

---

# DRACO: A Denoising-Reconstruction Autoencoder for Cryo-EM

---

Yingjun Shen<sup>1,2\*</sup> Haizhao Dai<sup>1,2\*</sup> Qihe Chen<sup>1,2</sup> Yan Zeng<sup>1,2</sup>

Jiakai Zhang<sup>1,2</sup> Yuan Pei<sup>1,3</sup> Jingyi Yu<sup>1</sup>

<sup>1</sup>School of Information Science and Technology, ShanghaiTech University.

<sup>2</sup>Cellverse Co, Ltd.

<sup>3</sup>Human Institute, ShanghaiTech University.

{shenyj2022,daihzh2023}@shanghaitech.edu.cn

{chenqh2024,zengyan2024,zhangjk,peiyuan,yujingyi}@shanghaitech.edu.cn

## Abstract

Foundation models in computer vision have demonstrated exceptional performance in zero-shot and few-shot tasks by extracting multi-purpose features from large-scale datasets through self-supervised pre-training methods. However, these models often overlook the severe corruption in cryogenic electron microscopy (cryo-EM) images by high-level noises. We introduce DRACO, a Denoising-Reconstruction Autoencoder for Cryo-EM, inspired by the Noise2Noise (N2N) approach. By processing cryo-EM movies into odd and even images and treating them as independent noisy observations, we apply a denoising-reconstruction hybrid training scheme. We mask both images to create denoising and reconstruction tasks. For DRACO’s pre-training, the quality of the dataset is essential, we hence build a high-quality, diverse dataset from an uncurated public database, including over 270,000 movies or micrographs. After pre-training, DRACO naturally serves as a generalizable cryo-EM image denoiser and a foundation model for various cryo-EM downstream tasks. DRACO demonstrates the best performance in denoising, micrograph curation, and particle picking tasks compared to state-of-the-art baselines.

## 1 Introduction

Foundation models in computer vision have demonstrated remarkable capabilities in zero-shot and few-shot tasks. These models learn to extract multi-purpose visual features from large-scale, diverse datasets through text-guided [1, 2, 3] or self-supervised [4, 5] pre-training methods such as masked image modeling (MIM) [6]. The features can then be applied to various downstream tasks. For instance, DINOv2 [5] is trained on a large-scale curated dataset and shows significant performance improvements in classification, retrieval, segmentation, etc. The success of vision foundation models has stimulated advances across various scientific disciplines. Due to the diverse modalities of scientific imaging, training domain-specific foundation models [7, 8, 9, 10] is essential to meet specific demands. For example, the UNI [7] foundation model for tissue imaging is pre-trained on 100 million images for 34 representative clinical downstream tasks.

In structural biology, cryogenic electron microscopy (cryo-EM) stands as a pivotal bio-imaging technique [11]. Unlike optical imaging methods, cryo-EM possesses several distinctive characteristics: first, cryo-EM utilizes high-energy electron beams as its illumination source [12] and direct detector

---

\*The authors contributed equally to this work.

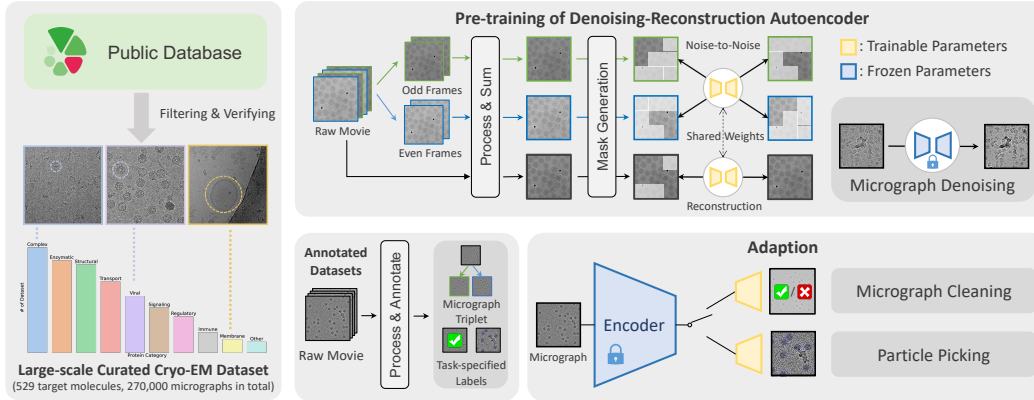


Figure 1: **Overview of DRACO.** For pre-training, we construct a large-scale curated dataset containing 529 types of protein data with over 270,000 cryo-EM movies or micrographs. Based on this, we present DRACO, a denoising-reconstruction autoencoder for cryo-EM. A pre-trained DRACO naturally serves as a generalizable cryo-EM image denoiser and a foundation for various downstream model adaptations such as micrograph curation and particle picking.

device (DDD) captures a continuous multi-frame sequence, often called a *movie* [13]. To mitigate specimen damage during exposure, the electron dose per frame is restricted, leading to *extremely low signal-to-noise ratios (SNR)* in the captured images. Second, cryo-EM employs motion correction to counteract blurring induced by specimen drift during exposure to ultimately obtain a sharper single micrograph [13]. Last, the acquired images comprise hundreds of thousands of target protein particles with diverse poses. To resolve 3D structures, researchers utilize a pipeline composed of multiple downstream tasks including micrograph curation, micrograph denoising, particle picking, pose estimation, and ultimately high-resolution reconstruction.

In line with denoising autoencoders [14], which consider “robustness to partial destruction of the input” a good representation criterion, existing self-supervised learning methods, such as masked autoencoders (MAE) [6], have been successful in learning expressive representations by reconstructing the missing patches from a partially masked image. However, in cryo-EM, these methods overlook the severe corruption caused by pixel-level random noise, leading to degraded performance. To be more robust to noises, DMAE [15] reconstructs the clean image from the masked one that is further corrupted by synthetic Gaussian noise. Nevertheless, cryo-EM clean reference images are impossible to obtain due to the fragile biological specimens, posing a significant challenge.

In this paper, we present DRACO, a Denoising-Reconstruction Autoencoder for Cryo-EM, as shown in Figure 1. Inspired by Noise2Noise (N2N) [16], which learns to denoise images using only paired noisy images, we divide the original movie into two sub-movies based on odd and even frame numbers, processing them into odd and even images. We treat them as two independent, noisy observations of the underlying true signal, thus the idea of N2N can apply. During training, we partially mask both images, creating masked regions and unmasked regions corresponding to denoising and reconstruction tasks: in the unmasked region, the odd noisy patch learns to recover the even noisy patch, and vice versa. In the masked region, we introduce relatively low-noise images from the complete movie to supervise the reconstructed results. This denoising-reconstruction hybrid training scheme achieves the robust feature extraction of noisy cryo-EM images.

The quality of training samples is crucial for the general-purpose feature extraction of cryo-EM images. Direct access to the public database leads to varying data quality, inconsistent data formats, or missing annotations. Therefore, we construct a large-scale, high-quality, and diverse single-particle cryo-EM image dataset by curating and manually processing 529 sets of data from EMPIAR [17], obtaining over 270,000 cryo-EM movies or micrographs in total. After pre-training, DRACO naturally serves as a generalizable cryo-EM image denoiser and a foundation for various downstream model adaptations. We hence explore the performance of DRACO on three downstream tasks: micrograph curation, denoising, and particle picking. Extensive experiments show that DRACO outperforms the state-of-the-art baselines in all downstream tasks. We will release code, pre-trained/fine-tuned models, and the large-scale curated dataset.

## 2 Related Work

Our work aims to extend the vision foundation model to the field of cryo-EM. We therefore only discuss the most relevant works in respective fields.

**Vision foundation models in computer vision.** Vision foundation models are pre-trained on large-scale image datasets [18, 19] using self-supervised learning methods [20], aimed at extracting general visual signals rapidly adaptable to various downstream visual tasks [21, 22, 23, 24]. Techniques for pre-training vision foundation models, such as contrastive learning [25, 26, 1, 27] and self-distillation [28, 5], focus on aligning features across different models or modalities, while another method, masked image modeling [6, 29, 30, 4], reconstructs features from masked images to capture high-level visual semantics. However, these existing vision foundation models are not directly applicable to cryo-EM imaging. In particular, their application to cryo-EM imaging is limited by the high noise levels in micrographs, which degrade signal capture. Therefore, we propose a denoising-reconstruction pre-training framework that is robust to highly noisy cryo-EM micrographs, making it suitable for specific downstream tasks in cryo-EM.

**Vision foundation models in life science.** The remarkable success of foundation models has extended to various life science imaging domains, including applications in retinal [10], fluorescence microscopy [9], histopathology [7, 8], and radiology imaging [31]. These models have shown considerable effectiveness in tasks such as disease diagnosis, lesion detection, and image restoration within these fields. In contrast to these domains, which benefit from extensive and well-curated datasets supporting model training, the cryo-EM field lacks such resources. To fill this gap, we have developed a well-curated, large-scale dataset specifically designed to support the training of cryo-EM foundation models, ensuring that they can be effectively applied in this specialized field.

**Cryo-EM image denoising.** To tackle the issues of low SNR and complex noise patterns in cryo-EM images, traditional denoising techniques often employ noise models like the Poisson-Gaussian model [12, 32] and rely on filtering methods [33, 34, 35] to denoise. However, these methods oversimplify the noise patterns, which can lead to the loss of high-frequency signal details. Recently, NT2C [36] uses Generative Adversarial Network to learn the noise patterns for denoising, but it requires simulated datasets as the clean references. Another series of learning-based methods [37, 38] make full use of multi-frame data by generating odd and even images for denoising based on Noise2Noise (N2N) [16, 39] framework. These methods do not require clean images for denoising, but they still suffer from small-scale datasets and network architectures, which limit their generalizability. In this paper, we propose DRACO, pre-trained on a large-scale curated dataset, that can naturally serve as a generalizable denoiser for cryo-EM micrographs.

**Downstream tasks in single particle analysis.** An effective foundational model can benefit downstream tasks in cryo-EM, including micrograph curation and particle picking. Micrograph curation aims to ensure that only high-quality images are selected for further analysis, yet current methods rely heavily on manual inspection [40, 41]. Particle picking involves identifying and extracting representative particles from micrographs, which is a critical task in the cryo-EM single particle analysis (SPA) reconstruction pipeline. Traditional methods [42, 43, 44], such as template matching [45, 46] and difference of Gaussians (DoG) method [47], heavily rely on prior information and require substantial *ad hoc* post-processing. Learning-based models [48, 49, 50], such as Topaz-Picking [51], crYOLO [52], and CryoTransformer [53], offer more streamlined processes but still face challenges in generalizability due to the limited data scale. Our DRACO, pre-trained on large-scale cryo-EM image datasets, can effectively adapt to these tasks and demonstrate strong generalization capabilities.

## 3 Preliminary: Imaging Formation Model

Cryo-EM uses a Direct Detector Device (DDD) camera [32] for their notably higher detective quantum efficiency (DQE) compared to traditional cameras. This allows recording the micrograph as a multi-frame movie rather than a single integrated exposure. In this setup, a movie is a series of continuous multi-frame images, denoted as  $\mathcal{M} = \{\hat{I}_i\}_{i=1}^M$ , where each frame  $\hat{I}_i$  is an independent observation of the true signal  $I$ .

Ideally, the imaging process in cryo-EM involves two main steps: 1) projecting the 3D density volume of the region of interest  $V(x, y, z) : \mathbb{R}^3 \rightarrow \mathbb{R}$  along the  $z$ -axis via weak-phase object approximation [54], and 2) modulating the projection image with the Point Spread Function (PSF)  $g$  of the cryo-EM

optical lens, expressed as:

$$I = g * \int V(x, y, z) dz. \quad (1)$$

where  $I$  is considered as the true signal.

However, in practice, the captured frames suffer from extremely low SNR due to the limited electron dosage and the high sensitivity of DDD. The main noise source is Poisson (shot) noise from the detector, denoted as  $\text{Poisson}(I)$ , arising from the inherent uncertainty of the electron measurement [12]. We assume that the additional noise types like heat, readout, and dark current noise are collectively modeled as a zero-mean Gaussian noise  $\mathcal{G}$  with an unknown variance  $\sigma^2$  [42]:

$$\hat{I} = \text{Poisson}(I) + \mathcal{G}. \quad (2)$$

As the number of observations increases, the average of these movie frames converges to the true signal:

$$\mathbb{E}[\hat{I}] = I \approx \frac{1}{M} \sum_{i=1}^M \hat{I}_i. \quad (3)$$

We define the image noise  $\epsilon$  as the difference between the captured and clean images for each frame. Thus, the expectation (mean value) of noise distribution is:

$$\mathbb{E}[\epsilon] = \mathbb{E}[\hat{I}_i - I] = 0. \quad (4)$$

Thus, we derive that the cryo-EM image noise is zero-mean. This conclusion is also aligned with existing cryo-EM reconstruction methods [42, 41], which directly assume that the noise distribution is an additive zero-mean Gaussian yet can still achieve high-resolution reconstruction. We are well aware that this derivation is relatively trivial compared to an actual theoretical analysis [12], but this gives us intuitive guidance to integrate the N2N [16] idea: learning to denoise images from solely noisy image pairs, into our pre-training framework.

**Movie to micrograph triplets.** Given an  $M$ -frame movie  $\mathcal{M}$ , we divide it into odd frames  $\mathcal{M}^o = \{I_{2i-1}\}_{i=1}^{\lceil M/2 \rceil}$  and even frames  $\mathcal{M}^e = \{I_{2i}\}_{i=1}^{\lfloor M/2 \rfloor}$ . An off-the-shelf motion correction method [13] is then applied to correct cross-frame drifts in  $\mathcal{M}$ ,  $\mathcal{M}^o$ , and  $\mathcal{M}^e$ . By summing up the frames within each subset, we generate three micrographs with the same shape: the original micrograph  $\hat{\mathbf{I}}$ , odd micrograph  $\hat{\mathbf{I}}^o$ , and even micrograph  $\hat{\mathbf{I}}^e$ . As aforementioned, all these micrographs are expected to reflect the true signal  $I$  but corrupted by noises.

## 4 Denoising-reconstruction Autoencoder

We introduce DRACO, a denoising-reconstruction autoencoder for cryo-EM, as illustrated in Figure 2. Different from existing masked imaging modeling methods, our model uses paired odd-even micrographs as inputs for the denoising target on visible patches. Further, we utilize the original micrographs as the additional supervision signal for reconstruction on masked patches.

**Masking.** Following the standard scheme in Vision Transformer (ViT) [55], each micrograph in a triplet, consisting of one original, one odd, and one even micrograph from the same movie, is divided into regular non-overlapping patches. We create patch sets  $\{\mathbf{x}_i^o\}_{i=1}^N$  for the odd,  $\{\mathbf{x}_i^e\}_{i=1}^N$  for the even, and  $\{\mathbf{x}_i\}_{i=1}^N$  for the original micrograph, where  $N$  represents the number of patches. For the odd and the even micrographs used as inputs in our model, we generate two sets of binary masks,  $\{\mathbf{m}_i^o\}_{i=1}^N$  and  $\{\mathbf{m}_i^e\}_{i=1}^N$ , with a mask ratio  $\gamma$ . Here,  $\mathbf{m}_i = 1$  means the  $i$ -th patch is masked, and 0 means unmasked. Additionally, we ensure that a patch can be 1) only visible in one of them, or 2) masked in both. This strategy ensures that each visible patch has no information sharing of its corresponding patch on the other micrograph. Notably, this requires  $\gamma \geq 0.5$  for each input micrograph.

**Network architecture.** For DRACO’s pre-training, we employ a ViT-based encoder-decoder architecture following the MAE framework [6]. Positional embeddings are first added to the input patches, which are then masked to select only visible patches, denoted as  $\{\mathbf{v}_i\}_{i=1}^{\gamma N}$ . The encoder,  $G_{\text{enc}}$ , is a ViT that transforms these visible patches from either odd or even micrographs into latent features. To




Figure 2: **The pipeline of DRACO.** Given a pair of partially masked odd and even micrographs, the encoder takes odd-visible patches and even-visible patches as inputs. The unmasked latent patches are combined with masked latent patches together to generate the latent representation  $\mathbf{z}_i$ . Then the latent representation passes through the decoder to generate predicted patches. The N2N loss is applied to odd-visible predicted patches with corresponding even input patches, and vice versa. The reconstruction loss is applied to both invisible predicted patches with higher SNR input patches.

align with the original unmasked size, zeros are padded to the latent features based on the positions indicated by the corresponding masks before the encoder outputs them:

$$\{\mathbf{z}_i\}_{i=1}^N = G_{\text{enc}}\left(\{\mathbf{v}_i^o\}_{i=1}^{\gamma N}, \{\mathbf{m}_i^o\}_{i=1}^N; \theta_{\text{enc}}\right), \quad \{\mathbf{z}_i^e\}_{i=1}^N = G_{\text{enc}}\left(\{\mathbf{v}_i^e\}_{i=1}^{\gamma N}, \{\mathbf{m}_i^e\}_{i=1}^N; \theta_{\text{enc}}\right). \quad (5)$$

Each of the latent features from odd and even micrographs retain part of the original micrograph's information. To reconstruct the micrograph, we generate the latent representation  $\{\mathbf{z}_i\}_{i=1}^N$  for further processing by the decoder:

$$\mathbf{z}_i = (1 - \mathbf{m}_i^o) \cdot \mathbf{z}_i^o + (1 - \mathbf{m}_i^e) \cdot \mathbf{z}_i^e + \mathbf{m}_i^o \cdot \mathbf{m}_i^e \cdot [\text{MASK}], \quad (6)$$

where the [MASK] token is a shared learnable embedding representing the masked patch to be predicted. Finally, the decoder  $G_{\text{dec}}$  takes the latent representation  $\{\mathbf{z}_i\}_{i=1}^N$  with another positional embedding added as input and predicts all masked and visible patches to reconstruct the complete micrograph:

$$\{\mathbf{y}_i\}_{i=1}^N = G_{\text{dec}}(\{\mathbf{z}_i\}_{i=1}^N; \theta_{\text{dec}}). \quad (7)$$

**Denosing target.** Inspired by N2N and Topaz-Denoise [37], which predict denoised images only from paired noisy images, we introduce an image-denoising target on visible patches. For any patch  $\mathbf{x}_i$  visible in only the odd micrographs, the model predicts its counterpart in the even micrographs, and vice versa, as illustrated in Figure 2. As mentioned earlier, the expectation (mean) values of the odd and even micrographs are the true signal. Therefore, following Topaz-Denoise, we employ a patch-wise L2 loss function for each visible patch, aiming for regressing the true signal:

$$\mathcal{L}_{\text{N2N}}(\mathbf{x}_i^o, \mathbf{x}_i^e, \mathbf{y}_i) = \begin{cases} \|\mathbf{x}_i^e - \mathbf{y}_i\|_2^2, & \text{if } \mathbf{m}_i^o = 0, \\ \|\mathbf{x}_i^o - \mathbf{y}_i\|_2^2, & \text{if } \mathbf{m}_i^e = 0. \end{cases} \quad (8)$$

**Reconstruction target.** For masked patches, we let the decoder predict the pixel value of patches from the original micrograph with higher SNR compared to odd and even micrographs for better reconstruction quality. Thus, the reconstruction loss is:

$$\mathcal{L}_{\text{recon}}(\mathbf{x}_i, \mathbf{y}_i) = \|\mathbf{x}_i - \mathbf{y}_i\|_2^2, \quad \text{if } \mathbf{m}_i^o \cdot \mathbf{m}_i^e = 1. \quad (9)$$

**Training objective.** During training, we combine the N2N loss with the reconstruction loss:

$$\mathcal{L} = \mathcal{L}_{\text{N2N}} + \lambda \mathcal{L}_{\text{recon}}, \quad (10)$$

where  $\lambda$  is a hyper-parameter set to 1.0 in all our experiments.

## 5 Experiments

**Large-scale curated dataset for pre-training.** The effectiveness and robustness of DRACO depend heavily on the quality and scale of the cryo-EM pre-training dataset. However, direct access to public databases like the Electron Microscopy Public Image Archive (EMPIAR) [17] results in variations in data quality, inconsistent data formats, and inaccurate or even missing annotations. To overcome these challenges, we have developed a data generation workflow. First, we select datasets with reported resolutions better than 10 Å, ensuring high-quality data acquisition. Next, we collect the raw data, including metadata, movies, and micrographs, from pre-defined high-quality datasets available on EMPIAR. Finally, we re-process the raw data using cryoSPARC [41] through a custom processing pipeline designed to exclude low-quality micrographs and movies, generate annotations for downstream tasks, and verify the resolutions of the reconstructed results. This workflow has allowed us to compile a large-scale, curated cryo-EM dataset containing over 270,000 raw micrographs and more than 50,000 raw movies from 529 verified single-particle cryo-EM datasets, occupying approximately 25 TB of disk storage in total. Details of the cryoSPARC processing pipeline are provided in Appendix C.1.

**Data augmentation.** Each micrograph in a triplet goes through the same data augmentation process: randomly cropping to between 1/16 and 1/4 of the original size (typically  $4096 \times 4096$ ), resizing to  $256 \times 256$ , applying random horizontal and vertical flips, and finally normalizing based on the mean and standard deviation computed from the original micrograph. We randomly crop each micrograph 16 times within a single epoch for fully utilization.

**Pre-training details.** We explore the performance of DRACO using two ViT architectures for the encoder, ViT-B and ViT-L, denoted as **DRACO-B** and **DRACO-L**, respectively. The decoder of DRACO uses 8 Transformer blocks with embedding dimension 512, followed by a three-layer convolution neck and a linear projection layer with an output dimension  $16 \times 16$ , which is also the patch size of the input. The mask ratio for the one input micrograph is 0.75 by default. To fully utilize our large-scale curated dataset, we warm up DRACO on the original 270,000 micrographs based on the MAE training scheme for 200 epochs. Then we adopt our novel denoising-reconstruction pre-training for 400 epochs. The warm-up stage takes 6 hours, and the pre-training stage takes 16 hours on a GPU cluster with 64 NVIDIA A800 GPUs, requiring approximately 80 GB of memory for a batch size of 4096.

### 5.1 Particle Picking

Particle picking aims to accurately locate particles in highly noisy micrographs, which is directly related to the resolution of the final reconstructed result. For adaption, we conduct supervised fine-tuning based on the ViT-based object detection framework Detectron2 [21] on our curated dataset. The results show that DRACO is capable of accurately detecting particles of various shapes and sizes across three challenging datasets.

**Dataset.** To create the annotated dataset for adaptation, we employ the standard workflow of cryoSPARC to generate a high-quality annotated dataset containing over 80,000 full micrographs and approximately 8 million particles across 46 types of protein. Detailed descriptions of the annotation workflow can be found in Appendix C.2. The test dataset includes three full micrograph sets with EMPIAR ID 1) 10081: Human HCN1 channel protein [56], the protein structure has a similar shape to ice, which could lead to false positive picking, 2) 10350: LetB transport protein [57], this kind of protein tends to aggregate together, posing challenges in accurate picking in crowded area, and 3) 10407: 70S ribosome [58], the micrographs are in the extremely low SNR.

**Baseline and metrics.** We compare DRACO with existing state-of-the-art learning-based methods for generalized particle picking, including **Topaz** [51], **crYOLO** [52], and **CryoTransformer** [53]. For ViT-based baselines, including **MAE**, **DRACO-B** and **DRACO-L** as aforementioned. We integrate them into Detectron2 framework by loading pre-trained weights of DRACO’s encoder into Detectron2’s encoder. Additionally, to show the effectiveness of pre-trained encoder, we compare with **Detectron2** trained from scratch. More configuration details can be found in Appendix D.1. We evaluate baselines in terms of conventional metrics including precision, recall, and F1 score, and the resolution of the resolved 3D structure from the picked particles, which is also a crucial metric. We process the particles selected by each method with the standard cryoSPARC workflow and finally




Figure 3: **Visualization of particle picking results.** We show the picking results of DRACO and baselines on the test datasets range from small transport proteins to huge ribosomes. Blue, red, and yellow circles denote true positives, false positives, and false negatives, respectively.

Table 1: **Particle picking results.** We report the precision, recall, F1 score, and resolution on each test dataset among all baselines. The resolution is obtained from the default cryoSPARC workflow. We compare DRACO with existing state-of-the-art methods and consistently achieve the best F1 score and resolution of reconstructed results.

Method	Human HCNI				70S ribosome				LetB			
	Precision ( $\uparrow$ )	Recall ( $\uparrow$ )	F1 score ( $\uparrow$ )	Res. ( $\downarrow$ )	Precision	Recall	F1 score	Res.	Precision	Recall	F1 score	Res.
Topaz	0.462	<b>0.956</b>	0.623	4.20	0.362	<b>0.943</b>	0.523	2.80	0.518	0.761	0.617	3.67
crYOLO	<b>0.818</b>	0.748	0.782	4.15	0.602	0.869	0.711	2.78	0.632	0.163	0.224	4.62
CryoTransformer	0.475	0.910	0.624	4.13	0.517	0.887	0.654	2.79	0.429	0.706	0.534	3.67
Detectron	0.392	0.834	0.533	4.50	0.668	0.901	0.767	2.85	0.589	0.804	0.680	3.86
MAE	0.703	0.649	0.675	4.32	0.712	0.876	0.786	2.84	0.591	<b>0.805</b>	0.682	4.03
<b>DRACO-B</b>	0.768	0.799	0.793	4.03	0.732	0.905	0.810	2.61	0.637	0.779	0.701	3.55
<b>DRACO-L</b>	0.830	0.802	<b>0.816</b>	<b>3.90</b>	<b>0.803</b>	0.846	<b>0.824</b>	<b>2.51</b>	<b>0.678</b>	0.780	<b>0.725</b>	<b>3.53</b>

produce 3D reconstruction density maps and the corresponding resolution, as described in Appendix C.3.

**Results.** As illustrated in Figure 3, both Topaz and CryoTransformer tend to pick a larger number of particles, but this often results in many false positives. In contrast, crYOLO achieves higher precision in picking, yet exhibits a higher number of false negatives. Detectron2 trained from scratch and pre-trained MAE both have difficulties in distinguishing signal and noise, leading to a lack of generalizability. In contrast, DRACO effectively identifies correct particles, surpassing the performance of all baselines, as demonstrated in Table 1.

## 5.2 Micrograph Denoising

Once pre-trained, our model can naturally serve as a generalizable denoiser by directly predicting every patch of input noisy micrograph without any further fine-tuning.

**Baseline and metrics.** To evaluate the effectiveness of DRACO on the denoising task, we first compare DRACO with the standard MAE trained on the 270,000 micrographs from our large-scale curated dataset with ViT-B as the backbone, denoted as **MAE**. We further compare with a popular denoising method **Topaz-Denoise** [37] in cryo-EM. For a fair comparison, we train Topaz-Denoise on our odd-even micrograph dataset for 100 epochs with default settings. Last, we compare with the traditional method **Low-pass** filtering that has already been integrated into commercial software cryoSPARC [41]. We utilize the same protocol used in cryoSPARC and set the low-pass cutoff resolution to 20 Å. As cryo-EM micrographs lack clean ground truth, following Topaz-Denoise, we employ an SNR calculation method that involves 20 manually annotated signal-background region pairs as references. For each  $i$ -th pair, we calculate the mean and variance for both the signal region

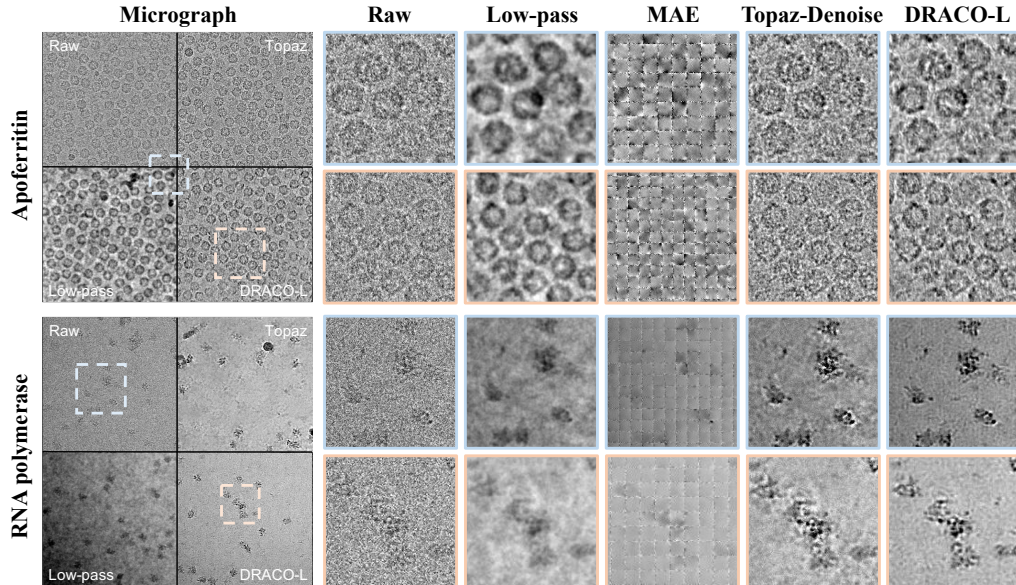


Figure 4: **Qualitative comparison results of micrograph denoising.** We visualize the denoising results of DRACO and state-of-the-art baselines. Our results show the most significant SNR improvement without the loss of the particle structure details. In contrast, Low-pass leads to a severe blur on particles, MAE introduces severe patch-wise artifacts and Topaz only shows either minor SNR improvements or blurred results.

Table 2: **Quantitative comparison of denoising results.** We report the SNR calculated with Equation 11. On our test dataset, DRACO outperforms all other baselines.

	Human Apoferritin	HA Trimer	Phage MS2	RNA polymerase
Raw	-10.01	-6.52	-12.52	-4.69
Low-pass	-2.18	-0.84	-6.71	3.09
Topaz-Denoise	-5.67	-0.83	-6.93	8.66
MAE	-0.31	-1.85	-8.45	1.27
<b>DRACO-B</b>	1.92	<b>3.69</b>	-0.13	10.13
<b>DRACO-L</b>	<b>2.01</b>	3.33	<b>0.23</b>	<b>12.21</b>

$r_i^s$  and the background region  $r_i^b$ , yielding  $\mu_i^s, v_i^s$  for the signal, and  $\mu_i^b, v_i^b$  for the background. The average SNR is then computed in dB as follows:

$$\text{SNR} = \frac{10}{N} \sum_{i=1}^N \log_{10} \frac{(\mu_i^s - \mu_i^b)^2}{v_i^b} \quad (11)$$

**Dataset.** To evaluate the denoising capabilities of the baseline models, we select four original micrograph datasets as test datasets, which are *excluded* from the training set. These sets are Human Apoferritin (EMPIAR-10421) [59], HA Trimer (EMPIAR-10096) [60], Phage MS2 (EMPIAR-10075) [61], and RNA polymerase (EMPIAR-11521) [62]. For each dataset, 5 micrographs are selected and 20 signal-background pair regions are labeled in total. Specifically, these signal and background regions are chosen close together to ensure similar background signals across both regions.

**Results.** The quantitative experiments show that DRACO achieves significant performance improvements in terms of SNR after denoising compared with state-of-the-art methods, as shown in Table 2. In Figure 4, the standard MAE can only recover smooth contours of the particle with severe artifacts. Low-pass filtering smooths both signal and background noise, but the background noise is still relatively high and the structure information of particles is corrupted. Topaz sometimes fails to effectively denoise micrographs but over-smooth them instead, which affects the generalizability. In contrast, DRACO outperforms all baselines in retaining the original particle signals with the lowest background noises, showing the best generalizability. We additionally reconstruct 3D density map and corresponding resolution using the denoised particles generated by each method. To ensure



Table 3: **Quantitative comparison of reconstruction using denoised particles.** We report the resolution obtained from the standard cryoSPARC workflow. On our test dataset, DRACO outperforms all other baselines.

	Human Apoferritin	HA Trimer	Phage MS2	RNA polymerase
Low-pass	2.63	<b>2.06</b>	3.46	2.75
Topaz-Denoise	2.34	3.06	2.52	2.93
MAE	2.77	2.15	3.78	2.81
<b>DRACO-B</b>	<b>2.05</b>	2.10	<b>2.51</b>	<b>2.56</b>

Table 4: **Quantitative comparison of micrograph curation.** Miffi employs its own general model, while ResNet18 is trained from scratch. DRACO reports the best results on all four classification metrics.

Method	Accuracy	Precision	Recall	F1 score
Miffi	0.836	0.899	0.845	0.871
ResNet18	0.938	0.923	0.960	0.940
MAE	0.904	0.927	0.892	0.909
<b>DRACO-B</b>	0.963	<b>0.976</b>	0.953	0.964
<b>DRACO-L</b>	<b>0.983</b>	<b>0.976</b>	<b>0.992</b>	<b>0.984</b>

that comparisons reflect only the impact of denoising quality, we fix the locations and poses of the picked particles, which were determined using the cryoSPARC workflow. The results are shown in Table 3. DRACO consistently achieves the highest resolution in most cases, demonstrating its effectiveness in preserving more high-frequency signals while effectively reducing background noise. We demonstrate additional denoising results in Figure 5.

### 5.3 Micrograph Curation

A modern cryo-EM can capture thousands of micrographs in a day. However, the quality of captured micrographs is unverified. Low-quality micrographs may arise from artifacts such as empty sample, ice crystals, ethane contamination, severe drifting, etc. [40]. Low-quality micrographs can negatively contribute to the final reconstruction results. A reliable automated micrograph curation method can significantly improve the efficiency of the data processing pipeline, resulting in shorter processing time and improved final resolution. We show that DRACO can easily adapt to this 2-class classification task by linear probing, and achieving the best performance compared to the state-of-the-art method. Similar to [6], we freeze the encoder backbone and train an extra linear classification head.

**Dataset.** We manually annotate 1,194 micrographs from original micrograph datasets, assigning a binary label (accept or reject) to each to indicate quality. The dataset comprises 617 high-quality and 577 low-quality micrographs. We divided these micrographs into training and evaluation datasets using an 80%/20% split ratio.

**Baseline and metric.** We compare **DRACO-B** and **DRACO-L** with linear probing against several baselines: a small **ResNet18** [63]; an existing supervised method **Miffi** [40], which has been trained on 45,000 annotated data; and the standard **MAE** with linear probing both pre-trained and adapted on our datasets. The ResNet18 is trained from scratch to show the effectiveness of our pre-training strategy. More configuration details is provided in Appendix D.2. We report the widely used metrics for classification including precision, recall, F1 score, and accuracy on our test dataset. As Miffi predicts multi-labels on low-quality micrographs, we consider them all as rejections for our metric calculations.

**Results.** As shown in Table 4, Miffi, limited by insufficient training data, lacks generalizability on the test dataset. Both ResNet and MAE face difficulties in effectively separating noise from signal for classification for accurate classification. In contrast, DRACO extracts global information from noisy images more effectively, resulting in higher classification accuracy and demonstrating better generalizability compared to other methods.

### 5.4 Ablation study

We compare the performance of networks with different parameter sizes in particle picking and micrograph curation tasks, as shown in Table 1 and Table 4. The results demonstrate that our method

Table 5: **Evaluation of mask ratios.** We demonstrate the performance of DRACO with different mask ratios. The result shows that at the 0.75 mask ratio, DRACO achieves the best performance.

Mask Ratio	Micrograph Curation				Denoising	
	Accuracy	Precision	Recall	F1 Score	RNA polymerase	HA Trimer
0.5	0.954	0.968	0.945	0.956	10.29	2.57
0.625	0.930	0.960	0.909	0.933	<b>10.49</b>	2.80
0.75	<b>0.963</b>	0.976	<b>0.953</b>	<b>0.964</b>	10.13	<b>3.69</b>
0.875	0.958	<b>0.984</b>	0.939	0.961	9.59	2.28

Table 6: **Evaluation of loss function.** We demonstrate the performance of DRACO with different training objective on the 70S ribosome dataset. The result shows that with both loss, DRACO achieves the best performance.

Training scheme	Particle Picking				Denoising
	Precision( $\uparrow$ )	Recall( $\uparrow$ )	F1 Score( $\uparrow$ )	Res.( $\downarrow$ )	SNR( $\uparrow$ )
<b>DRACO-B w/o N2N</b>	0.712	0.876	0.786	2.84	-4.94
<b>DRACO-B w/o recon</b>	0.713	0.817	0.761	2.85	-4.22
<b>DRACO-B</b>	<b>0.732</b>	<b>0.905</b>	<b>0.810</b>	<b>2.61</b>	<b>-2.86</b>

can effectively scale up. Additionally, we evaluate the impact of different mask ratios on denoising and micrograph curation performance, as shown in Table 5. DRACO achieves the highest SNR and curation metric at a 0.75 mask ratio, thus we choose 0.75 as our default mask ratio. We have also conducted an additional ablation study of loss design. Specifically, we remove either the N2N loss (**w/o N2N**) or the reconstruction loss (**w/o recon**) from the training and evaluate the resulting models on particle picking and denoising tasks, as shown in Table 6. The result shows that both N2N and reconstruction losses improve performance. Without the N2N loss, DRACO struggles to distinguish between signal and noise in micrographs. Without the reconstruction loss, DRACO loses its ability to extract general features. We show additional visualization of their denoising results (Figure 6) in the Appendix A.

## 6 Discussion

**Limitations.** As the first attempt to achieve robust feature extraction for cryo-EM via a novel denoising-reconstruction autoencoder, our work presents opportunities for future enhancements. First, our method relies heavily on the performance of motion correction algorithms. This can be improved by designing a more comprehensive denoising task for raw noisy movies. Second, although we have collected what we believe to be the largest curated dataset for cryo-EM, it focuses primarily on mainstream single-particle datasets. To enhance dataset diversity, other types of cryo-EM datasets, such as cryo-electron tomography (cryo-ET)[64] datasets, should also be included. Finally, our current approach only supports various micrograph-level downstream tasks. For particle-level tasks, such as pose estimation, a more fine-grained yet robust feature extraction is required. This can be achieved by developing a particle-level version of DRACO.

**Conclusion.** We have introduced DRACO, a foundation model designed specifically for cryo-EM image processing, supported by a unique denoising-reconstruction pre-training framework to enable robust feature extraction for cryo-EM micrographs. We have constructed a diverse and high-quality cryo-EM image dataset from the uncuration public database, comprising over 270,000 movies and micrographs. After pre-training, our model’s versatility is evidenced by its superior adaptation performance across multiple downstream tasks, including denoising, micrograph curation, and particle picking. All code, pre-trained model weights, and datasets will be made publicly available for further research and model development.

## 7 Acknowledgement

This work was supported by HPC Platform of ShanghaiTech University. We thank Zhenyang Xu for shaping the paper.

## References

- [1] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *ArXiv*, abs/2204.06125, 2022.
- [2] Yanghao Li, Haoqi Fan, Ronghang Hu, Christoph Feichtenhofer, and Kaiming He. Scaling language-image pre-training via masking. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 23390–23400, 2022.
- [3] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *Trans. Mach. Learn. Res.*, 2022, 2022.
- [4] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *International Conference on Learning Representations (ICLR)*, 2022.
- [5] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Q. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba, Russ Howes, Po-Yao (Bernie) Huang, Shang-Wen Li, Ishan Misra, Michael G. Rabbat, Vasu Sharma, Gabriel Synnaeve, Huijiao Xu, Hervé Jégou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision. *ArXiv*, abs/2304.07193, 2023.
- [6] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15979–15988, 2021.
- [7] Richard J. Chen, Tong Ding, Ming Y. Lu, Drew F. K. Williamson, Guillaume Jaume, Andrew H. Song, Bowen Chen, Andrew Zhang, Daniel Shao, Muhammad Shaban, Mane Williams, Lukas Oldenburg, Luca L. Weishaupt, Judy J. Wang, Anurag Vaidya, Long Phi Le, Georg Gerber, Sharifa Sahai, Walt Williams, and Faisal Mahmood. Towards a general-purpose foundation model for computational pathology. *Nature medicine*, 2024.
- [8] Ming Y. Lu, Bowen Chen, Drew F. K. Williamson, Richard J. Chen, Ivy Liang, Tong Ding, Guillaume Jaume, Igor Odintsov, Long Phi Le, Georg Gerber, Anil V. Parwani, Andrew Zhang, and Faisal Mahmood. A visual-language foundation model for computational pathology. *Nature medicine*, 2024.
- [9] Chenxi Ma, Weimin Tan, Ruian He, and Bo Yan. Pretraining a foundation model for generalizable fluorescence microscopy-based image restoration. *Nature Methods*, pages 1–10, 2024.
- [10] Yukun Zhou, Mark A Chia, Siegfried K. Wagner, Murat S. Ayhan, Dominic J. Williamson, Robbert R. Struyven, Timing Liu, Moucheng Xu, Mateo G. Lozano, Peter Woodward-Court, et al. A foundation model for generalizable disease detection from retinal images. *Nature*, 622:156 – 163, 2023.
- [11] Werner Kühlbrandt. The resolution revolution. *Science*, 343(6178):1443–1444, 2014.
- [12] Miloš Vulović, Raimond B.G. Ravelli, Lucas J. van Vliet, Abraham J. Koster, Ivan Lazić, Uwe Lücken, Hans Rullgård, Ozan Öktem, and Bernd Rieger. Image formation modeling in cryo-electron microscopy. *Journal of Structural Biology*, 183(1):19–32, 2013.
- [13] Shawn Q Zheng, Eugene Palovcak, Jean-Paul Armache, Kliment A Verba, Yifan Cheng, and David A Agard. Motioncor2: anisotropic correction of beam-induced motion for improved cryo-electron microscopy. *Nature methods*, 14(4):331–332, 2017.
- [14] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103, 2008.
- [15] QuanLin Wu, Hang Ye, Yuntian Gu, Huishuai Zhang, Liwei Wang, and Di He. Denoising masked autoencoders help robust classification. *arXiv preprint arXiv:2210.06983*, 2022.
- [16] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2noise: Learning image restoration without clean data. *ArXiv*, abs/1803.04189, 2018.
- [17] Andrii Iudin, Paul K. Korir, Sriram Somasundharam, Simone Weyand, Cesare Cattavittello, Néli José da Fonseca, Osman Salih, Gerard J. Kleywegt, and Ardan Patwardhan. Empiar: the electron microscopy public image archive. *Nucleic Acids Research*, 51:D1503 – D1511, 2022.
- [18] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Kumar Gupta. Revisiting unreasonable effectiveness of data in deep learning era. pages 843–852, 2017.
- [19] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. Laion-5b: An open large-scale dataset for training next generation image-text models. *ArXiv*, abs/2210.08402, 2022.

- [20] Randall Balestriero, Mark Ibrahim, Vlad Sobal, Ari S. Morcos, Shashank Shekhar, Tom Goldstein, Florian Bordes, Adrien Bardes, Grégoire Mialon, Yuandong Tian, Avi Schwarzschild, Andrew Gordon Wilson, Jonas Geiping, Quentin Garrido, Pierre Fernandez, Amir Bar, Hamed Pirsiavash, Yann LeCun, and Micah Goldblum. A cookbook of self-supervised learning. *ArXiv*, abs/2304.12210, 2023.
- [21] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. In *European Conference on Computer Vision*, pages 280–296. Springer, 2022.
- [22] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross B. Girshick. Segment anything. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3992–4003, 2023.
- [23] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jérôme Revaud. Dust3r: Geometric 3d vision made easy. *ArXiv*, abs/2312.14132, 2023.
- [24] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *ArXiv*, abs/2108.07258, 2021.
- [25] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. *ArXiv*, abs/2002.05709, 2020.
- [26] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *ArXiv*, abs/1807.03748, 2018.
- [27] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, and Furu Wei. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *ArXiv*, abs/2208.10442, 2022.
- [28] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9630–9640, 2021.
- [29] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *ArXiv*, abs/2106.08254, 2021.
- [30] Zhiliang Peng, Li Dong, Hangbo Bao, Qixiang Ye, and Furu Wei. Beit v2: Masked image modeling with vector-quantized visual tokenizers. *ArXiv*, abs/2208.06366, 2022.
- [31] Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. Towards generalist foundation model for radiology. *ArXiv*, abs/2308.02463, 2023.
- [32] Ludwig Reimer and Helmut Kohl. Transmission electron microscopy. 2008.
- [33] Xinrui Huang, Sha Li, and Song Gao. Applying a modified wavelet shrinkage filter to improve cryo-electron microscopy imaging. *Journal of computational biology : a journal of computational molecular cell biology*, 25 9:1050–1058, 2018.
- [34] Wen Jiang, Matthew L. Baker, Qiu Wu, Chandrajit L. Bajaj, and Wah Chiu. Applications of a bilateral denoising filter in biological electron microscopy. *Journal of structural biology*, 144 1-2:114–22, 2003.
- [35] Dai Wei and Chang cheng Yin. An optimized locally adaptive non-local means denoising filter for cryo-electron microscopy data. *Journal of structural biology*, 172 3:211–8, 2010.
- [36] Hongjia Li, Hui Zhang, Xiaohua Wan, Zhidong Yang, Chengmin Li, Jintao Li, Renmin Han, Ping Zhu, and Fa Zhang. Noise-transfer2clean: denoising cryo-em images based on noise modeling and transfer. *Bioinformatics*, 38(7):2022–2029, 2022.
- [37] Tristan Bepler, Kotaro Kelley, Alex J Noble, and Bonnie Berger. Topaz-denoise: general deep denoising models for cryoem and cryoet. *Nature communications*, 11(1):5208, 2020.
- [38] Tim-Oliver Buchholz, Mareike Jordan, Gaia Pigo, and Florian Jug. Cryo-care: Content-aware image restoration for cryo-transmission electron microscopy data. In *2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019)*, pages 502–506. IEEE, 2019.
- [39] Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2void - learning denoising from single noisy images. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2124–2132, 2018.
- [40] Da Xu and Nozomi Ando. Miffi: Improving the accuracy of cnn-based cryo-em micrograph filtering with fine-tuning and fourier space information. *Journal of Structural Biology*, page 108072, 2024.
- [41] Ali Punjani, John L Rubinstein, David J Fleet, and Marcus A Brubaker. cryosparc: algorithms for rapid unsupervised cryo-em structure determination. *Nature methods*, 14(3):290–296, 2017.
- [42] Jasenko Zivanov, Takanori Nakane, Björn O Forsberg, Dari Kimanius, Wim JH Hagen, Erik Lindahl, and Sjors HW Scheres. New tools for automated high-resolution cryo-em structure determination in relion-3. *elife*, 7:e42166, 2018.

- [43] Guang Tang, Liwei Peng, Philip R Baldwin, Deepinder S Mann, Wen Jiang, Ian Rees, and Steven J Ludtke. Eman2: an extensible image processing suite for electron microscopy. *Journal of structural biology*, 157(1):38–46, 2007.
- [44] Dimitry Tegunov and Patrick Cramer. Real-time cryo-electron microscopy data preprocessing with warp. *Nature methods*, 16(11):1146–1152, 2019.
- [45] AM Roseman. Findem—a fast, efficient program for automatic selection of particles from electron micrographs. *Journal of structural biology*, 145(1-2):91–99, 2004.
- [46] K Zhang, M Li, and F Sun. Gautomatch: an efficient and convenient gpu-based automatic particle selection program. *Unpublished manuscript*, 2011.
- [47] NR Voss, CK Yoshioka, M Radermacher, CS Potter, and B Carragher. Dog picker and tiltpicker: software tools to facilitate particle selection in single particle electron microscopy. *Journal of structural biology*, 166(2):205–213, 2009.
- [48] Feng Wang, Huichao Gong, Gaochao Liu, Meijing Li, Chuangye Yan, Tian Xia, Xueming Li, and Jianyang Zeng. Deeppicker: A deep learning approach for fully automated particle picking in cryo-em. *Journal of structural biology*, 195(3):325–336, 2016.
- [49] Yanan Zhu, Qi Ouyang, and Youdong Mao. A deep convolutional neural network approach to single-particle recognition in cryo-electron microscopy. *BMC bioinformatics*, 18:1–10, 2017.
- [50] Chentianye Xu, Xueying Zhan, and Min Xu. Cryomae: Few-shot cryo-em particle picking with masked autoencoders. *arXiv preprint arXiv:2404.10178*, 2024.
- [51] Tristan Bepler, Andrew Morin, Micah Rapp, Julia Brasch, Lawrence Shapiro, Alex J Noble, and Bonnie Berger. Positive-unlabeled convolutional neural networks for particle picking in cryo-electron micrographs. *Nature methods*, 16(11):1153–1160, 2019.
- [52] Thorsten Wagner, Felipe Merino, Markus Stabrin, Toshio Moriya, Claudia Antoni, Amir Apelbaum, Philine Hagel, Oleg Sitsel, Tobias Raisch, Daniel Prumbaum, et al. Sphire-cryolo is a fast and accurate fully automated particle picker for cryo-em. *Communications biology*, 2(1):218, 2019.
- [53] Ashwin Dhakal, Rajan Gyawali, Ligu Wang, and Jianlin Cheng. Cryotransformer: a transformer model for picking protein particles from cryo-em micrographs. *Bioinformatics*, 40, 2023.
- [54] Miloš Vulović, Lenard M Voortman, Lucas J van Vliet, and Bernd Rieger. When to use the projection assumption and the weak-phase object approximation in phase contrast cryo-em. *Ultramicroscopy*, 136:61–66, 2014.
- [55] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ArXiv*, abs/2010.11929, 2020.
- [56] Chia-Hsueh Lee and Roderick MacKinnon. Structures of the human hcn1 hyperpolarization-activated channel. *Cell*, 168(1):111–120, 2017.
- [57] Georgia L Isom, Nicolas Coudray, Mark R MacRae, Collin T McManus, Damian C Ekiert, and Gira Bhabha. Letb structure reveals a tunnel for lipid transport across the bacterial envelope. *Cell*, 181(3):653–664, 2020.
- [58] David Nicholson, Thomas A Edwards, Alex J O’Neill, and Neil A Ranson. Structure of the 70s ribosome from the human pathogen acinetobacter baumannii in complex with clinically relevant antibiotics. *Structure*, 28(10):1087–1100, 2020.
- [59] Yong Zi Tan and John L Rubinstein. Through-grid wicking enables high-speed cryoem specimen preparation. *Acta Crystallographica Section D: Structural Biology*, 76(11):1092–1103, 2020.
- [60] Yong Zi Tan, Philip R Baldwin, Joseph H Davis, James R Williamson, Clinton S Potter, Bridget Carragher, and Dmitry Lyumkis. Addressing preferred specimen orientation in single-particle cryo-em through tilting. *Nature methods*, 14(8):793–796, 2017.
- [61] Roman I Koning, Josue Gomez-Blanco, Inara Akopjana, Javier Vargas, Andris Kazaks, Kaspars Tars, José María Carazo, and Abraham J Koster. Asymmetric cryo-em reconstruction of phage ms2 reveals genome structure in situ. *Nature communications*, 7(1):12524, 2016.
- [62] Andreas U Mueller, James Chen, Mengyu Wu, Courtney Chiu, B Tracy Nixon, Elizabeth A Campbell, and Seth A Darst. A general mechanism for transcription bubble nucleation in bacteria. *Proceedings of the National Academy of Sciences*, 120(14):e2220874120, 2023.
- [63] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [64] EV Orlova and Helen R Saibil. Structural analysis of macromolecular assemblies by electron microscopy. *Chemical reviews*, 111(12):7710–7748, 2011.

- [65] Florian KM Schur, Martin Obr, Wim JH Hagen, William Wan, Arjen J Jakobi, Joanna M Kirkpatrick, Carsten Sachse, Hans-Georg Kräusslich, and John AG Briggs. An atomic model of hiv-1 capsid-sp1 reveals structures regulating assembly and maturation. *Science*, 353(6298):506–508, 2016.
- [66] Emdb—the electron microscopy data bank. *Nucleic Acids Research*, 52(D1):D456–D465, 2024.
- [67] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9992–10002, 2021.

## A Additional Results of DRACO

We demonstrate that DRACO achieves the highest visual denoising quality in terms of both signal preservation and noise removal, as shown in Figure 5. This figure serves as an extension to Figure 4 in our main paper. Furthermore, we visualize the denoising results from our ablation study on different mask ratios in Figure 6. We observe that at a 0.75 mask ratio, DRACO achieves the best denoising results, as supported by Table 2 in our main paper. Lastly, we visualize the reconstruction ability at a 0.75 mask ratio in Figure 7. DRACO demonstrates comparable reconstruction capability to MAE, while significantly better denoising results on visible patches.




Figure 5: **Additional denoising results.** We have conducted additional experiments on datasets of membrane proteins and bacteriophages. DRACO achieves the highest visual denoising quality by optimally balancing signal preservation and noise reduction.




Figure 6: **Visualization of DRACO's results on different mask ratios.** At a 0.75 mask ratio, DRACO achieves the best trade-off between signal preservation and background noise removal.




Figure 7: **Additional results on image reconstruction.** We present the reconstruction results at various image resolutions while maintaining a consistent mask ratio of 0.75. DRACO demonstrates enhanced detail preservation on the visible patches compared to MAE.

## B Zero-shot Capability on Cryo-ET

Though DRACO has not been pre-trained on cryo-ET datasets, we found that DRACO can be directly applied on cryo-ET tilt series. Here, we demonstrate that DRACO is capable of denoising the unseen HIV tilt series [65], as shown in Figure 8. Specifically, we evaluate DRACO on both the tilt series and the volume slices, showing that DRACO effectively removes background noise and achieves higher contrast in the volume. Furthermore, we assess DRACO’s performance in the context of reconstruction. We compare the slices from both tomograms reconstructed using original and denoised tilt series. Additionally, we compare these with the denoised slice from the original tomogram. The result shows that DRACO can improve the contrast of slices before or after the reconstruction.






Figure 8: **Denoising cryo-ET HIV tilt series with DRACO.** Figure (a) and (b) show the HIV tilt series before and after DRACO’s denoising process. Using IMOD, we reconstruct 3D volumes of HIV from both the original and denoised series, showing their slice in Figures (c) and (d). Note the horizontal stripes in these images, which are artifacts due to the missing wedge issue in cryo-ET. Figure (e) shows a denoised slice from Figure (c) by DRACO.

## C Workflow of CryoSPARC

### C.1 Pre-training Dataset Details

EMPIAR [17] is a public archive for storing raw cryo-EM images and 3D reconstructions from vEM and XT experiments. It currently contains over 2,000 entries, totaling more than 2 PB of data. The EMDB [66] is an archive of 3D reconstructions derived from cryo-EM experiments, many of which supplement the experimental information of the EMDB-related EMPIAR entries, such as sample preparation and reconstruction processes.

In the field of 3D electron microscopy, cryo-EM micrographs undergo complex image processing to achieve 3D reconstructions with specified resolutions. Experts can model protein molecules accurately on 3D reconstructions with resolutions better than 3 Å. We define the quality of datasets from the perspective of structural biology; hence, high-quality datasets should produce high-resolution 3D reconstructions suitable for detailed analysis.

When constructing the pre-training dataset, we utilize the REST API provided by EMPIAR to obtain metadata for each dataset, such as experiment type, EMDB ID, and image classification. We specifically filter for EMPIAR entries that are experimentally linked to EMDB, prioritizing those with reconstruction resolutions better than 10 Å. Subsequently, we collect as many datasets as possible that contain single-frame micrographs. For datasets that include multiple frames, we further process them by separating the frames into odd and even micrographs.

Cryo-EM raw data consists of multi-frame recordings known as movies, which capture the number of electrons. After motion correction [13], these movies yield single frames referred to as micrographs. We download as many single-frame micrographs as possible, categorized by image type. Additionally,

we process multi-frame movies into single-frame micrographs using cryoSPARC’s Patch Motion Correction. For each multi-frame movie, we separately process the complete frames, odd frames, and even frames to generate three types of single-frame micrographs for DRACO learning.

## C.2 Particle Picking Dataset Details.

We filter and generate a dataset comprising approximately 80,000 single-frame micrographs annotated with about 8 million particles, following a process based on the cryo-EM single particle analysis reconstruction pipeline software, cryoSPARC [41]. Each EMPIAR public dataset [17] comes with a solved 3D density map, available on EMDDB [66]. Using the “create template” step in cryoSPARC, we project this map into 50 diverse poses to generate high-quality templates for template picking. Subsequent rounds of the “2D classification” step are employed to eliminate potential false positives. Finally, using these particles, we reconstruct results whose resolution did not differ by more than 20% from the reported resolution. This method allowed us to collect a high-quality annotated particle dataset.

## C.3 3D Reconstruction Pipeline

The reconstruction process is also based on cryoSPARC. After picking the particles, the standard reconstruction workflow consists of 2D classification, ab initio reconstruction and homogeneous refinement. 2D classification aims to remove any false positives in picked particles. Ab initio reconstruction can create an initial 3D model from a certain set of particles. Based on this initial model, homogeneous refinement can achieve a high-resolution result. The final resolution is determined by the Fourier shell correlation (FSC) curve. The specific method involves dividing the particles into two random halves, each undergoing homogeneous reconstruction. After reconstruction, we perform a cross-correlation on each Fourier shell in the frequency domain of two reconstructed 3D density maps. The final resolution is determined using the standard threshold of 0.143 on the FSC curve.

# D Downstream tasks settings

## D.1 Particle Picking Settings

**Particle picking baselines.** We use the Topaz [51] general model with its “resnet16u64” backbone for our baseline, picking particles that score higher than 0.0 as the final results. We use the crYOLO [52] general model for our baseline from its official website. Similarly, we use CryoTransformer’s [53] open-source general model on Github, choosing particles with scores ranging from the 25th to the 100th percentile.

**Detectron2 configurations.** For particle picking, we employ the Faster R-CNN framework within Detectron2, to fit the particle picking task with our curated dataset. The configurations for both ViT-B and ViT-L include the standard feature pyramid and window attention [67]. We also set the non-maximum suppression threshold of the region proposal network to 0.6 and adjust the pooling size of the box pooler in the region of interest network to 14. Given that the particles are mostly square, the aspect ratio of the anchors is fixed to 1.0. All other settings remain at their default values. The data augmentation goes the same process as described in the pre-training stage. We fine-tune Detectron2-based particle picking model on 64 NVIDIA A800 GPUs for 100 epochs with a batch size of 256, requiring approximately 9 hours and consuming around 100GB of memory. After fine-tuning, we process the test dataset and pick particles with scores higher than 0.1 as the final results.

## D.2 Micrograph Curation Settings

For micrograph curation, each ViT-based model undergoes a linear probing phase, while a ResNet18 [63] is trained from scratch. We employ the miffi\_v1 [40] general model of Miffi to inference on test datasets. All the model trains 50 epochs with a batch size of 128 on a single NVIDIA RTX 3090 GPU, taking about 10 minutes and utilizing around 8GB of memory.

## NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and precede the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

**The checklist answers are an integral part of your paper submission.** They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- **Delete this instruction block, but keep the section heading "NeurIPS paper checklist",**
- **Keep the checklist subsection headings, questions/answers and guidelines below.**
- **Do not modify the questions and only use the provided macros for your answers.**

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We believe that the abstract and introduction accurately reflect the development of our model, its training methodology, its effectiveness on various downstream tasks, and the creation of a curated dataset to support this model.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We have listed out the limitations of our work in Discussion.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We have provided the image and the noise model use in Section 3.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We have clearly described the data curation pipeline, data preprocessing, network architecture, and adaptations to downstream tasks in Experiments and Appendix A.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

## 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We do not provide code and data in submission. However, the code and data will be released upon acceptance.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.

- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We have stated the pre-training details in Section 5 and the adaptation details in Appendix B.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: We do not provide error bars or other appropriate information about the statistical significance of the experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We have stated the pre-training computer resources in Experiments and the adaptation computer resources in Appendix B.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

#### 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: Our research are all conducted in line with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

#### 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We have discussed the positive and negative societal impacts in Introduction and Discussion.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

#### 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We do not describe safeguards in our paper, but we will consider to prevent our model and data from misuse when we release our code and dataset.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All of our code and datasets we built upon are in public with proper licenses.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

## 13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: All of code and datasets we built upon are in public with proper licenses.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

## 14. Crowdsourcing and Research with Human Subjects



Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: We do not have crowdsourcing experiments in our paper.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

**15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: We do not have crowdsourcing experiments in our paper.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.