}_{I,p})} \mathbb{E}_{(x,y)\sim\gamma}[\|x - y\|].$$

Comparing probability distributions

- KL divergence, Jensen-Shannon divergence, total variation, MMD...
- The Earth-Mover's distance or Wasserstein-1 distance

$$W(\mathbb{P}_{\mathsf{data}}, \mathbb{P}_{I,p}) = \inf_{\gamma \in \prod(\mathbb{P}_{\mathsf{data}}, \mathbb{P}_{I,p})} \mathbb{E}_{(x,y)\sim\gamma}[\|x - y\|].$$

- The Earth-Mover's distance or Wasserstein-1 distance using Kantorovich-Rubinstein duality

$$W(\mathbb{P}_{\mathsf{data}}, \mathbb{P}_{I,p}) = \sup_{\|f\|_L \leq 1} \mathbb{E}_{x\sim\mathbb{P}_{\mathsf{data}}}[f(x)] - \mathbb{E}_{x\sim\mathbb{P}_{I,p}}[f(x)].$$

Comparing probability distributions

- KL divergence, Jensen-Shannon divergence, total variation, MMD...
- The Earth-Mover's distance or Wasserstein-1 distance

$$W(\mathbb{P}_{\mathsf{data}}, \mathbb{P}_{I,p}) = \inf_{\gamma \in \prod(\mathbb{P}_{\mathsf{data}}, \mathbb{P}_{I,p})} \mathbb{E}_{(x,y)\sim\gamma}[\|x - y\|].$$

- The Earth-Mover's distance or Wasserstein-1 distance using Kantorovich-Rubinstein duality

$$W(\mathbb{P}_{\mathsf{data}}, \mathbb{P}_{I,p}) = \sup_{\|f\|_L \leq 1} \mathbb{E}_{x\sim\mathbb{P}_{\mathsf{data}}}[f(x)] - \mathbb{E}_{x\sim\mathbb{P}_{I,p}}[f(x)].$$

- Earth-Mover's distance has nicer properties than Jensen-Shannon divergence or KL divergence in learning distributions, specifically for distributions supported on lower dimensional manifolds.

Image representation

We use the Hartley transform of the images, which is a real representation closely related to Fourier transform and defined as:

$$\widetilde{I} = \mathcal{H}(I) = \mathrm{real}\{\mathcal{F}(I)\} - \mathrm{imag}\{\mathcal{F}(I)\},$$

- Projection slice theorem in Hartley domain: $\widetilde{\zeta}_\ell(\xi) = \widetilde{I}(\xi, \theta_\ell)$

[1] A. Klug and R. Crowther, "Three-dimensional image reconstruction from the viewpoint of information theory," Nature, vol. 238, no. 5365, pp. 435–440, 1972.

[2] Zhao and Singer, "Fourier–Bessel rotational invariant eigenimages", JOSA A, vol. 30, pp. 871-877, 2013

[3] Zhao, Shkolnisky, and Singer, "Fast steerable principal component analysis", IEEE TCI, vol. 2, pp. 1-12, 2016

Image representation

We use the Hartley transform of the images, which is a real representation closely related to Fourier transform and defined as:

$$\widetilde{I} = \mathcal{H}(I) = \mathrm{real}\{\mathcal{F}(I)\} - \mathrm{imag}\{\mathcal{F}(I)\},$$

- Projection slice theorem in Hartley domain: $\widetilde{\zeta}_\ell(\xi) = \widetilde{I}(\xi, \theta_\ell)$
- Truncated expansion of $\widetilde{I}$ on Hartley-Bessel (HB) basis

$$\widetilde{I}(\xi, \theta) \approx \sum_{k=-k_{\max}}^{k_{\max}} \sum_{q=1}^{p_k} c_{k,q} u_s^{k,q}(\xi, \theta).$$

[1] A. Klug and R. Crowther, "Three-dimensional image reconstruction from the viewpoint of information theory," Nature, vol. 238, no. 5365, pp. 435–440, 1972.

[2] Zhao and Singer, "Fourier–Bessel rotational invariant eigenimages", JOSA A, vol. 30, pp. 871-877, 2013

[3] Zhao, Shkolnisky, and Singer, "Fast steerable principal component analysis", IEEE TCI, vol. 2, pp. 1-12, 2016

Image representation

We use the Hartley transform of the images, which is a real representation closely related to Fourier transform and defined as:

$$\widetilde{I} = \mathcal{H}(I) = \mathrm{real}\{\mathcal{F}(I)\} - \mathrm{imag}\{\mathcal{F}(I)\},$$

- Projection slice theorem in Hartley domain: $\widetilde{\zeta}_\ell(\xi) = \widetilde{I}(\xi, \theta_\ell)$
- Truncated expansion of $\widetilde{I}$ on Hartley-Bessel (HB) basis

$$\widetilde{I}(\xi, \theta) \approx \sum_{k=-k_{\max}}^{k_{\max}} \sum_{q=1}^{p_k} c_{k,q} u_s^{k,q}(\xi, \theta).$$

- Rewrite the forward model in transformed domain as:

$$\widetilde{\zeta}_\ell = H_{\theta_\ell} c + \widetilde{\varepsilon}_\ell, \quad \theta_\ell \sim p, \quad \ell \in \{1, 2, ..., L\}$$

[1] A. Klug and R. Crowther, "Three-dimensional image reconstruction from the viewpoint of information theory," Nature, vol. 238, no. 5365, pp. 435–440, 1972.

[2] Zhao and Singer, "Fourier–Bessel rotational invariant eigenimages", JOSA A, vol. 30, pp. 871-877, 2013

[3] Zhao, Shkolnisky, and Singer, "Fast steerable principal component analysis", IEEE TCI, vol. 2, pp. 1-12, 2016

Image reconstruction from truncated expansion coefficients

From the truncated expansion coefficient c, we can reconstruct the image in the spatial domain,

$$I(r, \varphi) = \sum_{k=-k_{\max}}^{k_{\max}} \sum_{q=1}^{p_k} c_{k,q} \, \mathcal{H}\left(u_s^{k,q}\right)(r, \varphi)$$

where

$$\mathcal{H}\left(u_s^{k,q}\right)(r, \varphi) = \frac{2\sqrt{2\pi} s (-1)^{(q+l)} R_{k,q} J_k(2\pi s r)}{(2\pi s r)^2 - R_{k,q}^2} \cos(k\varphi + \frac{\pi}{4}),$$

and $l = \frac{k+1}{2}$ for odd k and $l = \frac{k}{2}$ for even k.

An adversarial learning based reconstruction

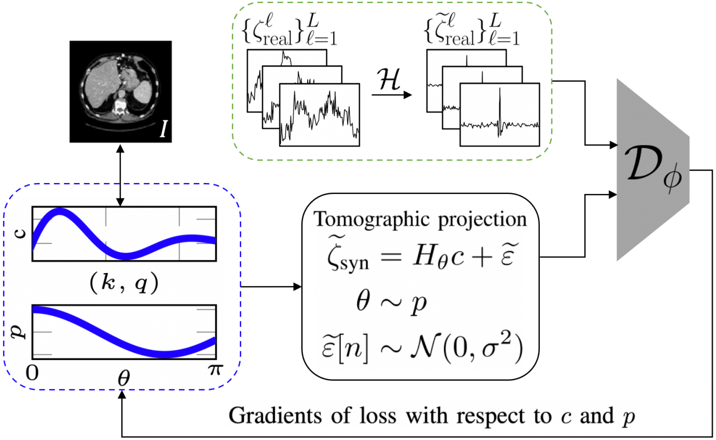

Wasserstein generative adversarial learning

- Follow the Wasserstein GAN, the loss is

$$\mathcal{L}(c,p,\phi) = \sum_{b=1}^{B} \mathcal{D}_\phi\left(\widetilde{\zeta}_{\text{real}}^{b}\right) - \mathcal{D}_\phi\left(\widetilde{\zeta}_{\text{syn}}^{b}\right) + \lambda \mathrm{GP}(\widetilde{\zeta}_{\text{int}}^{b})$$

- Generator loss:

$$\begin{aligned}\mathcal{L}_G(c,p) &= -\sum_{b=1}^{B} \mathcal{D}_\phi(H_{\theta_b} c + \widetilde{\varepsilon}_b), \\ &= -\sum_{b=1}^{B} \sum_{t=1}^{N_\theta} \delta(\theta_t - \theta_b) \mathcal{D}_\phi(H_{\theta_t} c + \widetilde{\varepsilon}_b),\ \theta_b \sim p.\end{aligned}$$

Wasserstein generative adversarial learning

- Follow the Wasserstein GAN, the loss is

$$\mathcal{L}(c,p,\phi)=\sum_{b=1}^{B}\mathcal{D}_\phi\left(\widetilde{\zeta}^b_{\text{real}}\right)-\mathcal{D}_\phi\left(\widetilde{\zeta}^b_{\text{syn}}\right)+\lambda\text{GP}(\widetilde{\zeta}^b_{\text{int}})$$

- Generator loss:

$$\begin{aligned}\mathcal{L}_G(c,p)&=-\sum_{b=1}^{B}\mathcal{D}_\phi(H_{\theta_b}c+\widetilde{\varepsilon}_b),\\ &=-\sum_{b=1}^{B}\sum_{t=1}^{N_\theta}\delta(\theta_t-\theta_b)\mathcal{D}_\phi(H_{\theta_t}c+\widetilde{\varepsilon}_b),\ \theta_b\sim p.\end{aligned}$$

- $\mathcal{L}_G$ is non-differentiable with respect to p. How to allow backpropagation?

Approximation of the generator loss

- Discretize projection angle domain to N_θ equal bins.

[1] E. Jang et al., Categorical reparameterization with Gumbel-softmax, ICLR, 2017

Approximation of the generator loss

- Discretize projection angle domain to N_θ equal bins.
- Use Gumbel-Softmax[1] to allow backpropagation through samples from discrete distributions.

[1] E. Jang et al., Categorical reparameterization with Gumbel-softmax, ICLR, 2017

Approximation of the generator loss

- Discretize projection angle domain to N_θ equal bins.
- Use Gumbel-Softmax[1] to allow backpropagation through samples from discrete distributions.
- Approximate $\mathcal{L}_G$ such that it is differentiable in both c and p.

[1] E. Jang et al., Categorical reparameterization with Gumbel-softmax, ICLR, 2017

Approximation of the generator loss

- The Gumbel-Max trick:

$$\delta(\theta_t - \theta_b) = \text{one hot}\left(\underset{i}{\text{argmax}}[g_{b,i} + \log(p[i])]\right),\ g_{b,i} \sim \text{Gumbel}(0,1)$$

[1] E. Jang et al., Categorical reparameterization with Gumbel-softmax, ICLR, 2017

Approximation of the generator loss

- The Gumbel-Max trick:

$$\delta(\theta_t - \theta_b) = \text{one hot}\left(\underset{i}{\operatorname{argmax}}[g_{b,i} + \log(p[i])]\right),\ g_{b,i} \sim \mathrm{Gumbel}(0,1)$$

- Use softmax function as a continuous, differentiable approximation to argmax[1]: $r_{i,b}(p) = \frac{\exp\left((g_{b,i}+\log(p[i]))/\tau\right)}{\sum_{j=1}^{N_\theta} \exp\left((g_{b,j}+\log(p[j]))/\tau\right)}$

[1] E. Jang et al., Categorical reparameterization with Gumbel-softmax, ICLR, 2017

Approximation of the generator loss

- The Gumbel-Max trick:

$$\delta(\theta_t - \theta_b) = \text{one hot}\left(\underset{i}{\operatorname{argmax}}[g_{b,i} + \log(p[i])]\right),\ g_{b,i} \sim \mathrm{Gumbel}(0,1)$$

- Use softmax function as a continuous, differentiable approximation to argmax[1]: $r_{i,b}(p) = \frac{\exp\left((g_{b,i}+\log(p[i]))/\tau\right)}{\sum_{j=1}^{N_\theta} \exp\left((g_{b,j}+\log(p[j]))/\tau\right)}$
- The approximate generator loss becomes:

$$\mathcal{L}_G(c,p) \approx -\sum_{b=1}^{B}\sum_{i=1}^{N_\theta} r_{i,b}(p)\mathcal{D}_\phi(H_{\theta_i}c + \widetilde{\varepsilon}_b).$$

[1] E. Jang et al., Categorical reparameterization with Gumbel-softmax, ICLR, 2017

Approximation of the generator loss

- The Gumbel-Max trick:

$$\delta(\theta_t - \theta_b) = \text{one hot}\left(\underset{i}{\operatorname{argmax}}[g_{b,i} + \log(p[i])]\right),\; g_{b,i} \sim \mathrm{Gumbel}(0,1)$$

- Use softmax function as a continuous, differentiable approximation to argmax[1]: $r_{i,b}(p) = \frac{\exp\left((g_{b,i}+\log(p[i]))/\tau\right)}{\sum_{j=1}^{N_\theta} \exp\left((g_{b,j}+\log(p[j]))/\tau\right)}$
- The approximate generator loss becomes:

$$\mathcal{L}_G(c,p) \approx -\sum_{b=1}^{B}\sum_{i=1}^{N_\theta} r_{i,b}(p)\mathcal{D}_\phi(H_{\theta_i}c + \widetilde{\varepsilon}_b).$$

- Differentiable with respect to both c and p.

[1] E. Jang et al., Categorical reparameterization with Gumbel-softmax, ICLR, 2017

Numerical results

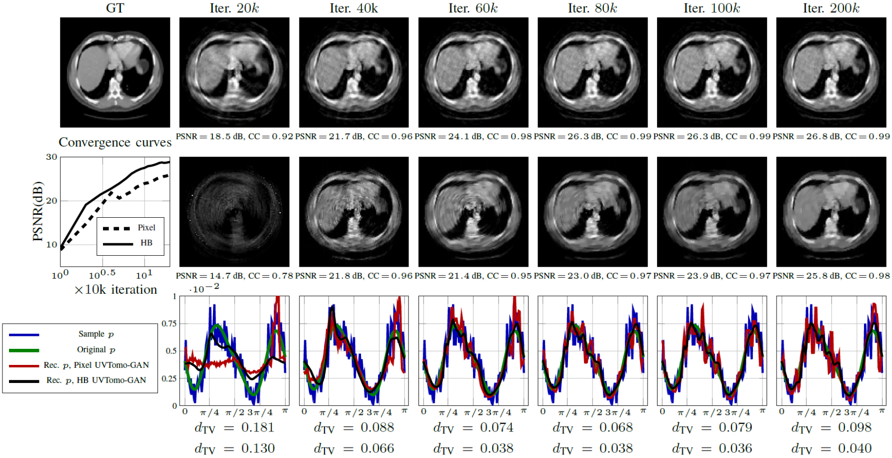

Numerical results–reconstructions from clean projections

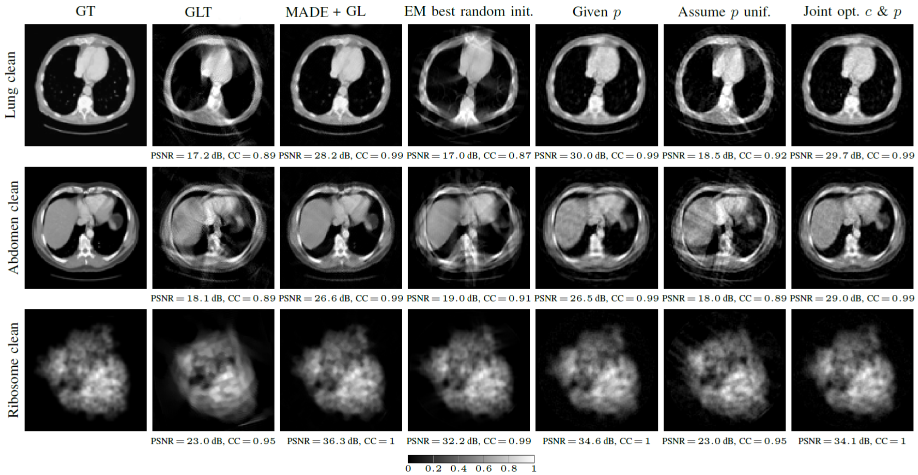

Numerical results–at SNR= 3

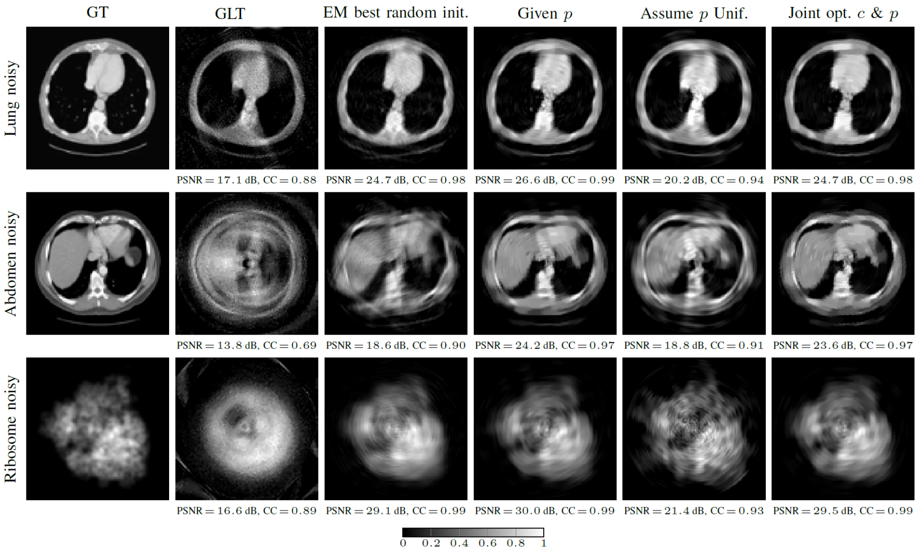

Numerical results–viewing angle distribution

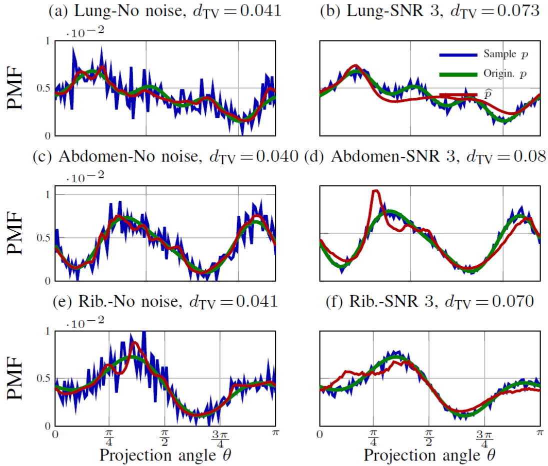

Summary and discussions

- We introduce an adversarial learning approach to jointly reconstruct image and viewing angle distribution from random 2D tomographic projections.
- To update the viewing angle distribution via backpropagation, we use Gumbel-softmax approximation of samples from categorical distribution.
- Future directions: Extend to 3D reconstruction, understand the sample complexity, adapt to different measurement uncertainty.

References

- H. Shuai, M. Zehni, I. Dokmanic, and Z. Zhao. “Orthogonal Matrix Retrieval with Spatial Consensus for 3D Unknown-View Tomography” *SIAM Journal on Imaging Sciences*, **16** (3), 1398-1439, 2023

- M. Zehni and Z. Zhao. “An adversarial learning based approach for unknown view tomographic reconstruction.” *IEEE TCI*, 2022

- M. Zehni and Z. Zhao. “UVTomo-GAN: An adversarial learning based approach for unknown view X-ray tomographic reconstruction.” *ISBI*, 2021

- M. Zehni, H. Shuai, I. Dokmanic, and Z. Zhao. “3D unknown view tomography via rotation invariants” *ICASSP*, 2020.

Thank You!

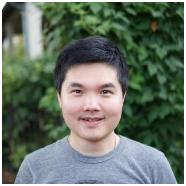

ALFRED P. SLOAN
FOUNDATION