

# GEM: 3D GAUSSIAN SPLATTING FOR EFFICIENT AND ACCURATE CRYO-EM RECONSTRUCTION

Anonymous authors

Paper under double-blind review

## ABSTRACT

Cryo-electron microscopy (cryo-EM) has become a central tool for high-resolution structural biology, yet the massive scale of datasets (often exceeding 100k particle images) renders 3D reconstruction both computationally expensive and memory intensive. Traditional Fourier-space methods are efficient but lose fidelity due to repeated transforms, while recent real-space approaches based on neural radiance fields (NeRFs) improve accuracy but incur cubic memory and computation overhead. Therefore, we introduce GEM, a novel cryo-EM reconstruction framework built on 3D Gaussian Splatting (3DGS) that operates directly in real-space while maintaining high efficiency. Instead of modeling the entire density volume, GEM represents proteins with compact 3D Gaussians, each parameterized by only 11 values. To further improve the training efficiency, we designed a novel gradient computation to 3D Gaussians that contribute to each voxel. This design substantially reduced both memory footprint and training cost. On standard cryo-EM benchmarks, GEM achieves up to  $48\times$  faster training and  $12\times$  lower memory usage compared to state-of-the-art methods, while improving local resolution by as much as 38.8%. These results establish GEM as a practical and scalable paradigm for cryo-EM reconstruction, unifying speed, efficiency, and high-resolution accuracy.

## 1 INTRODUCTION

Cryo-electron microscopy (cryo-EM) (Bai et al., 2015a; Milne et al., 2013b) is a transformative biological imaging technique that enables the determination of macromolecular structures at near-atomic resolution, and its significance was recognized with the 2017 Nobel Prize in Chemistry (Pre). In cryo-EM, biological samples are rapidly frozen in vitreous ice and imaged under an electron microscope, producing thousands to millions of noisy two-dimensional projection images of individual particles. Each projection corresponds to a different orientation of the same underlying macromolecule. By aggregating these projections, the three-dimensional electron density of the protein can be reconstructed (Murata & Wolf, 2018). The cryo-EM imaging model naturally aligns with the Fourier slice theorem, where each projection corresponds to a central slice of the 3D Fourier transform of the density, enabling efficient reconstruction in Fourier space (Zhong et al., 2021; Punjani et al., 2017; Scheres, 2012; Tang et al., 2007).

Despite its efficiency, the reliance on Fourier transforms introduces information loss in the traditional methods, limiting achievable resolution. Recent works propose adapting neural radiance fields (NeRFs) to cryo-EM by modeling protein densities directly in real-space (Huang et al., 2024; Qu et al., 2025; Liu et al., 2023), using representations such as MLPs (Liu et al., 2023), Transformers (Huang et al., 2024), or multi-resolution hash encoding (Qu et al., 2025). These approaches remove the dependence on Fourier transforms in the Fourier slice theorem (see Appendix 6.1 for more details) and demonstrate improved resolution compared to traditional Fourier-space pipelines. However, direct application of NeRF to cryo-EM remains impractical. Unlike view synthesis tasks, cryo-EM requires reconstruction of the full 3D density

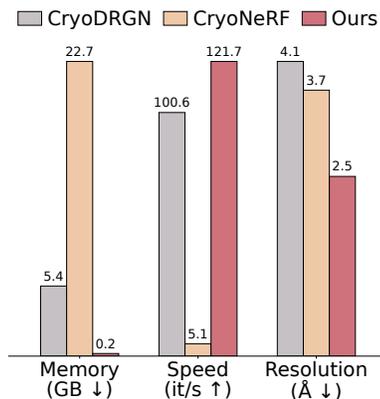


Figure 1: GEM achieves lower memory usage, faster speed, and higher reconstruction resolution compared to existing approaches.

rather than rendering sparse pixels. And each training step involves applying the contrast transfer function (CTF) across entire projection images, resulting in cubic scaling of memory and computation. Even for small proteins, this requires millions of sampled points and gradients, making NeRF-based cryo-EM reconstruction prohibitively slow and memory-intensive, especially on commodity GPUs.

To address these challenges, we propose *GEM*, a novel framework that introduces 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023; Zha et al., 2024) for cryo-EM reconstruction. Unlike NeRF-based approaches that model density implicitly, *GEM* adopts an explicit representation using a set of 3D Gaussians, each parameterized by only 11 values. This representation is both compact and efficient: (i) 3D Gaussians are instantiated only at nonzero-density regions, and during projection, (ii) each 3D Gaussian contributes to multiple pixels and can be rasterized in parallel. Furthermore, (iii) gradient computation is restricted to the subset of Gaussians influencing a given pixel, in contrast to NeRF methods that require activating the entire network for each pixel. This design eliminates the cubic memory footprint of NeRFs and yields substantial acceleration in training.

We validate *GEM* on four widely used cryo-EM datasets under diverse evaluation protocols. As shown in Figure 1, *GEM* achieves substantially faster training and lower memory usage than the real-space NeRF-based pipeline CryoNeRF, while also surpassing both the Fourier-space baseline CryoDRGN and the real-space baseline CryoNeRF in reconstruction quality. Across all datasets, *GEM* consistently attains gold-standard Fourier shell correlation (GSFSC) resolutions better than 3 Å, outperforming existing methods. Additional evaluations further confirm its resolution advantage, and compared to CryoNeRF (Qu et al., 2025), *GEM* not only improves quality but also delivers significantly higher efficiency, even surpassing the Fourier-based CryoDRGN.

Our main contributions are summarized as follows:

- We introduce *GEM*, a novel framework that incorporates 3D Gaussian Splatting into cryo-EM reconstruction, providing an explicit, efficient, and accurate representation of protein density that avoids the information loss of Fourier-based methods and the cubic overhead of NeRF-based real-space approaches.
- By explicitly modeling protein density with 3D Gaussians placed only in nonzero-density regions and enabling parallel rasterization across pixels, *GEM* achieves up to 48× faster training speed compared to existing methods.
- We design an efficient gradient computation strategy that restricts updates to only the 3D Gaussians contributing to each pixel, thereby eliminating the cubic memory footprint and computation overhead of prior real-space methods.
- Extensive experiments on four widely used cryo-EM datasets under diverse evaluation protocols demonstrate that *GEM* consistently outperforms existing approaches, often reaching resolutions close to the instrumental limitation.

## 2 RELATED WORK

**Cryo-EM reconstruction.** Classical single-particle cryo-EM pipelines formulate reconstruction in Fourier space via the Fourier slice theorem, solving for the 3D density with iterative refinement (Zhong et al., 2021; Punjani et al., 2017; Scheres, 2012; Tang et al., 2007). This formulation enables efficient global solvers and robust regularization schemes and has underpinned much of the field’s progress (Milne et al., 2013a; Murata & Wolf, 2018; Bai et al., 2015b). Quality is typically assessed using gold-standard Fourier shell correlation (GSFSC) (Van Heel & Schatz, 2005), local resolution (Glaeser et al., 2021) and directional consistency analyses (Huang et al., 2024). Despite their maturity, Fourier-space methods involve repeated FFTs and interpolation steps and operate on band-limited representations, which can incur information loss. Our work targets these pain points by adopting 3D Gaussian Splatting as an explicit, efficient real-space representation while remaining compatible with standard cryo-EM pipelines and evaluation protocols.

**Real-space cryo-EM reconstruction.** A growing line of work reconstructs density directly in Euclidean space using neural implicit fields and related parameterizations. Neural-field approaches, which often use MLPs with Fourier or hash encodings, or Transformer-based fields, optimize a continuous density so that its simulated projections match observed images, thereby avoiding

frequency-domain interpolation and enabling real-space regularization (Huang et al., 2024; Qu et al., 2025; Liu et al., 2023; Herreros et al., 2025; Palmer & Aylett, 2022; Lu et al., 2022). These methods have demonstrated improved resolution on several benchmarks. However, directly training neural fields for cryo-EM is computationally demanding: unlike view synthesis, we must reconstruct full volumes rather than sparse rays, apply CTFs per micrograph, and backpropagate through dense sampling, leading to memory/time that scale roughly with the volume (and number of samples) instead of the number of visible pixels. This cubic-like footprint makes large proteins and commodity GPUs challenging and motivates representations whose gradients remain localized to the subset of primitives that actually contribute to each pixel.

**3D Gaussian splatting.** 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023; Zha et al., 2024; Wu et al., 2024; Fei et al., 2024; Yu et al., 2024; Ye et al., 2025) represents scenes as explicit sets of anisotropic Gaussians rendered by fast rasterization, yielding state-of-the-art real-time view synthesis with memory that scales with the number of Gaussians rather than the discretized volume (Kerbl et al., 2023). For biological settings, recent work proposes rectified radiative Gaussian splatting to improve fidelity under line/slice integrals (Zha et al., 2024; 2022; Yu et al., 2025), suggesting that splatting is a promising alternative to implicit fields when projections are the fundamental measurements. Adapting 3DGS to cryo-EM introduces two key advantages: (i) sparsity, since Gaussians need only occupy non-empty regions of the macromolecule, and (ii) locality, since gradients for a pixel involve only the Gaussians that cover it, avoiding the cubic memory of dense neural fields. Our framework instantiates these ideas for single-particle cryo-EM by coupling a CTF-aware, differentiable projection of Gaussians with joint optimization over density and imaging latents.

### 3 METHOD

#### 3.1 CRYO-EM RECONSTRUCTION

Cryo-EM reconstruction (Elmlund & Elmlund, 2015) aims to recover the 3D structure of a protein from a large collection of noisy 2D particle images. In practice, each image is acquired as a projection of a protein particle illuminated by an electron beam, modulated by the microscope’s contrast transfer function (CTF) and corrupted by background noise.

**Image Formation Model.** Let  $\mathcal{I} = \{I_i\}_{i=1}^N$  denote a dataset of  $N$  particle images and  $\mathcal{V} = \{V_i\}_{i=1}^N$  the corresponding particles. Each image  $I_i \in \mathbb{R}^{d \times d}$  in Figure 3 is a projection of an underlying 3D density map  $V_i \in \mathbb{R}^{d \times d \times d}$  following the imaging model that can be expressed as

$$I_i = C_i * \text{Proj}(V_i; \phi_i, t_i) + \eta_i, \quad (1)$$

where  $C_i$  is the CTF determined by microscope settings and  $*$  is the convolution operation.  $\text{Proj}(V) = \int_{-\infty}^0 V dz$  denotes the projection operator with rotation angles  $\phi_i$  and translation  $t_i$ , and  $\eta_i$  represents additive noise.

**Homogeneous Reconstruction.** In this work, we focus on homogeneous reconstruction, where all particles, e.g.  $V_1$ ,  $V_2$ , and  $V_3$  in Figure 3, share the same underlying structure:

$$V_i = V, \forall V_i \in \mathcal{V}. \quad (2)$$

The goal is then to recover the shared 3D density  $V$  from all images and their associated parameters:

$$\hat{V} = \text{Reconstruct}(\mathcal{I}; \Phi, \mathcal{T}, \mathcal{C}), \quad (3)$$

where  $\Phi = \{\phi_i\}_{i=1}^N$ ,  $\mathcal{T} = \{t_i\}_{i=1}^N$ , and  $\mathcal{C} = \{C_i\}_{i=1}^N$  denote the set of rotations, translations and CTFs, which are known parameters before the reconstruction.

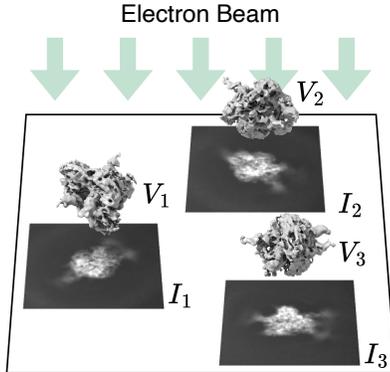


Figure 2: Cryo-EM imaging process. Randomly oriented proteins are illuminated by the electron beam, leaving traces on the sensor as particle images  $I_1$ ,  $I_2$ , and  $I_3$ . In this work we focus on homogeneous reconstruction, where all particles  $V_1$ ,  $V_2$ , and  $V_3$  share the same underlying 3D structure and differ only in orientation (i.e., multiple copies of the same object).

**Fourier Slice Theorem for Reconstruction:** Fourier-space methods leverage the Fourier slice theorem and therefore only need to predict a slice in the Fourier domain, i.e.,  $\text{Slice}(\mathcal{F}(\hat{V}_i); \phi_i, t_i)$ , without instantiating the full protein density  $\hat{V}$ . **Real-Space Cryo-EM Reconstruction:** In contrast, real-space methods instantiate the entire volume  $\hat{V}$  and perform projection following Equation 1, which requires cubic memory to store both the predicted values and their gradients.

### 3.2 3D GAUSSIAN SPLATTING FOR RECONSTRUCTION

We present the GEM training framework in Figure 4. First, protein density can be represented as a set of 3D Gaussians (Figure 3a). To optimize the randomly initialized 3D Gaussian-based representation, we can project 3D Gaussian representations to 2D space. The 3D Gaussian representations can then be optimized by minimizing the differences between the projected 2D images and the captured 2D images by the real microscope. During training, we also introduced novel efficient gradient computation to reduce memory cost. After training, the reconstructed protein density can be queried from the trained 3D Gaussians (Figure 3b).

**Protein Density represented by 3D Gaussians.** Following previous 3DGS literature (Kerbl et al., 2023; Zha et al., 2024), we model the protein density as a set of 3D Gaussians  $\mathcal{G} = \{G_j\}_{j=1}^M$ , where each Gaussian is parameterized by its center  $\mathbf{p}_j$ , covariance matrix  $\Sigma_j$ , and density coefficient  $\rho_j$ . Each kernel  $G_j$  defines a local Gaussian-shaped density field:

$$G_j(\mathbf{x} | \rho_j, \mathbf{p}_j, \Sigma_j) = \rho_j \cdot \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{p}_j)^\top \Sigma_j^{-1}(\mathbf{x} - \mathbf{p}_j)\right), \quad \Sigma_j = \mathbf{R}_j \mathbf{S}_j \mathbf{S}_j^\top \mathbf{R}_j^\top. \quad (4)$$

where the covariance  $\Sigma_j$  is further decomposed into a rotation  $\mathbf{R}_j$  and a scale matrix  $\mathbf{S}_j$ . The overall density of protein  $V_i$  at position  $\mathbf{x}$  is then given by the summation of all Gaussian kernels:

$$\hat{V}_i(\mathbf{x}) = \sum_{j=1}^M G_j(\mathbf{x} | \rho_j, \mathbf{p}_j, \Sigma_j). \quad (5)$$

Compared to the original 3DGS formulation, which incorporates view-dependent effects for neural rendering, our variant adopts an isotropic, density-based, density-based formulation consistent with the cryo-EM imaging model, where reconstruction relies on where reconstruction relies on direct density integration along the electron beam.

**Fast 2D Projection of 3D Gaussian representations via Explicit Integration for .** Owing to the property of 3D Gaussians, the projection can be expressed in a closed form, avoiding querying density values and then aggregating them. Let  $\hat{V}_i(\mathbf{x})$  denote the protein density for view  $i$ . The volume is defined as a sum of 3D 3D Gaussians following Equation 5, and the integration along the optical axis  $z$  yields a 2D Gaussian mixture:

$$\hat{I}_i(\mathbf{r}) = \sum_{j=1}^M \rho_j \frac{1}{2\pi |\hat{\Sigma}_j|^{1/2}} \exp\left(-\frac{1}{2}(\hat{\mathbf{x}} - \hat{\mathbf{p}}_j)^\top \hat{\Sigma}_j^{-1}(\hat{\mathbf{x}} - \hat{\mathbf{p}}_j)\right) = \sum_{j=1}^M G_j^2, \quad (6)$$

where  $\hat{\mathbf{x}} \in \mathbb{R}^2$ ,  $\hat{\mathbf{p}}_j \in \mathbb{R}^2$ , and  $\hat{\Sigma}_j \in \mathbb{R}^{2 \times 2}$  are the in-plane (mean, covariance) obtained by dropping the third row/column of their 3D counterparts  $\mathbf{x}$ ,  $\mathbf{p}_j$ , and  $\Sigma_j$ , respectively. This closed-form marginalization enables efficient CUDA kernels and removes the need to explicitly query  $\hat{V}_i$  on 3D grids before the projection, substantially reducing both compute and memory. Please see Appendix 6.2 for more detail of the derivation of the projection. As shown in Figure 4, Step1 applies Equation 6 to generate the projection from the protein density.

**Optimizing 3D GS representations** Given the projected image  $\hat{I}_i$ , we apply the CTF  $C_i$  in Fourier space and compare to the noisy measurement  $I_i$  via an  $\ell_2$  loss:

$$I_i^{\text{pred}} = \mathcal{F}^{-1}(\mathcal{F}(C_i) \cdot \mathcal{F}(\hat{I}_i)), \quad \mathcal{L}_i = \|I_i^{\text{pred}} - I_i\|_2^2, \quad (7)$$

where  $\mathcal{F}$  and  $\mathcal{F}^{-1}$  are Fourier and inverse Fourier transforms, respectively. In Equation 7, Fourier transform is utilized as the convolution in Equation 1 equals to the multiplication in Fourier space

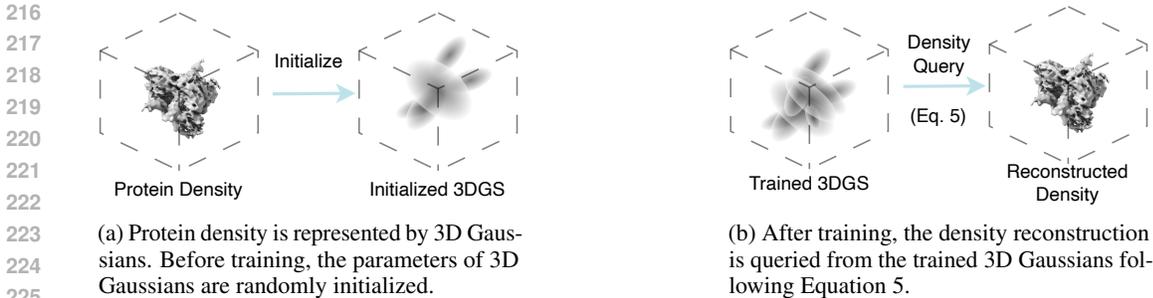


Figure 3: Initialization before training and density reconstruction after training.

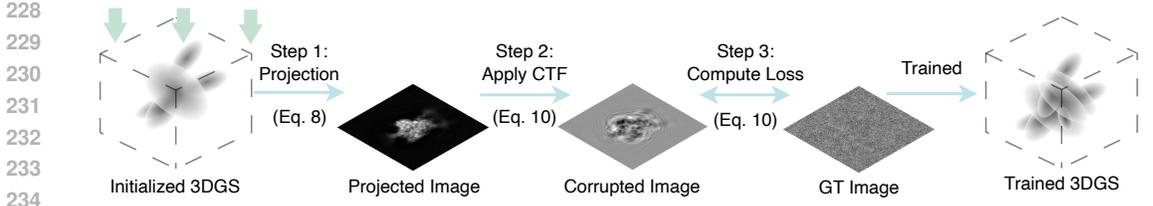


Figure 4: The overview of GEM training. The training begins with randomly initialized Gaussians. The 3D Gaussians are projected following Equation 6. Then the CTF is applied (Equation 7) to the projection and being compared with the noisy experimental image to calculate the loss (Equation 7).

(McGillem & Cooper, 1991). Although 3DGS is typically trained on clean natural images, we observe that the standard MSE loss is sufficient for noisy cryo-EM images. Equation 7 corresponds to Step 2 and 3 in Figure 4, where the CTF is first applied to the projected image, followed by computing the standard MSE loss between the CTF-corrected image and the experimental ground-truth image.

**Efficient Gradient Computation to Reduce Memory Overhead.** While our rasterization implementation follows the principle of Equation 6, directly evaluating all 3D Gaussians remains inefficient. To address this, we introduce a thresholding strategy that selects only the Gaussians with non-negligible contributions to a given ray  $\mathbf{r}$ . Formally,

$$\hat{I}_i(\mathbf{r}) = \sum_{j=1}^K G_j^2, \quad \text{where } G_j^2 > \tau, \tag{8}$$

where  $\tau$  is a predefined threshold. This pruning avoids unnecessary computation from Gaussians far from the ray, since their contribution decays quadratically with distance to the Gaussian center  $\mathbf{p}_j$ .

After selection, the remaining 3D Gaussians are sorted along the  $z$ -axis, and the accumulation in Equation 8 is performed from the lowest  $z$ -value (closest to the cryo-EM sensor) to the highest (farthest from the sensor). This ordering prioritizes the most influential 3D Gaussians during rendering and further improves computational efficiency.

## 4 EXPERIMENT

In this section, we evaluate the proposed method GEM using four widely adopted cryo-EM datasets: EMPIAR-10005 (TRPV1, Liao et al. (2013)), EMPIAR-10028 (*Plasmodium falciparum* 80S ribosome, Wong et al. (2014)), EMPIAR-10049 (Synaptic RAG1-RAG2 complex, Ru et al. (2015)), and EMPIAR-10076 (L17-depleted 50S ribosomal intermediates, Davis et al. (2016)).

To demonstrate the effectiveness of GEM, we compare it with two representative cryo-EM reconstruction approaches: (i) the conventional voxel-based method CryoSPARC (Punjani et al., 2017), and (ii) the neural network-based CryoDRGN (Zhong et al., 2021). Both of the methods reconstruct cryo-EM densities in Fourier space leveraging the Fourier slice theorem (Radon, 2005; Garces et al., 2011), in contrast to our GEM that operates in 3D Euclidean space. For CryoSPARC, we first perform *ab initio*

Table 1: Efficiency comparison of deep learning-based cryo-EM reconstruction methods. Speed (it/s) measures the number of images processed per second, an Memor (GB) reports the maximum GPU memory consumption during training. All experiments are conducted on a single NVIDIA RTX A6000 (48 GB). OOM indicates out-of-memory errors. **Our method achieves the highest training throughput while requiring substantially less memory across all datasets.**

	EMPIAR-10028		EMPIAR-10049		EMPIAR-10076		EMPIAR-10005	
	Speed (it/s) $\uparrow$	Memory (GB) $\downarrow$						
CryoDRGN	47.62	8.81	142.86	2.58	58.82	8.79	100.00	4.93
CryoNeRF	-	OOM	9.13	9.94	2.10	37.75	4.03	20.50
Ours	<b>94.10</b>	<b>1.54</b>	<b>160.49</b>	<b>0.50</b>	<b>96.60</b>	<b>0.74</b>	<b>108.11</b>	<b>0.51</b>

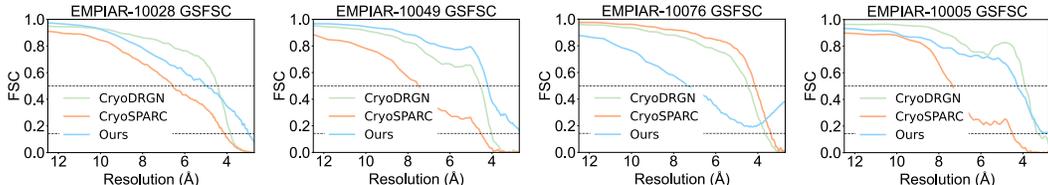


Figure 5: GSFSC of our GEM and baselines. The horizontal axis denotes resolution in ångströms (Å), and the vertical axis denotes the FSC value. The two dashed horizontal lines are FSC thresholds of 0.5 and 0.143. The final resolution is defined as the resolution at which the GSFSC curve first drops below the 0.143 threshold (smallest possible resolution if no intersection). **Intersection points further to the right corresponds to better reconstruction quality.**

Table 2: Final resolutions (in Å) obtained by each method according to the GSFSC 0.143 threshold. **Lower values indicate better resolution.**

	EMPIAR-10028	EMPIAR-10049	EMPIAR-10076	EMPIAR-10005
CryoSPARC	4.12 Å	4.53 Å	3.39 Å	4.50 Å
CryoDRGN	3.81 Å	4.02 Å	3.66 Å	3.27 Å
Ours	<b>2.98 Å</b>	<b>2.56 Å</b>	<b>2.62 Å</b>	<b>2.43 Å</b>

reconstruction, followed by homogeneous refinement. For CryoDRGN, we employ the `train_nn` function, which is designed for homogeneous reconstruction without modeling heterogeneity. In both cases, we use their default parameters. In the following sections, we first evaluate the reconstruction efficiency, and then examine the reconstructed protein densities from GEM, CryoSPARC and CryoDRGN with three mainstream evaluation protocols: gold-standard Fourier shell correlation (Harauz & van Heel (1986), Section 4.2), local resolution estimation (Adams et al. (2010), Section 4.3), and Fourier slice correlation (Huang et al. (2024), Section 4.4).

#### 4.1 EFFICIENCY COMPARISON

We first evaluate the efficiency of GEM against deep learning-based cryo-EM reconstruction methods. Specifically, we compare with CryoDRGN (Fourier-based) and CryoNeRF (NeRF-based). All experiments are conducted on a single NVIDIA RTX 6000 GPU, and we report both training throughput (images per second) and peak GPU memory usage. As shown in Table 1, GEM consistently achieves the best efficiency, with substantially lower memory usage and faster speed than the baselines. CryoNeRF suffers from cubic memory and computation overhead inherent to NeRF, resulting in very slow training and frequent out-of-memory (OOM) failures. CryoDRGN, which leverages the Fourier slice theorem with quadratic complexity, is faster than CryoNeRF but still significantly outperformed by GEM. Moreover, GEM exhibits more stable memory usage, as it depends primarily on the number of 3D Gaussians used to represent the protein density.

#### 4.2 GOLD-STANDARD FOURIER SHELL CORRELATION

Gold-standard Fourier shell correlation (GSFSC) (Van Heel & Schatz, 2005) is a commonly used evaluation to quantify the reconstruction quality of a cryo-EM reconstruction method on a dataset.

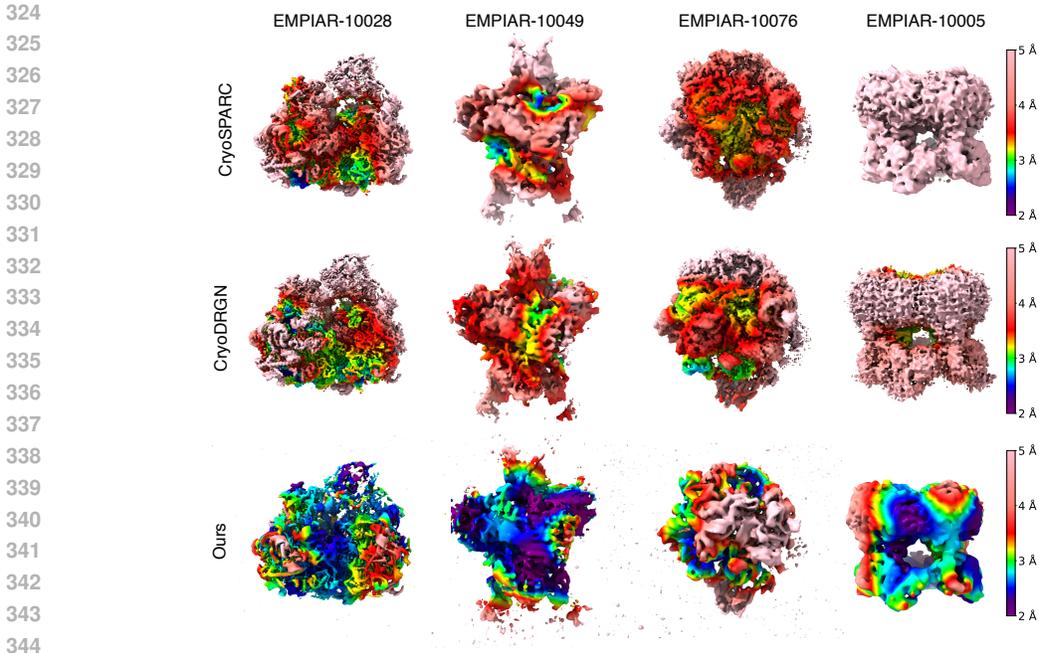


Figure 6: Local resolution maps of reconstructions from GEM, CryoSPARC, and CryoDRGN. Each map is colored according to the estimated local resolution, with the color scale ranging from 2 Å (blue, higher resolution) to 5 Å (red, lower resolution). **Better reconstructions are indicated by larger regions in blue-green.**

Table 3: Mean local resolution (in Å) of reconstructions on the four cryo-EM datasets. **Lower values correspond to higher resolution.**

	EMPIAR-10028	EMPIAR-10049	EMPIAR-10076	EMPIAR-10005
CryoSPARC	3.55 Å	3.57 Å	3.16 Å	5.54 Å
CryoDRGN	3.39 Å	3.22 Å	3.22 Å	3.38 Å
Ours	<b>2.87 Å</b>	<b>2.62 Å</b>	<b>2.99 Å</b>	<b>2.58 Å</b>

Given a cryo-EM particle image dataset  $\mathcal{I} = \{I_i\}_{i=1}^N$ , to evaluate a method with GSFSC, the dataset  $\mathcal{I}$  is first randomly splitted into two non-overlapping splits  $\mathcal{I}_1$  and  $\mathcal{I}_2$ . Then the method to test is applied on the two splits to produce two reconstructions  $X_1$  and  $X_2$ . The FSC value at spatial frequency shell  $r_i$  in the Fourier space is defined as:

$$FSC(r_i) = \frac{\sum_{r \in r_i} F_1(r) \cdot F_2(r)^*}{\sqrt{\sum_{r \in r_i} |F_1(r)|^2 \cdot \sum_{r \in r_i} |F_2(r)|^2}}, \tag{9}$$

where  $F_1(r)$  and  $F_2(r)$  denote the Fourier coefficients of  $X_1$  and  $X_2$  at frequency  $r$ , and  $*$  indicates complex conjugation. The summations are computed over all voxels in the shell  $r_i$ . As in the GSFSC protocol the particle images are split into two halves  $\mathcal{I}_1$  and  $\mathcal{I}_2$  and reconstructed independently, it ensures that the correlation reflects reproducible structural information rather than overfitting.

Since the spatial frequency  $r$  is inversely related to the 3D Euclidean space resolution, the GSFSC values can be plotted as a curve of FSC values versus resolutions (in Å). The resolution estimate is then determined by the point at which the GSFSC curve drops below a predefined threshold (commonly 0.143). This crossing point represents the highest spatial frequency (lowest resolution value) at which the reconstruction remains reliable, thereby providing a quantitative resolution assessment.

As shown in Figure 5 and Table 2, **●** GEM consistently achieves higher resolutions (i.e., lower Å values at the intersection with the 0.143 FSC threshold) than the baselines across all datasets. It is

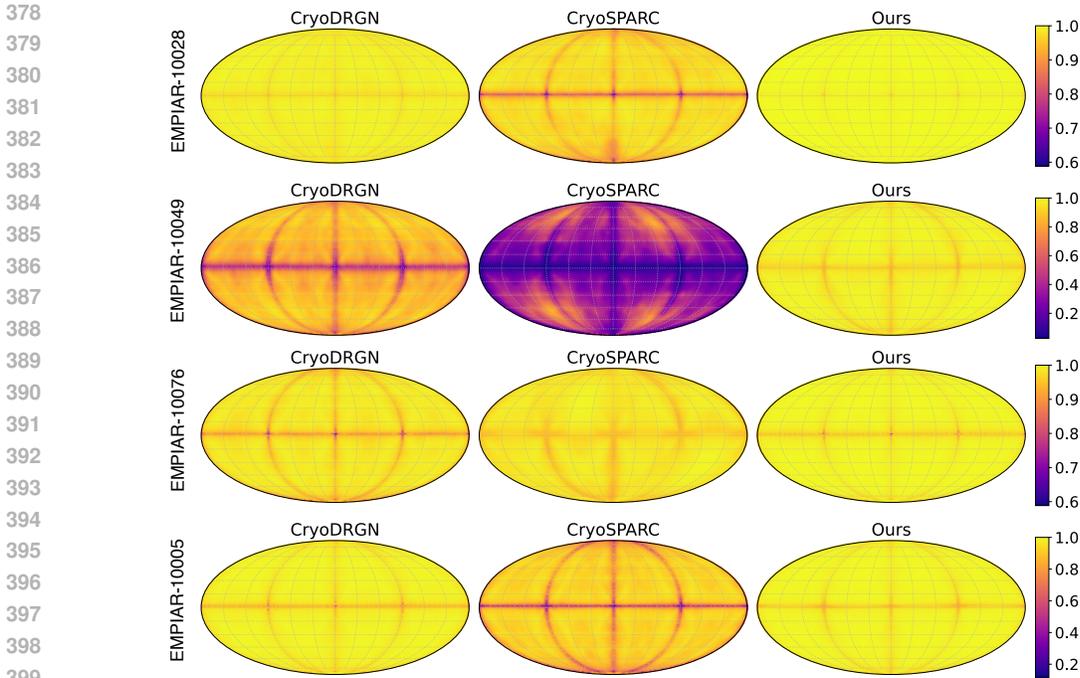


Figure 7: FSLC maps of reconstructions from GEM, CryoSPARC, and CryoDRGN across four datasets. Each map depicts the correlation values of slices sampled across different elevation and azimuth angles, with the color scale ranging from low correlation (purple) to high correlation (yellow). **More uniform yellow regions indicate higher and more isotropic directional resolution.**

Table 4: FSLC results on four cryo-EM datasets. Each row reports the mean FSLC value with its standard deviation across randomly sampled orientations. **Higher mean values indicate stronger directional consistency, and lower standard deviation reflects greater stability.**

	EMPIAR-10028	EMPIAR-10049	EMPIAR-10076	EMPIAR-10005
CryoSPARC	0.961 ± 0.030	0.348 ± 0.160	0.978 ± 0.012	0.886 ± 0.084
CryoDRGN	0.989 ± 0.005	0.814 ± 0.109	0.976 ± 0.020	0.974 ± 0.025
Ours	<b>0.998 ± 0.002</b>	<b>0.971 ± 0.022</b>	<b>0.992 ± 0.010</b>	<b>0.982 ± 0.019</b>

noteworthy that, according to the Nyquist-Shannon sampling theorem (Shannon, 2006), the ultimate attainable resolution is limited to twice the voxel size of the protein density map, corresponding to 2.46 Å and 2.62 Å for EMPIAR-10049 and EMPIAR-10028, respectively. On these two datasets, GEM achieves resolutions that approach this theoretical upper bound, highlighting its strong ability to recover fine structural details from cryo-EM images.

### 4.3 LOCAL RESOLUTION ESTIMATION

We then evaluate the reconstructions from all methods using the standard protocol of local resolution estimation. While the FSC provides a single global resolution value, the same principle can be localized to regions within the protein density, providing more detailed estimation of the quality of the reconstruction. Specifically, given two half maps reconstructed from independent data splits, a cubic subregion is extracted at the same location from each map. The FSC between the two cubes is then computed following Equation 9, and the local resolution is determined from the frequency at which the FSC curve falls below the 0.143 threshold. This estimated resolution is assigned to the voxel at the cube center. Repeating this process across the entire map yields a resolution map that depicts the spatial variability of the reconstruction quality.

As shown in Figure 6 and Table 3, GEM consistently achieves superior local resolution across all four datasets. The EMPIAR-10005 dataset poses a significant challenge to existing approaches due to its

Table 5: Ablation study on different parameters of 3DGS. "No Rotation" means the rotation of all 3D Gaussians are fixed at 0. "Isotropic Scale" indicates the scale matrix  $\mathbf{S}_i$  is identical across the three axis of the 3D Gaussians, resulting in a spherical shape of the 3D Gaussians.

	EMPIAR-10028	EMPIAR-10049	EMPIAR-10076	EMPIAR-10005
① No Rotation	3.71 Å	4.23 Å	5.99 Å	4.15 Å
② Isotropic Scaling	4.02 Å	5.53 Å	12.80 Å	4.21 Å
① + ②	10.67 Å	6.81 Å	13.54 Å	7.95 Å
Original	<b>2.98 Å</b>	<b>2.56 Å</b>	<b>2.62 Å</b>	<b>2.43 Å</b>

high symmetry and structural complexity, leading CryoSPARC and CryoDRGN to achieve average local resolutions of approximately 5 Å and 3 Å, respectively. In contrast, GEM attains an average local resolution of 2.28 Å, which approaches the Nyquist-Shannon limit of 2.43 Å, thereby demonstrating its strong capability to recover fine structural details.

#### 4.4 FOURIER SLICE CORRELATION

We further validate the reconstructions using a recently proposed evaluation protocol, Fourier slice correlation (FSLC, Huang et al. (2024)). FSLC serves as an intermediate-level metric that quantifies directional resolution anisotropy by measuring the similarity between corresponding slices extracted from two reconstructed protein densities (Tan et al., 2017; Huang et al., 2024). In practice, slices are sampled at uniformly distributed elevation angles ( $0^\circ - 180^\circ$ ) and azimuth angles ( $0^\circ - 360^\circ$ ).

As shown in Figure 7 and Table 4, GEM consistently achieves higher FSLC values, indicating stronger agreement across a wide range of rotation angles. While all methods exhibit relatively lower correlation around the azimuthal regions of the elevation axis, GEM still outperforms the baselines, demonstrating superior robustness in challenging orientations. Moreover, our method attains the lowest standard deviation across the four datasets, further highlighting its stability and consistency with respect to directional variations.

#### 4.5 DESIGN CHOICES OF 3DGS FOR CRYO-EM

While 3DGS was developed for reconstruction of natural scenes, it is unclear which set of parameters are essential for cryo-EM reconstruction. We therefore ablate per-Gaussian rotation and anisotropic scaling. Concretely, we evaluate: (i) No Rotation, fix  $\mathbf{R}_i = \mathbf{I}$ ; (ii) Isotropic Scaling, enforce  $\mathbf{S}_i = s_i \mathbf{I}$ , where  $s_i$  is the distance of the  $i$ -th 3D Gaussian to its nearest neighbor; (iii) No Rotation and Isotropic Scaling (both constraints).

As shown in Table 5 (lower Å is better), ① all ablations degrade resolution across datasets. ② Enforcing isotropic scaling is consistently more harmful than removing rotation (e.g., on EMPIAR-10076, 2.62 Å to 12.80 Å vs. 2.62 Å to 5.99 Å), and combining both constraints cause the most harm to the performance. These results indicate that both per-Gaussian rotation and anisotropic scaling are critical for accurately modeling cryo-EM densities, validating our choice to retain the full 3DGS parameterization in GEM.

## 5 CONCLUSION

In this work, we introduced GEM, a novel framework that leverages 3D Gaussian Splatting (3DGS) for cryo-EM reconstruction. By explicitly representing protein densities with 3D Gaussians, GEM avoids the information loss of Fourier-space methods and the cubic memory and computation overhead of NeRF-based real-space approaches, thereby achieving high resolution and efficiency. Through comprehensive experiments, GEM shows that it delivers up to  $48\times$  faster training,  $12\times$  lower memory consumption, and up to 38.8% improvement in local resolution over prior state-of-the-art methods, often approaching the physical resolution limit of the microscope. Overall, GEM establishes 3DGS as a practical and scalable paradigm for cryo-EM, unifying efficiency with high-resolution accuracy and offering broad benefits to the structural biology community.

486 ETHICS STATEMENT  
487

488 We strictly follow the ICLR Code of Ethics. This work relies solely on existing open-source cryo-EM  
489 datasets and does not involve any biological experiments. As such, it does not generate any harmful  
490 insights or biological outcomes.  
491

492 REPRODUCIBILITY STATEMENT  
493

494 Reproducibility is a core principle of this work. All datasets used are publicly available and can be  
495 accessed through their official repositories. The code, along with detailed instructions for its use, will  
496 be released as open source upon acceptance to ensure full reproducibility of our results.  
497

498 THE USE OF LARGE LANGUAGE MODELS (LLMs)  
499

500 To enhance clarity and readability, we employed OpenAI’s GPT-5 and GPT-5-thinking models exclu-  
501 sively as language polishing tools. Their role was limited to proofreading, grammatical correction,  
502 and stylistic refinement—functions comparable to those of conventional grammar checkers and  
503 dictionaries. These tools did not contribute any new scientific content or ideas, and their usage is  
504 consistent with standard practices in manuscript preparation.  
505

506 REFERENCES  
507

508 Press release: The 2017 Nobel Prize in Chemistry.  
509

510 Paul D Adams, Pavel V Afonine, Gábor Bunkóczi, Vincent B Chen, Ian W Davis, Nathaniel Echols,  
511 Jeffrey J Headd, L-W Hung, Gary J Kapral, Ralf W Grosse-Kunstleve, et al. Phenix: a compre-  
512 hensive python-based system for macromolecular structure solution. *Biological crystallography*, 66  
513 (2):213–221, 2010.  
514

515 Xiao-Chen Bai, Greg McMullan, and Sjors HW Scheres. How cryo-em is revolutionizing structural  
516 biology. *Trends in biochemical sciences*, 40(1):49–57, 2015a.  
517

518 Yang Bai, Daniel B Müller, Girish Srinivas, Ruben Garrido-Oter, Eva Potthoff, Matthias Rott, Nina  
519 Dombrowski, Philipp C Münch, Stijn Spaepen, Mitja Remus-Emsermann, et al. Functional overlap  
520 of the arabidopsis leaf and root microbiota. *Nature*, 528(7582):364–369, 2015b.  
521

522 Ronald N Bracewell. Strip integration in radio astronomy. *Australian Journal of Physics*, 9(2):  
523 198–217, 1956.

524 Joseph H Davis, Yong Zi Tan, Bridget Carragher, Clinton S Potter, Dmitry Lyumkis, and James R  
525 Williamson. Modular assembly of the bacterial large ribosomal subunit. *Cell*, 167(6):1610–1622,  
526 2016.  
527

528 Dominika Elmlund and Hans Elmlund. Cryogenic electron microscopy and single-particle analysis.  
529 *Annual review of biochemistry*, 84(1):499–517, 2015.  
530

531 Ben Fei, Jingyi Xu, Rui Zhang, Qingyuan Zhou, Weidong Yang, and Ying He. 3d gaussian splatting  
532 as new era: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 2024.

533 Daissy H Garces, William T Rhodes, and Nestor M Peña. Projection-slice theorem: a compact  
534 notation. *Journal of the Optical Society of America A*, 28(5):766–769, 2011.  
535

536 Robert M Glaeser, Eva Nogales, and Wah Chiu. *Single-particle Cryo-EM of biological macro-*  
537 *molecules*. IOP publishing, 2021.  
538

539 George Harauz and Marin van Heel. Exact filters for general geometry three dimensional reconstruc-  
tion. *Optik.*, 73(4):146–156, 1986.

- 540 David Herreros, Carlos Perez Mata, Chari Noddings, Deli Irene, James Krieger, David A Agard,  
541 Ming-Daw Tsai, Carlos Oscar Sanchez Sorzano, and Jose Maria Carazo. Real-space heterogeneous  
542 reconstruction, refinement, and disentanglement of cryo-em conformational states with hetsiren.  
543 *Nature communications*, 16(1):3751, 2025.
- 544 Yue Huang, Chengguang Zhu, Xiaokang Yang, and Manhwa Liu. High-resolution real-space recon-  
545 struction of cryo-em structures using a neural field network. *Nature Machine Intelligence*, 6(8):  
546 892–903, 2024.
- 547 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting  
548 for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.
- 549 Maofu Liao, Erhu Cao, David Julius, and Yifan Cheng. Structure of the trpv1 ion channel determined  
550 by electron cryo-microscopy. *Nature*, 504(7478):107–112, 2013.
- 551 Xinhang Liu, Yan Zeng, Yifan Qin, Hao Li, Jiakai Zhang, Lan Xu, and Jingyi Yu. Cryoformer:  
552 Continuous heterogeneous cryo-em reconstruction using transformer-based neural representations.  
553 *arXiv preprint arXiv:2303.16254*, 2023.
- 554 Yonggang Lu, Jiakuan Liu, Li Zhu, Bianlan Zhang, and Jing He. 3d reconstruction from cryo-em  
555 projection images using two spherical embeddings. *Communications Biology*, 5(1):304, 2022.
- 556 Clare D McGillem and George R Cooper. Continuous and discrete signal and system analysis. (*No*  
557 *Title*), 1991.
- 558 Iain Milne, Gordon Stephen, Micha Bayer, Peter JA Cock, Leighton Pritchard, Linda Cardle, Paul D  
559 Shaw, and David Marshall. Using tablet for visual exploration of second-generation sequencing  
560 data. *Briefings in bioinformatics*, 14(2):193–202, 2013a.
- 561 Jacqueline LS Milne, Mario J Borgnia, Alberto Bartesaghi, Erin EH Tran, Lesley A Earl, David M  
562 Schauder, Jeffrey Lengyel, Jason Pierson, Ardan Patwardhan, and Sriram Subramaniam. Cryo-  
563 electron microscopy—a primer for the non-microscopist. *The FEBS journal*, 280(1):28–45, 2013b.
- 564 Kazuyoshi Murata and Matthias Wolf. Cryo-electron microscopy for structural analysis of dynamic  
565 biological macromolecules. *Biochimica et Biophysica Acta (BBA)-General Subjects*, 1862(2):  
566 324–334, 2018.
- 567 Colin M Palmer and Christopher HS Aylett. Real space in cryo-em: the future is local. *Biological*  
568 *Crystallography*, 78(2):136–143, 2022.
- 569 Ali Punjani, John L Rubinstein, David J Fleet, and Marcus A Brubaker. cryosparc: algorithms for  
570 rapid unsupervised cryo-em structure determination. *Nature methods*, 14(3):290–296, 2017.
- 571 Huaizhi Qu, Xiao Wang, Yuanyuan Zhang, Sheng Wang, William Stafford Noble, and Tianlong Chen.  
572 Cryonerf: reconstruction of homogeneous and heterogeneous cryo-em structures using neural  
573 radiance field. *bioRxiv*, pp. 2025–01, 2025.
- 574 Johann Radon. 1.1 über die bestimmung von funktionen durch ihre integralwerte längs gewisser  
575 mannigfaltigkeiten. *Classic papers in modern diagnostic radiology*, 5(21):124, 2005.
- 576 Heng Ru, Melissa G Chambers, Tian-Min Fu, Alexander B Tong, Maofu Liao, and Hao Wu. Molecular  
577 mechanism of v (d) j recombination from synaptic rag1-rag2 complex structures. *Cell*, 163(5):  
578 1138–1152, 2015.
- 579 Sjors HW Scheres. Relion: implementation of a bayesian approach to cryo-em structure determination.  
580 *Journal of structural biology*, 180(3):519–530, 2012.
- 581 Claude E Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21,  
582 2006.
- 583 Yong Zi Tan, Philip R Baldwin, Joseph H Davis, James R Williamson, Clinton S Potter, Bridget  
584 Carragher, and Dmitry Lyumkis. Addressing preferred specimen orientation in single-particle  
585 cryo-em through tilting. *Nature methods*, 14(8):793–796, 2017.
- 586  
587  
588  
589  
590  
591  
592  
593

- 594 Guang Tang, Liwei Peng, Philip R Baldwin, Deepinder S Mann, Wen Jiang, Ian Rees, and Steven J  
595 Ludtke. Eman2: an extensible image processing suite for electron microscopy. *Journal of structural*  
596 *biology*, 157(1):38–46, 2007.
- 597  
598 Marin Van Heel and Michael Schatz. Fourier shell correlation threshold criteria. *Journal of structural*  
599 *biology*, 151(3):250–262, 2005.
- 600  
601 Wilson Wong, Xiao-chen Bai, Alan Brown, Israel S Fernandez, Eric Hanssen, Melanie Condron,  
602 Yan Hong Tan, Jake Baum, and Sjors HW Scheres. Cryo-em structure of the plasmodium  
603 falciparum 80s ribosome bound to the anti-protozoan drug emetine. *elife*, 3:e03080, 2014.
- 604  
605 Tong Wu, Yu-Jie Yuan, Ling-Xiao Zhang, Jie Yang, Yan-Pei Cao, Ling-Qi Yan, and Lin Gao. Recent  
606 advances in 3d gaussian splatting. *Computational Visual Media*, 10(4):613–642, 2024.
- 607  
608 Vickie Ye, Ruilong Li, Justin Kerr, Matias Turkulainen, Brent Yi, Zhuoyang Pan, Otto Seiskari,  
609 Jianbo Ye, Jeffrey Hu, Matthew Tancik, et al. gsplat: An open-source library for gaussian splatting.  
610 *Journal of Machine Learning Research*, 26(34):1–17, 2025.
- 611  
612 Weihao Yu, Yuanhao Cai, Ruyi Zha, Zhiwen Fan, Chenxin Li, and Yixuan Yuan. X<sup>2</sup>-gaussian:  
613 4d radiative gaussian splatting for continuous-time tomographic reconstruction. *arXiv preprint*  
614 *arXiv:2503.21779*, 2025.
- 615  
616 Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free  
617 3d gaussian splatting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*  
618 *recognition*, pp. 19447–19456, 2024.
- 619  
620 Ruyi Zha, Yanhao Zhang, and Hongdong Li. Naf: neural attenuation fields for sparse-view cbct  
621 reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted*  
622 *Intervention*, pp. 442–452. Springer, 2022.
- 623  
624 Ruyi Zha, Tao Jun Lin, Yuanhao Cai, Jiwen Cao, Yanhao Zhang, and Hongdong Li. R<sup>2</sup>-  
625 gaussian: Rectifying radiative gaussian splatting for tomographic reconstruction. *arXiv preprint*  
626 *arXiv:2405.20693*, 2024.
- 627  
628 Ellen D Zhong, Tristan Bepler, Bonnie Berger, and Joseph H Davis. Cryodrgn: reconstruction of  
629 heterogeneous cryo-em structures using neural networks. *Nature methods*, 18(2):176–185, 2021.
- 630  
631 Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. Ewa volume splatting. In  
632 *Proceedings Visualization, 2001. VIS'01.*, pp. 29–538. IEEE, 2001.
- 633  
634 Matthias Zwicker, Hanspeter Pfister, Jeroen Van Baar, and Markus Gross. Ewa splatting. *IEEE*  
635 *Transactions on Visualization and Computer Graphics*, 8(3):223–238, 2002.
- 636  
637  
638  
639  
640  
641  
642  
643  
644  
645  
646  
647

648 6 APPENDIX

649 6.1 DESCRIPTION OF BASELINE METHODS

650 **Fourier Slice Theorem for Reconstruction.** A common strategy in cryo-EM reconstruction  
 651 (Zhong et al., 2021; Punjani et al., 2017) is to leverage the Fourier slice theorem (Bracewell, 1956),  
 652 which states that the Fourier transform of a 2D projection of a 3D density  $V_i$  corresponds to a central  
 653 slice of its 3D Fourier transform perpendicular to the projection direction. Formally,  
 654

$$655 \mathcal{F}(\text{Proj}(V_i; \phi_i, t_i)) = \text{Slice}(\mathcal{F}(V_i); \phi_i, t_i), \quad (10)$$

656 where  $\mathcal{F}(\cdot)$  denotes the Fourier transform and  $\text{Slice}(\cdot; \phi_i, t_i)$  extracts the 2D slice orthogonal to the  
 657 projection direction specified by  $(\phi_i, t_i)$ . This theorem implies that reconstructing the protein density  
 658 reduces to learning to predict individual slices in Fourier space, thereby significantly lowering the  
 659 memory footprint during training.  
 660

661 **Real-Space Cryo-EM Reconstruction.** The reconstruction formulation in Equation 3 closely  
 662 parallels 3D reconstruction problems in computer vision, which has motivated recent work to perform  
 663 cryo-EM reconstruction directly in real space using neural radiance fields (NeRFs) (Huang et al., 2024;  
 664 Qu et al., 2025). Unlike Fourier-based methods, this approach avoids repeated Fourier transforms  
 665 and can in principle achieve higher resolution. However, it requires predicting the full density  $\hat{V}_i$   
 666 before supervision can be applied, since the contrast transfer function (CTF)  $C_i$  must operate on the  
 667 complete projection image due to the convolution operation.  
 668

669 Concretely, given a set of predefined coordinates  $D$ , each particle’s pose is applied as

$$670 D_i = \text{Rotate}(D; \phi_i) + t_i, \quad (11)$$

671 and the resulting coordinates are used to query the NeRF, producing the predicted density volume  $\hat{V}_i$ .  
 672 The model is then trained by projecting  $\hat{V}_i$  under the same pose and computing the loss between the  
 673 simulated projection  $\hat{I}_i$  and the experimental image  $I_i$ :  
 674

$$675 \mathcal{L}_i = \text{MSE}(I_i, \text{Proj}(\hat{V}_i; \phi_i, t_i)). \quad (12)$$

676 Because the entire density  $\hat{V}_i$  must be generated at each iteration, this method incurs cubic memory  
 677 and computation cost, making training slow and resource-intensive despite its resolution advantage.  
 678

681 6.2 DERIVATION OF 3DGS PROJECTION

682 The  $j$ -th 3D Gaussian is defined as:

$$683 G_j(\mathbf{x} | \rho_j, \mathbf{p}_j, \Sigma_j) = \rho_j \cdot \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{p}_j)^\top \Sigma_j^{-1}(\mathbf{x} - \mathbf{p}_j)\right). \quad (13)$$

684 Therefore for a query position  $\mathbf{x}$  in the protein density, its value is defined as

$$685 \hat{V}_i(\mathbf{x}) = \sum_{j=1}^M G_j(\mathbf{x} | \rho_j, \mathbf{p}_j, \Sigma_j). \quad (14)$$

686 To avoid the query of the entire protein density before the projection, we aim to explicitly derive the  
 687 expression of the projected image  $\hat{I}_i$ . Since each pixel in the image  $\hat{I}_i$  corresponds to an electron  
 688 beam  $\mathbf{r}$ , we denote the projection of this specific pixel as  $\hat{I}_i(\mathbf{r})$ .  
 689

690 Following the definition, the projection of the pixel can be defined as:

$$691 \hat{I}_i(\mathbf{r}) = \int \hat{V}_i(\mathbf{x}) dz = \int \sum_{j=1}^M G_j(\mathbf{x} | \rho_j, \mathbf{p}_j, \Sigma_j) dz \quad (15)$$

$$692 = \sum_{j=1}^M \int G_j(\mathbf{x} | \rho_j, \mathbf{p}_j, \Sigma_j) dz. \quad (16)$$

From the standard 3DGS for natural images, we have

$$\hat{I}_r(\mathbf{r}) \approx \sum_{j=1}^M \int G_j(\mathbf{x}|\rho_j, \mathbf{p}_j, \underbrace{\mathbf{J}_j \mathbf{W} \boldsymbol{\Sigma}_j \mathbf{W}^\top \mathbf{J}_j^\top}_{\tilde{\boldsymbol{\Sigma}}_i}) dz, \quad (17)$$

where  $\tilde{\boldsymbol{\Sigma}}_i$  is the new Gaussian covariance matrix controlled by local approximation matrix  $\mathbf{J}_j$  and viewing transformation matrix  $\mathbf{W}$  that only corresponds to the rotation  $\phi$ . The projection of natural images follows the pinhole camera model, while in cryo-EM the projection rays are parallel, therefore the Jacobian matrix satisfies  $\mathbf{J}_j = \mathbf{I}$ , which gives  $\tilde{\boldsymbol{\Sigma}}_j = \boldsymbol{\Sigma}_j$ . Thus following (Zwicker et al., 2002; 2001) we have

$$\hat{I}_i(\mathbf{r}) = \sum_{j=1}^M \rho_j (2\pi)^{\frac{3}{2}} |\tilde{\boldsymbol{\Sigma}}_j|^{\frac{1}{2}} \int \frac{1}{(2\pi)^{\frac{3}{2}} |\tilde{\boldsymbol{\Sigma}}_j|^{\frac{1}{2}}} \exp\left(-\frac{1}{2} (\mathbf{x} - \tilde{\mathbf{p}}_j)^\top \tilde{\boldsymbol{\Sigma}}_j^{-1} (\mathbf{x} - \tilde{\mathbf{p}}_j)\right) dz \quad (18)$$

$$= \sum_{j=1}^M \rho_j (2\pi)^{\frac{3}{2}} |\boldsymbol{\Sigma}_j|^{\frac{1}{2}} \int \frac{1}{(2\pi)^{\frac{3}{2}} |\boldsymbol{\Sigma}_j|^{\frac{1}{2}}} \exp\left(-\frac{1}{2} (\mathbf{x} - \mathbf{p}_j)^\top \boldsymbol{\Sigma}_j^{-1} (\mathbf{x} - \mathbf{p}_j)\right) dz \quad (19)$$

$$= \sum_{j=1}^M \rho_j \frac{1}{2\pi |\hat{\boldsymbol{\Sigma}}_j|^{1/2}} \exp\left(-\frac{1}{2} (\hat{\mathbf{x}} - \hat{\mathbf{p}}_j)^\top \hat{\boldsymbol{\Sigma}}_j^{-1} (\hat{\mathbf{x}} - \hat{\mathbf{p}}_j)\right). \quad (20)$$

### 6.3 DESCRIPTION OF DATASETS

The characteristics of datasets used in this paper are summarized as follows:

- **EMPIAR-10005** (TRPV1): A well-characterized protein in the vertebrate TRP family, and is frequently used to investigate fundamental TRP channel functions and structures.
- **EMPIAR-10028** (*Pf* 80S ribosome): Contains a complex composition of the antibiotic and *Pf* ribosome and is commonly used for evaluation of the reconstruction resolution.
- **EMPIAR-10049** (Synaptic RAG1–RAG2 complex): Exhibits substantial compositional and conformational heterogeneity and is therefore challenging for high-resolution reconstruction.
- **EMPIAR-10076** (L17-depleted 50S ribosomal intermediates): Comprises complex intermediate assembly states that can cause reconstruction quality degradation.