

CryoDECO: Deconstructing Compositional and Conformational Heterogeneity in Cryo-EM with Foundation Model Priors

Yang Yan^{1,2,6}, Yanwanyu Xi^{1,2,6}, Shiqi Fan^{1,2}, Ziyun Tang^{3,4,5}, Fajie Yuan^{2,*}, and Huaizong Shen^{3,4,5,*}

¹Zhejiang University, Hangzhou, Zhejiang, China

²School of Engineering, Westlake University, Hangzhou, Zhejiang, China

³Zhejiang Key Laboratory of Structural Biology, School of Life Sciences, Westlake University; Hangzhou, Zhejiang, China

⁴Westlake Laboratory of Life Sciences and Biomedicine, Hangzhou, Zhejiang, China

⁵Institute of Biology, Westlake Institute for Advanced Study, Hangzhou, Zhejiang, China

⁶These authors contributed equally to this work.

*Correspondence: F. Yuan (yuanfajie@westlake.edu.cn), H. Shen (shenhuaizong@westlake.edu.cn).

Abstract

Resolving compositional and conformational heterogeneity remains a fundamental bottleneck in single-particle cryo-EM. This challenge stems from a circular dependency: classification requires reliable references, while reference generation requires accurate classification. Current deep learning methods often resort to blind stochastic initialization, frequently becoming trapped in local minima within complex optimization landscapes. To address this, we present CryoDECO, a framework that breaks this deadlock by integrating representation priors from the pretrained cryo-EM foundation model, Cryo-IEF. By projecting particle images onto a semantically structured manifold, CryoDECO bypasses the *tabula rasa* search phase of standard *ab initio* reconstruction. We demonstrate that this prior-informed strategy robustly disentangles extreme compositional heterogeneity, successfully resolving 100 distinct structures from a single simulated mixture while simultaneously mapping complex conformational dynamics. CryoDECO offers a robust, high-throughput solution for systems structural biology by replacing *ab initio* reconstruction with prior-guided optimization.

Main

Cryo-electron microscopy (cryo-EM) has fundamentally transformed structural biology, moving beyond the static determination of purified macromolecules to offer a direct window into the dynamic structural landscape within complex biological samples^{1,2}. The standard Single-Particle Analysis (SPA) workflow reconstructs 3D density maps by averaging millions of 2D projection images^{3–5}; however, this relies on a strict, often unrealistic assumption of sample homogeneity. In biological reality, structural heterogeneity is ubiquitous, manifesting in two distinct but equally challenging forms: compositional heterogeneity, arising from the presence of distinct macromolecular complexes, and conformational heterogeneity, driven by the intrinsic flexibility and continuous dynamics of biological machinery. Conventional pipelines frequently treat these variations as noise, conflating distinct states into blurred consensus maps that obscure true biological mechanisms.

For emerging frontiers such as time-resolved studies⁶ and the analysis of native cell extracts^{7–9}, heterogeneity is not a technical artifact but a rich data signal. In these scenarios, the objective shifts from simple averaging to “Shotgun Cryo-EM”. Here, we redefine this term as a high-throughput structural biology paradigm: rather than relying on prior biochemical purification, we utilize *in silico* sorting to simultaneously identify, classify, and reconstruct distinct macromolecular structures directly from complex mixtures—ranging from crude cell lysates to complex environmental samples. Realizing this potential, however, presents a formidable optimization challenge known as the *ab initio* circular dependency. Accurate particle classification requires 3D structural references, but

generating those references requires the particles to be classified first. While recent deep learning approaches attempt to resolve this by jointly estimating poses and latent conformations, they operate as *tabula rasa* learners^{13,15}. Initialized from random distributions, these models must simultaneously learn the complex physics of image formation and the underlying biological structures from scratch. This necessitates a "blind" search in a high-dimensional landscape, often resulting in entrapment in local minima and a failure to disentangle complex mixtures.

To overcome this initialization bottleneck, we introduce CryoDECO (DEconstructing COmpositional and COnformational heterogeneity). This framework fundamentally alters the reconstruction paradigm by leveraging representation priors from a pretrained foundation model, Cryo-IEF¹⁶. By treating the foundation model as a "semantic interpreter" of noisy projections, our approach extracts structurally meaningful features before the reconstruction loop begins. These learned priors provide a robust geometric baseline that bypasses the "blind" search phase, effectively breaking the circular dependency inherent in *ab initio* reconstruction and guiding the optimization toward biologically valid solutions.

Foundation Model Priors Break the Optimization Deadlock

The core strategy of CryoDECO involves injecting strong inductive bias to constrain the vast search space of structural reconstruction. We achieve this using a specialized encoder-decoder framework in which the encoder projects noisy particle images onto a semantically structured manifold derived from the Cryo-IEF foundation model (Fig. 1a and Methods)¹⁶.

To ensure a truly universal representation of biological heterogeneity, we expanded the foundation model’s training corpus more than twofold—scaling from 65 million to approximately 134 million particle images. This unprecedented scale equips the encoder with robust, generalized priors that break the optimization deadlock from the very onset. Consequently, on the benchmark CryoBench Ribosome dataset (four particle species), CryoDECO converges substantially more rapidly and resolves clearer, better-separated clusters than CryoDRGN-AI, a leading *ab initio* method (Fig. 1b,c and Extended Data Fig. 2).

To operationalize this massive prior without incurring prohibitive computational costs, we optimized the encoder backbone, transitioning from a ViT-Base to a ViT-Small architecture. Empirical benchmarking confirms that this lightweight design significantly accelerates training throughput while retaining the necessary feature extraction fidelity (Extended Data Fig. 1b). Complementing the encoder, the decoder functions as a generative neural field that synthesizes 3D density maps. We optimized this module by inserting Layer Normalization (LayerNorm) within the Multi-Layer Perceptron (MLP) blocks. This structural adjustment stabilizes gradient flow, yielding significantly reduced training loss and faster convergence compared to standard MLP architectures (Extended Data Fig. 1a).

Finally, the interface between the foundation model and the reconstruction network—the latent vector (z)—serves as a critical topological hyperparameter that must align with the complexity of the biological target. We observe that discrete compositional heterogeneity demands a high-dimensional interface (e.g., $z = 128$) to ensure

orthogonality between disjoint structural states. Conversely, simple conformational dynamics necessitate a tight bottleneck (e.g., $z = 4$) to enforce manifold smoothness, while complex conformational heterogeneity benefits from an intermediate capacity (e.g., $z = 64$) to resolve high-dimensional non-linear motions without overfitting.

Resolving Extreme Compositional Complexity

We rigorously evaluated the capacity of CryoDECO to handle extreme compositional heterogeneity using the Ribosemby and Tomotwin-100 datasets¹⁷, which contain simulated particle images from 16 and 100 different structures, respectively. This scale of heterogeneity substantially recapitulates the structural complexity of unpurified cell lysates. Benchmarking against CryoDRGN-AI and CryoDRGN2 revealed the decisive advantage of our prior-informed strategy¹⁵. In the absence of strong inductive biases, baseline methods struggled to navigate the optimization landscape, often collapsing distinct structural states into ambiguous averages. In contrast, CryoDECO effectively utilized its "structural compass" to project particles from 100 different species into sharp, well-separated clusters (Fig. 2e,f).

Quantitatively, the foundation model priors conferred a fundamental performance advantage. On the challenging Tomotwin-100 dataset, CryoDECO achieved a clustering Adjusted Rand Index (ARI) of 0.622 and an Adjusted Mutual Information (AMI) of 0.853 (Table 1). This represents a substantial margin over CryoDRGN-AI (ARI: 0.086; AMI: 0.275) and CryoDRGN2 (ARI: 0.116; AMI: 0.374), which were unable to resolve the high-cardinality mixture. Furthermore, our results empirically validate the topological hypothesis

proposed in the previous section: classification accuracy improved dramatically as the latent bottleneck was expanded, rising from 41.3% at $z = 4$ to 94.0% at $z = 128$ (Extended Data Fig. 3). This confirms that a high-dimensional interface is essential for ensuring orthogonality when resolving discrete, disjoint structures. Finally, analysis of training dynamics shows that CryoDECO not only reaches higher asymptotic k-NN scores but does so with significantly faster convergence, demonstrating that the encoder's priors effectively shortcut the initial "blind" search phase of reconstruction.

In Silico Purification of Real-World Mixtures

To demonstrate the framework's utility in realistic experimental scenarios, we applied CryoDECO to the "EM ladder" dataset (EMPIAR-11693)¹⁸, a mixture of four purified macromolecular complexes: Apoferritin, β -galactosidase, Tobacco Mosaic Virus (TMV), and PP7 virus-like particles. The benchmark method, CryoDRGN-AI, exhibited significant limitations; its encoder extracted a feature space with substantial overlap between distinct species, failing to isolate the PP7 particles entirely (Extended Data Fig. 4).

In stark contrast, CryoDECO's foundation model priors successfully disentangled the mixture, projecting particles into four distinct, well-separated clusters within the latent space (Fig. 3a). 2D averaging results of the resolved clusters of particles reveal features of homogeneous species, confirming the accuracy of the classification (Fig. 3c). By utilizing these clusters for particle subset selection, we performed independent 3D reconstruction and refinement in CryoSPARC using the model output maps as starting templates. This workflow yielded high-resolution maps for all four targets: 3.00 Å for β -gal, 2.48 Å for

ApoF, 2.83 Å for TMV, and 3.38 Å for PP7 (Fig. 3b). Crucially, these resolutions outperform those reported in the original processing of this dataset, confirming that CryoDECO effectively performs "in silico purification," sorting heterogeneous mixtures into structural populations of sufficient purity to support atomic-resolution reconstruction.

Shotgun Cryo-EM in Native Cell Extracts

CryoDECO operationalizes the concept of "Shotgun Cryo-EM". Unlike historical usage of the term involving hybrid mass-spectrometry workflows, our results establish shotgun cryo-EM as a fully image-based modality for the 'digital biochemistry' of unpurified samples. To validate this paradigm, we resolved protein communities directly from a fraction of native *Chaetomium thermophilum* lysates⁸. This dataset represents a significant challenge, requiring the identification of stable structures against a complex, endogenous background—a task previously reliant on laborious manual 2D classification.

CryoDECO successfully realized this shotgun cryo-EM workflow by deconvolving the complex mixture into four discrete clusters within the latent space (Fig. 4a). 2D averaging results of the resolved clusters of particles reveal features of homogeneous species, confirming the accuracy of the classification (Fig. 4c). Subsequent ab initio reconstruction and structure refinement recovered high-resolution maps for all four expected targets: the Pyruvate Dehydrogenase Complex core (PDHc, 3.65 Å), Fatty Acid Synthase (FAS, 4.02 Å), the Oxoglutarate Dehydrogenase Complex E2 core (OGDHc, 3.86 Å), and the pre-60S ribosomal subunit (Pre-60S, 4.01 Å). Crucially, CryoDECO not only

automated the identification process but also achieved higher resolutions for all four structures compared to the original manual study.

The necessity of our foundation model priors for enabling shotgun cryo-EM was starkly illustrated by the benchmark comparison: CryoDRGN-AI failed to disentangle the feature space, resulting in contaminated clusters and the complete loss of the OGDHc and PDHc complexes (Extended Data Fig. 6). These results confirm that the "blind" initialization of conventional methods is insufficient for unpurified biological extracts, establishing CryoDECO's prior-informed strategy as the foundational technology for the era of systems structural biology.

Mapping Conformational Heterogeneity

Beyond discrete compositional heterogeneity, CryoDECO is equally adept at modeling the continuous conformational landscapes that define protein function. Using the IgG-1D and IgG-RL datasets from CryoBench, we demonstrated that the topology of the latent manifold can be tuned to the nature of the motion (Extended Data Fig.3). A tight bottleneck ($z = 4$) was sufficient to resolve the simple one-dimensional motion of IgG-1D, visualizing a clear circular manifold (Fig. 5b,c). However, the complex, randomized flexible linker in IgG-RL required an expanded dimension ($z = 64$) to capture the non-linear high-dimensional motion (Fig. 5e,f). This capability extends to real-world datasets. For the tri-snRNP spliceosome complex, traversing the latent manifold visualized a smooth, continuous transition characterized by a counter-clockwise rotation of the head region (Fig. 6a-d).

Similarly, analysis of the $\alpha V\beta 8$ integrin dataset captured the flexibility of the protein's extended arm (Fig. 6e-h).

Discussion

Current cryo-EM analysis pipelines operate predominantly as signal processing tasks, relying on iterative alignment from random initializations. In this work, we demonstrate that the initialization strategy from a "blind" random start is a primary bottleneck for resolving heterogeneity. CryoDECO represents a paradigm shift by redefining reconstruction as a semantic interpretation problem. By injecting inductive biases from the Cryo-IEF foundation model, we bridge the gap between the physics of image formation and the biological semantics of macromolecular structures. This "prior-informed" strategy does not merely accelerate convergence; it fundamentally alters the optimization landscape, allowing the algorithm to bypass the circular dependency that entraps traditional *ab initio* learners.

The empirical success on the Tomotwin-100 dataset and the native *Chaetomium thermophilum* extract underscores the transformative potential of CryoDECO for shotgun cryo-EM. The ability to resolve 100 distinct structures without ground-truth poses suggests that computational capacity need no longer be the limiting factor in analyzing complex mixtures. This capability effectively enables "in silico purification," allowing structural biologists to relax stringent biochemical purification requirements. By shifting the burden of separation from the wet lab to the computational pipeline, CryoDECO opens new

avenues for studying structural landscapes in near-native environments, capturing transient interactions and rare states that are often lost during traditional purification.

Despite these advances, challenges remain. Currently, the topology of the latent manifold—specifically the dimensionality of the latent interface (z)—acts as a critical hyperparameter that requires user selection based on the expected complexity of the sample. While discrete compositional heterogeneity benefits from high-dimensional orthogonality, continuous conformational motions require tighter bottlenecks to enforce smoothness. Future iterations of CryoDECO will focus on integrating adaptive manifold topology estimation. By automating the selection of latent dimensions, we aim to develop a fully self-configuring framework that can dynamically adjust its capacity to match the intrinsic complexity of the biological data.

In summary, CryoDECO provides a robust, prior-guided solution for the high-resolution analysis of samples exhibiting extreme compositional and conformational heterogeneity. By grounding 3D reconstruction in the generalized feature space of a foundation model, we establish a new standard for computational cryo-EM, moving the field closer to the ultimate goal of systems structural biology: the comprehensive, high-throughput structural analysis of unpurified native cellular machineries

Methods

Prior-Guided Heterogeneous Reconstruction

The CryoDECO framework adopts an encoder-decoder architecture, but unlike traditional methods that learn features from scratch, it introduces a paradigm shift by injecting learned structural priors to guide the learning process. The encoder of CryoDECO is initialized with the pre-trained Cryo-IEF model¹⁶, a foundation model retrained on a vast and diverse dataset of approximately 130 million cryo-EM particles. This initialization is not merely for robust feature extraction, it serves to regularize the highly non-convex optimization landscape of heterogeneous reconstruction. In conventional approaches, random initialization often leads to unstable convergence and "blind" pose searches. By contrast, our encoder leverages the universal structural knowledge encapsulated in Cryo-IEF to map 2D images directly to a semantically structured latent manifold from the onset. Although the decoder is initialized from scratch, this structured input acts as a strong regularizer, ensuring that the optimization begins within a valid structural manifold, thereby reducing the complexity of the subsequent optimization.

Retraining of the Cryo-IEF Foundation Model

To construct a more robust structural prior for CryoDECO, we retrained the Cryo-IEF foundation model on an expanded corpus of approximately 134 million particle images. While the original Cryo-IEF used a ViT-Base backbone¹⁹, we employed a Vision Transformer Small (ViT-Small) architecture (12 layers, 384 embedding dimension, 12 heads) by default to optimize computational efficiency (Extended Data Fig. 1b).

We trained the model from scratch using the Momentum Contrast (MoCo v3) self-supervised learning framework²⁰, following the hyperparameter configuration of the original Cryo-IEF study¹⁶. We used the AdamW optimizer with a batch size of 2,048. The learning rate followed a cosine decay schedule with a base learning rate of 1.5×10^{-4} (scaled linearly with batch size: $\text{lr} \times \text{BatchSize} / 256$) and a 5-epoch linear warmup. The contrastive loss temperature τ was 0.5, and the momentum update parameter for the key encoder was set to 0.99.

Preprocessing and data augmentation matched the original protocol. Input particle images were resized to 224×224 pixels and normalized. During training, we generated views via random cropping, horizontal flipping, and Gaussian blurring to enforce invariance to noise and orientation. This retraining ensures that the encoder learns a generalized latent manifold capable of separating complex structural states before fine-tuning for specific reconstruction tasks.

Model Architecture

Encoder as a Structural Compass. The encoder is built upon a Vision Transformer (ViT-Small) architecture initialized with Cryo-IEF weights, serving as a "structural compass" for the downstream decoder. The backbone consists of a 12-layer transformer encoder with an embedding dimension of 384 and 12 attention heads, and it processes the input image using a patch size of 14×14 . Evaluation of ViT variants confirmed that the Small architecture provides the optimal trade-off between computational efficiency and representation performance (Extended Data Fig. 1b).

294

295 **Decoder as a neural field.** The decoder is a 3-layer residual MLP with a hidden dimension
 296 of 256, functioning as a coordinate-based neural field. To synthesize the density map, the
 297 network queries 3D spatial coordinates that are first transformed by the particle's pose
 298 parameters $\phi = (R, t)$. We employ Fourier feature positional encoding $\gamma(\cdot)$ to map these
 299 pose-transformed coordinates into high-dimensional features. The input to the MLP is
 300 constructed by adhering the structural information to the geometric framework: we
 301 concatenate the latent descriptor z (from the encoder) with the spatially encoded
 302 coordinates ($\text{cat}[z, \gamma(R\mathbf{x} + t)]$). This design ensures the MLP synthesizes fine structural
 303 details dictated by z at the precise orientation defined by ϕ . Incorporating Layer
 304 Normalization (LN) after each linear layer stabilized training and accelerated convergence
 305 compared to a standard architecture (Extended Data Fig. 1a).

306

307 **Latent manifold regularization.** The dimensionality of the latent space z acts as a critical
 308 topological regularizer that affects the model's ability to capture underlying structural
 309 heterogeneity (Extended Data Fig. 3). On real-world datasets, we tailor z to the intrinsic
 310 complexity of the heterogeneity. For discrete compositional heterogeneity, a high-
 311 dimensional space (e.g., $z = 128$) is necessary to provide sufficient orthogonality,
 312 allowing the model to encode disjoint structural states without forcing them into
 313 overlapping regions of the manifold. In the case of conformational dynamics, the optimal
 314 dimensionality depends on the system's degrees of freedom. We find that a tight
 315 information bottleneck (e.g., $z = 4$) is optimal for simple continuous motions, forcing the
 316 model to map dynamics onto a smooth, low-dimensional manifold rather than overfitting to

noise. However, conformational landscapes characterized by high entropy and multiple degrees of freedom—such as the randomized flexible linker in the IgG-RL dataset (Extended Data Fig. 3)—demand expanded latent capacity (e.g., $z = 64$) to accurately capture the complex non-linear feature space.

Optimization Protocol

CryoDECO's optimization progressively adapts the universal Cryo-IEF priors to specific experimental data using a three-stage protocol¹⁵.

Warm-up adaptation. The network first undergoes a brief adaptation phase using a subset of 10,000 images with fixed random poses. Our encoder leverages its pre-trained weights to project particle images onto a structured latent manifold, enabling the decoder to establish a valid coarse structural distribution before pose optimization begins. During this stage, network parameters are updated via gradient descent minimizing the reconstruction loss using the AdamW optimizer (learning rate 10^{-5} , batch size 64).

Global pose discovery (Hierarchical search). Pose estimation in *ab initio* reconstruction is challenging because it is coupled with unknown structural heterogeneity. A coarse-to-fine hierarchical search strategy is employed in this stage, parameterizing the search space on the $SO(3)$ group using the Hopf fibration and discretizing it into a multi-resolution grid (HEALPix). The hierarchical pose search is performed on the dataset to determine the optimal discrete orientation for each particle. Leveraging the structural priors from the pre-trained Cryo-IEF smooths the optimization landscape, enabling this iterative search to

converge more reliably toward the global optimum. In this stage batch size is reduced to 44 to accommodate memory constraints.

Refinement with SGD (Continuous optimization). Finally, we transition from discrete grid search to continuous optimization. In this stage, the pose parameters $\phi_i = (R_i, t_i)$ for each particle i are instantiated as learnable tensors stored in a lookup table. These variables are initialized with the best values from the hierarchical search. We then perform joint optimization of the network weights θ , latent codes z_i , and pose parameters ϕ_i using Stochastic Gradient Descent (SGD) with the Adam optimizer. Backpropagation flows directly into the pose lookup table, allowing the model to refine orientations and translations continuously beyond the resolution limits of the initial search grid. In this stage batch size is increased to 384.

Objective Functions

Cryo-IEF pre-training. The Cryo-IEF encoder backbone is pre-trained using the InfoNCE contrastive loss function²¹ with a temperature parameter $\tau = 0.5$. This objective maximizes the similarity between different augmented views of the same particle (positive pairs) while minimizing similarity with other particles (negative pairs), thereby driving the encoder to learn representations that are robust to noise and orientation. The loss for a query representation \mathbf{q} is defined as:

$$\mathcal{L}_{contrastive} = -\log \frac{\exp(\mathbf{q} \cdot \mathbf{k}^+ / \tau)}{\exp(\mathbf{q} \cdot \mathbf{k}^+ / \tau) + \sum_{\mathbf{k}^-} \exp(\mathbf{q} \cdot \mathbf{k}^- / \tau)} \quad (1)$$

where \mathbf{k}^+ represents the positive key (an augmented view of the same particle) and \mathbf{k}^- represents the set of negative keys (views of other particles in the batch).

Generative reconstruction. To bridge the gap between latent feature extraction and 3D structural determination, the full CryoDECO framework (encoder and decoder) is fine-tuned by minimizing a physics-based reconstruction loss. This loss quantifies the discrepancy between the experimental particle images and the theoretical projections generated by the network. The specific objective function is given by:

$$\mathcal{L}(\theta, \{\phi_i\}, \{z_i\}) = \sum_{i=1}^N \|\mathcal{H}[I_i] - \hat{\mathcal{C}}_i \odot T_{t_i} S_{R_i} \mathcal{V}_\theta[z_i]\|_2^2 \quad (2)$$

Here, θ denotes network parameters, and $\phi_i = (\mathbf{R}_i, \mathbf{t}_i)$ is the pose of the i -th particle, updated via pose search and then refined with SGD. The decoder \mathcal{V}_θ maps the latent embedding \mathbf{z}_i to a 3D density volume in the Fourier domain. The operator S_{R_i} extracts a central slice perpendicular to viewing direction \mathbf{R}_i . The phase-shift operator T_{t_i} corrects for in-plane translation \mathbf{t}_i , and $\hat{\mathcal{C}}_i$ is the Contrast Transfer Function (CTF). Minimizing this objective jointly optimizes the network for accurate density generation and infers the latent conformation \mathbf{z}_i and pose ϕ_i for each particle.

Data

Dataset for pre-training cryo-IEF. On the basis the 117 datasets (~65 million particles) used in previous work¹⁶, we further expanded the pre-training dataset for Cryo-IEF by incorporating an additional 308 publicly available cryo-EM datasets from EMPIAR²², as summarized in Extended Data Table 1. The detailed processing steps for each dataset can

be found in ¹⁶. The expanded dataset comprises approximately 134 million particle images, significantly enhancing the diversity and representativeness of the training data.

CryoBench datasets. The Ribosembly dataset contains 16 distinct structures of the ribosome assembly intermediate. It presents a discrete heterogeneity problem where structures share a common core that progressively acquires additional proteins and ribosomal RNA. The particle distribution across states is non-uniform, totaling 335,240 particles, mimicking realistic experimental conditions. The data was simulated at a signal-to-noise ratio (SNR) of 0.01, testing the method's ability to resolve and classify discrete compositional states under typical noise levels.

The Tomotwin-100 dataset represents an extreme case of compositional heterogeneity. This dataset contains a mixture of 100 different macromolecular complexes from a curated cellular catalogue, with molecular weights ranging from approximately 200 kDa to 4 MDa. It comprises 100,000 particle images at an SNR of 0.01. It challenges the method's capacity to perform ab initio reconstruction and classification in a highly complex mixture without a common structural core, simulating a scenario akin to analyzing cellular lysates.

The IgG-1D dataset features continuous conformational heterogeneity simulated from an IgG antibody (PDB: 1HZH). The heterogeneity is generated by a simple, one-dimensional rotation of a Fab domain around a dihedral angle, creating a continuous circular motion across 100 conformations. The dataset contains 100,000 particles simulated

at an SNR of 0.01. It serves as a diagnostic benchmark for a method's ability to recover a simple, continuous low-dimensional manifold of conformational change.

The IgG-RL dataset presents a more challenging and realistic form of conformational heterogeneity. It uses the same IgG structure but introduces flexibility by randomizing the conformation of a 5-residue peptide linker, causing one Fab domain to sample a wide range of random orientations. This dataset also contains 100,000 particles at an SNR of 0.01, testing the method's robustness in handling complex, non-linear motions.

EM Ladder dataset. Micrographs for the "EM ladder" dataset were downloaded from EMPIAR-11693¹⁸ and imported into cryoSPARC (v4.6.0)²³. We performed CTF estimation (CTFFIND4²⁴), particle picking (Blob Picker), particle extraction, and 2D classification. This process yielded 117,527 particle images with a box size of 400 pixels, which were then downsampled to a 128-pixel box size using the Fourier crop tool in cryoSPARC. 3D refinement of the four structures was performed separately using cryoSPARC's Homogeneous Refinement tool with appropriate symmetry settings: O for Apoferritin, D2 for β -galactosidase, and I for the PP7 virus-like particle. Before the final refinement, all particles are mapped back to their original 400-pixel box size in CryoSPARC. For TMV, helical refinement was used with a rise of 1.408 Å and a twist of 22.04°. The original study¹⁸ reported high-resolution reconstructions for Apoferritin (2.52 Å; 101,354 particles), β -galactosidase (3.09 Å; 5,696 particles), PP7 (3.40 Å; 2,351 particles), and TMV (2.85 Å; 4,972 particles).

Time-resolved Cryo-EM data of GroEL-GroES complexes. This dataset (EMPIAR-11481)⁶ contains 13,664 movies in total (2,609 at 13 ms, 1,416 at 50 ms, 5,445 at 200 ms, and 4,194 at 20 s), which were processed using cryoSPARC²³. After motion correction (Patch Motion Correction) and CTF estimation (Patch CTF Estimation), 5,486,868 total particles were picked (Blob Picker) and extracted from all micrographs with a box size of 480 pixels. Several rounds of 2D classification yielded a final set of 950,014 particles for heterogeneous reconstruction. These particles were downsampled to a 128-pixel box size and mapped back to 480 pixels before final refinement in cryoSPARC.

Cryo-EM data from a Chaetomium thermophilum native cell extract. This dataset (EMPIAR-10892)⁸ comprises 2,799 movies. The movies were processed using cryoSPARC²³, including motion correction (Patch Motion Correction) and CTF estimation (Patch CTF Estimation). From these, 1,846,404 particles were picked (Blob Picker) and extracted from all micrographs with a box size of 320 pixels. Several rounds of 2D classification and template-based picking yielded 78,100 particles for heterogeneous reconstruction. These particles were then downsampled to a 128-pixel box size for input into CryoDECO and mapped back to 320 pixels before final refinement in cryoSPARC.

Cryo-EM data of tri-snRNP. This dataset (EMPIAR-10073)²⁵ consists of 138,899 particles of the U4/U6.U5 tri-snRNP (tri-small nuclear ribonucleoprotein) complex from *Saccharomyces cerevisiae*. Particles were downsampled to a 128-pixel box size before analysis.

Cryo-EM data of $\alpha V\beta 8$ integrin. This dataset (EMPIAR-10345)²⁶ consists of 84,266 particles of the $\alpha V\beta 8$ integrin complex. These particles, selected after 2D classification, were downsampled to a 128-pixel box size.

Metrics

K-NN scores. We utilize the K-NN score to quantitatively evaluate how well the extracted features separate particles from different heterogeneous states. We employ a Leave-One-Out (LOO) evaluation strategy. For each data point (feature vector) \mathbf{f}_i in the entire dataset, we identify its k nearest neighbors by computing the dot product similarity S_{ij} with all other data points \mathbf{f}_j (where $j \neq i$). Instead of a simple majority vote, we perform a weighted classification. The contribution of each of the k neighbors is weighted based on its similarity score S_{ij} and a temperature parameter T . The score for a specific class c for sample i is computed by summing the weights of its neighbors belonging to that class:

$$\text{Score}_i(c) = \sum_{r=1}^k \exp(S_{ir}/T) \cdot \mathbf{1}(l_{ir} = c) \quad (3)$$

where l_{ir} is the label of the r -th nearest neighbor for sample i (found from the set where $j \neq i$), and $\mathbf{1}(\cdot)$ is the indicator function. In our experiments, we set $k = 10$ and $T = 1$. We report the Top-1 and Top-10 accuracy. Let y_i be the true class for sample i , and let $(\hat{y}_i^{(1)}, \hat{y}_i^{(2)}, \dots, \hat{y}_i^{(k')})$ be the list of predicted classes ranked by $\text{Score}_i(c)$, where $k' = \min(k, C)$ and C is the total number of classes. The Top-1 score is the standard LOO accuracy, checking if the highest-scoring predicted class matches the true class:

$$\text{Top-1 Score} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}(\hat{y}_i^{(1)} = y_i) \quad (4)$$

where N is the total number of samples in the dataset. The Top-10 score measures if the true class y_i is present within the top 10 ranked predictions:

$$\text{Top-10 Score} = \frac{1}{N} \sum_{i=1}^N \left[\sum_{j=1}^{\min(10, k, C)} \mathbf{1}(\hat{y}_i^{(j)} = y_i) \right] \quad (5)$$

Adjusted Rand Index (ARI). The Adjusted Rand Index (ARI) is used to measure the similarity between the true class assignments (ground truth U) and the predicted clustering results (V). Given a set of N samples, let $U = \{u_1, u_2, \dots, u_R\}$ be the true class assignments (with R classes), and $V = \{v_1, v_2, \dots, v_C\}$ be the predicted cluster assignments (with C clusters). We can construct a contingency table where n_{ij} represents the number of samples that are in both true class u_i and predicted cluster v_j . Let $a_i = \sum_j n_{ij}$ be the total number of samples in true class u_i , and $b_j = \sum_i n_{ij}$ be the total number of samples in predicted cluster v_j . The ARI is calculated as:

$$\text{ARI} = \frac{\sum_{ij} \binom{n_{ij}}{2} - \left[\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2} \right] / \binom{N}{2}}{\frac{1}{2} \left[\sum_i \binom{a_i}{2} + \sum_j \binom{b_j}{2} \right] - \left[\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2} \right] / \binom{N}{2}} \quad (6)$$

The ARI score ranges from -1 to 1. A score of 1 indicates perfect agreement between the true labels and the predicted clusters, while a score near 0 indicates that the clustering agreement is only at the level expected by chance.

Adjusted Mutual Information (AMI). The Adjusted Mutual Information (AMI) is another metric used to assess the agreement between two clusterings (true labels U and predicted labels V). The AMI is defined as:

$$AMI(U, V) = \frac{MI(U, V) - E[MI(U, V)]}{\max(H(U), H(V)) - E[MI(U, V)]} \quad (7)$$

where $MI(U, V)$ is the mutual information between the two clusterings, $H(U)$ and $H(V)$ are the entropies of the respective clusterings, and $E[MI(U, V)]$ is the expected mutual information between two random clusterings with the same number of clusters and data distribution. AMI scores typically range from 0 to 1. A score of 1 indicates a perfect match between the clusterings, while a score of 0 indicates that the mutual information between the two is only what would be expected by random chance.

CryoDRGN-AI Analysis. CryoDRGN-AI¹⁵ was used as a benchmark method for comparison with CryoDECO. Instead of using traditional encoder for feature extraction, CryoDRGN-AI employs a look-up table to directly learn a latent vector for each particle image. In our experiments, we followed the official guidelines provided in the CryoDRGN-AI repository (<https://github.com/ml-struct-bio/drgnai>) to set up and run the experiments. All hyperparameters were set to their default values unless specified otherwise.

Computational Resources. All particles were downsampled to a 128-pixel box size prior to training. Experiments were conducted using NVIDIA A40 (40 GB memory) and NVIDIA A100 (80 GB memory) GPUs. Approximate training times for CryoDECO for each dataset were as follows: CryoBench (100,000 particles; IgG-1D, IgG-RL, Tomotwin-

100): ~13 hours on 2x NVIDIA A40 GPUs. CryoBench (335,240 particles; Ribosembly):
 ~39 hours on 2x NVIDIA A40 GPUs. EM Ladder (111,733 particles; EMPIAR-11693):
 ~14 hours on 2x NVIDIA A40 GPUs. GroEL–GroES (959,834 particles; EMPIAR-11481):
 ~37 hours on 8x NVIDIA A100 GPUs. C. thermophilum (78,100 particles; EMPIAR-
 10892): ~9 hours on 2x NVIDIA A40 GPUs.

Data availability. The raw micrographs are available in the Electron Microscopy Public
 Image Archive <https://www.ebi.ac.uk/empirar/> under accession code EMPIAR-11693,
 EMPIAR-11481 and EMPIAR–10892. The IgG-D and IgG-RL datasets from CryoBench
 are available at <https://zenodo.org/records/11629428>, and the Ribosembly and Tomotwin-
 100 datasets from CryoBench are available at <https://zenodo.org/records/12528292>.

Code availability. The codes and pretrained models for Cryo-IEF are available at
<https://github.com/westlake-repl/Cryo-IEF>. The codes for CryoDECO are available at
<https://github.com/yanyang1998/CryoDECO>.

- 527 1. Yip, K. M., Fischer, N., Paknia, E., Chari, A. & Stark, H. Atomic-resolution protein
528 structure determination by cryo-EM. *Nature* **587**, 157–161 (2020).
- 529 2. Nakane, T. *et al.* Single-particle cryo-EM at atomic resolution. *Nature* **587**, 152–156
530 (2020).
- 531 3. Frank, J., Goldfarb, W., Eisenberg, D. & Baker, T. S. Reconstruction of glutamine
532 synthetase using computer averaging. *Ultramicroscopy* **3**, 283–290 (1978).
- 533 4. Cheng, Y. Single-Particle Cryo-EM at Crystallographic Resolution. *Cell* **161**, 450–457
534 (2015).
- 535 5. Cheng, Y. Single-particle cryo-EM—How did it get here and where will it go. *Science*
536 <https://doi.org/10.1126/science.aat4346> (2018) doi:10.1126/science.aat4346.
- 537 6. Torino, S., Dhurandhar, M., Stroobants, A., Claessens, R. & Efremov, R. G. Time-
538 resolved cryo-EM using a combination of droplet microfluidics with on-demand
539 jetting. *Nat. Methods* **20**, 1400–1408 (2023).
- 540 7. Tüting, C. *et al.* Cryo-EM snapshots of a native lysate provide structural insights into a
541 metabolon-embedded transacetylase reaction. *Nat. Commun.* **12**, 6933 (2021).
- 542 8. Skolidis, I. *et al.* Cryo-EM and artificial intelligence visualize endogenous protein
543 community members. *Structure* **30**, 575–589.e6 (2022).
- 544 9. Semchonok, D. A., Kyrilis, F. L., Hamdi, F. & Kastiris, P. L. Cryo-EM of a
545 heterogeneous biochemical fraction elucidates multiple protein complexes from a
546 multicellular thermophilic eukaryote. *J. Struct. Biol. X* **8**, 100094 (2023).
- 547 10. Punjani, A. & Fleet, D. J. 3D variability analysis: Resolving continuous flexibility and
548 discrete heterogeneity from single particle cryo-EM. *J. Struct. Biol.* **213**, 107702
549 (2021).
- 550 11. Punjani, A. & Fleet, D. J. 3DFlex: determining structure and motion of flexible
551 proteins from cryo-EM. *Nat. Methods* **20**, 860–870 (2023).
- 552 12. Zhong, E. D., Bepler, T., Berger, B. & Davis, J. H. CryoDRGN: reconstruction of
553 heterogeneous cryo-EM structures using neural networks. *Nat. Methods* **18**, 176–185
554 (2021).
- 555 13. Zhong, E. D., Lerer, A., Davis, J. H. & Berger, B. CryoDRGN2: Ab Initio Neural
556 Reconstruction of 3D Protein Structures From Real Cryo-EM Images. in 4066–4075
557 (2021).
- 558 14. Luo, Z., Ni, F., Wang, Q. & Ma, J. OPUS-DSD: deep structural disentanglement for
559 cryo-EM single-particle analysis. *Nat. Methods* **20**, 1729–1738 (2023).
- 560 15. Levy, A. *et al.* CryoDRGN-AI: neural ab initio reconstruction of challenging cryo-EM
561 and cryo-ET datasets. *Nat. Methods* **22**, 1486–1494 (2025).
- 562 16. Yan, Y., Fan, S., Yuan, F. & Shen, H. A comprehensive foundation model for cryo-
563 EM image processing. *Nat. Methods* 1–8 (2025) doi:10.1038/s41592-025-02916-8.
- 564 17. Jeon, M. *et al.* CryoBench: Diverse and challenging datasets for the heterogeneity
565 problem in cryo-EM. Preprint at <http://arxiv.org/abs/2408.05526> (2024).
- 566 18. Bobe, D., Kopylov, M., Miller, J., Owji, A. P. & Eng, E. T. Multi-species cryoEM
567 calibration and workflow verification standard. *Acta Crystallogr. Sect. F* **80**, 320–327
568 (2024).
- 569 19. Dosovitskiy, A. *et al.* An Image is Worth 16x16 Words: Transformers for Image
570 Recognition at Scale. in (2020).

- 571 20. Chen, X., Xie, S. & He, K. An Empirical Study of Training Self-Supervised Vision
572 Transformers. in *2021 IEEE/CVF International Conference on Computer Vision*
573 (*ICCV*) 9620–9629 (2021). doi:10.1109/ICCV48922.2021.00950.
- 574 21. Oord, A. van den, Li, Y. & Vinyals, O. Representation Learning with Contrastive
575 Predictive Coding. Preprint at <https://doi.org/10.48550/arXiv.1807.03748> (2019).
- 576 22. Iudin, A., Korir, P. K., Salavert-Torres, J., Kleywegt, G. J. & Patwardhan, A.
577 EMPIAR: a public archive for raw electron microscopy image data. *Nat. Methods* **13**,
578 387–388 (2016).
- 579 23. Punjani, A., Rubinstein, J. L., Fleet, D. J. & Brubaker, M. A. cryoSPARC: algorithms
580 for rapid unsupervised cryo-EM structure determination. *Nat. Methods* **14**, 290–296
581 (2017).
- 582 24. Rohou, A. & Grigorieff, N. CTFFIND4: Fast and accurate defocus estimation from
583 electron micrographs. *J. Struct. Biol.* **192**, 216–221 (2015).
- 584 25. Nguyen, T. H. D. *et al.* Cryo-EM structure of the yeast U4/U6.U5 tri-snRNP at 3.7 Å
585 resolution. *Nature* **530**, 298–302 (2016).
- 586 26. Campbell, M. G. *et al.* Cryo-EM Reveals Integrin-Mediated TGF-β Activation without
587 Release from Latent TGF-β. *Cell* **180**, 490-501.e16 (2020).
- 588 27. Pettersen, E. F. *et al.* UCSF ChimeraX: Structure visualization for researchers,
589 educators, and developers. *Protein Sci.* **30**, 70–82 (2021).
- 590

Figures

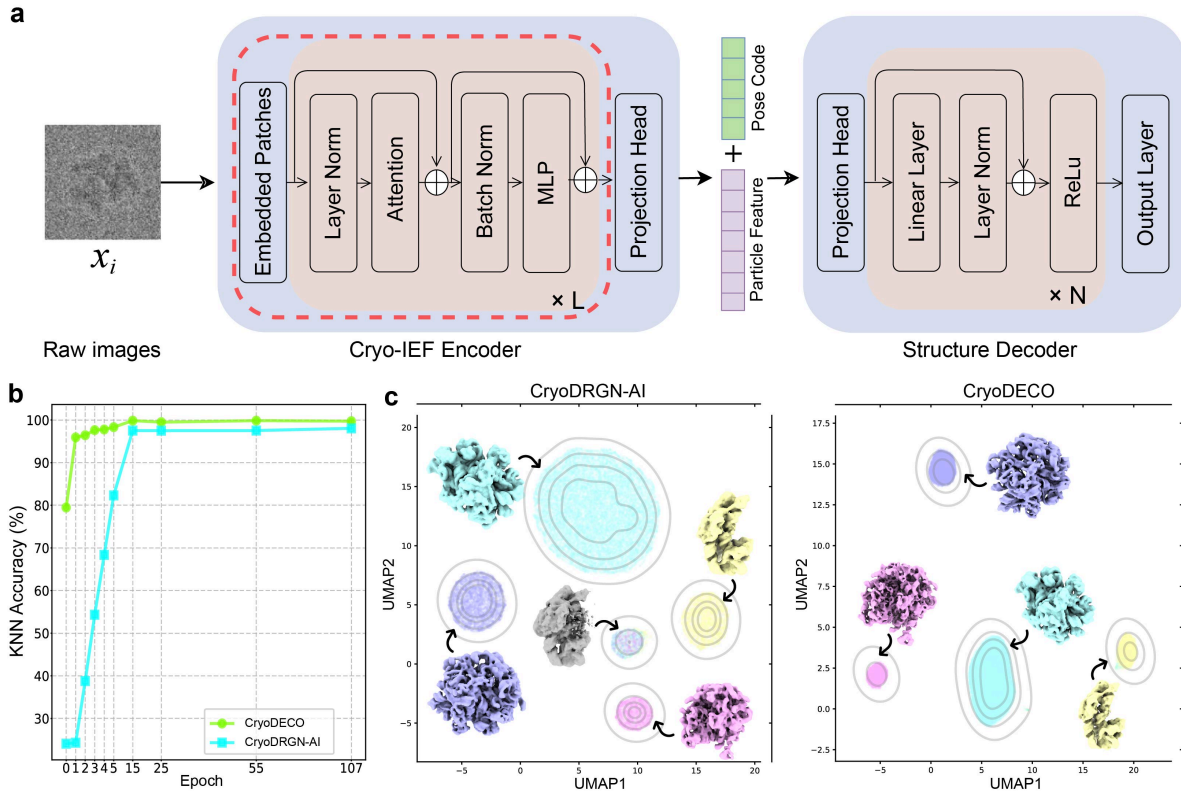


Fig. 1 | Overview and validation of the CryoDECO framework. (a) CryoDECO utilizes an encoder-decoder architecture. The encoder, a self-supervised pre-trained cryo-IEF model, extracts features from particle images. The decoder is a trainable module that reconstructs heterogeneous structures from these features and estimated poses. (b) Comparison of training dynamics (k-NN score vs. epoch) between CryoDECO and CryoDRGN-AI on the test dataset. (c) UMAP visualizations of the feature space at the final training epoch for CryoDRGN-AI and CryoDECO. Data points are color-coded by their ground-truth structures shown in Extended Data Fig 2. Insets show structures generated by the decoder from the center of each feature cluster.

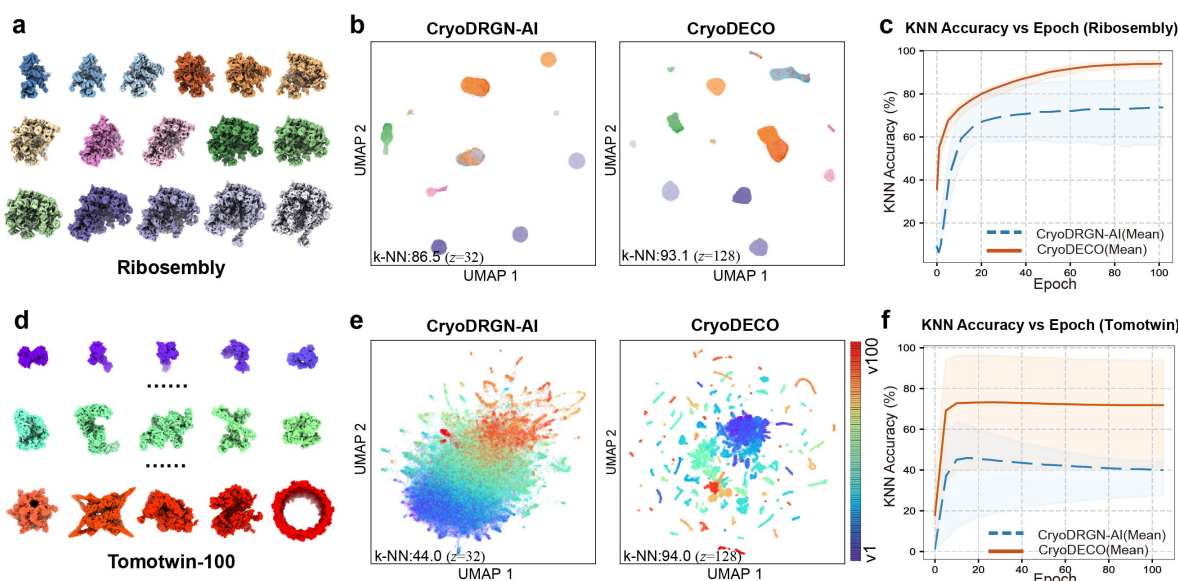


Fig. 2 | CryoDECO's performance on CryoBench compositional heterogeneity

datasets. (a) The 16 ground-truth structures in the Ribosembly dataset. (d) Examples of the 100 ground-truth structures in the Tomotwin-100 dataset. (b, e) UMAP visualizations of particle feature spaces from the Ribosembly (b) and Tomotwin-100 (e) datasets. Features from CryoDRGN-AI (left panels) are shown for comparison. Data points are color-coded by their ground-truth structures. Optimal latent dimensions (z) are shown (Ribosembly: $z = 32$ for CryoDRGN-AI, $z = 128$ for CryoDECO; Tomotwin-100: $z = 32$ for CryoDRGN-AI, $z = 128$ for CryoDECO). (c, f) k-NN scores during training for the Ribosembly (c) and Tomotwin-100 (f) datasets. Scores are averaged over three independent runs with varying z dimensions ($z=4, 16, 32, 64, 128$). CryoDECO (orange line) and CryoDRGN-AI (blue dashed line) are shown. Shaded areas represent the score range across all z dimensions.

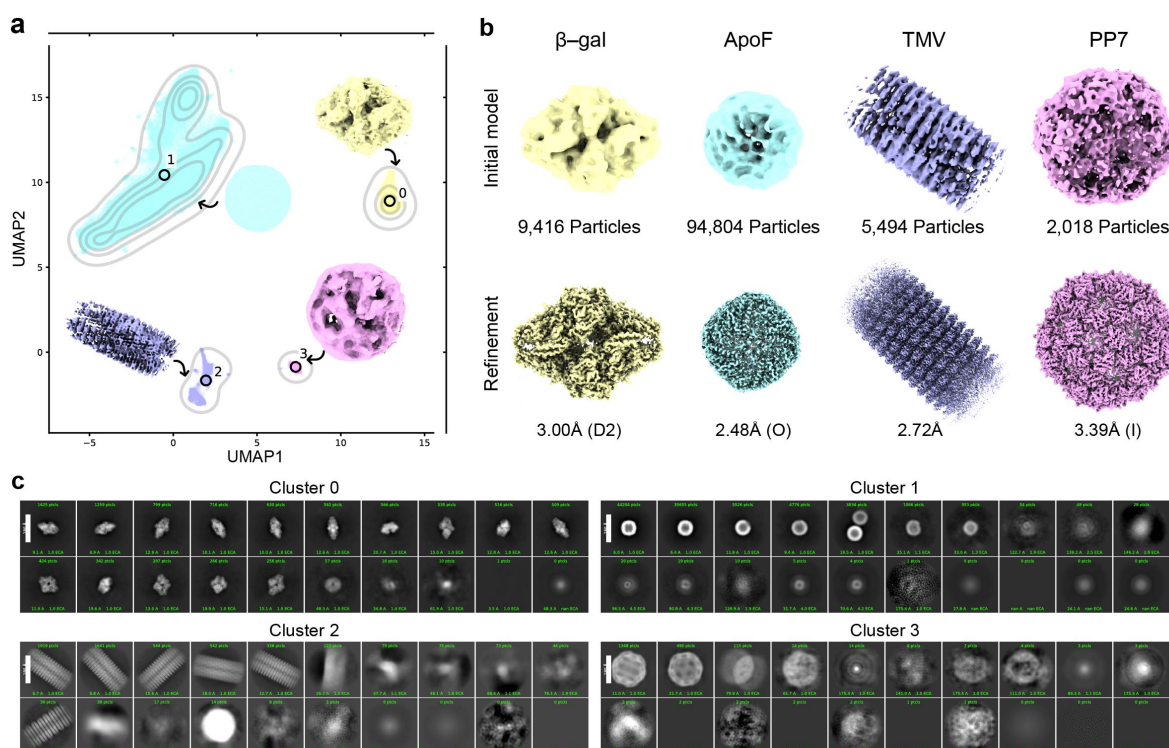


Fig. 3 | CryoDECO classifies four distinct structures from the EM ladder. (a) UMAP visualization of the particle feature space extracted by CryoDECO. Particles are clustered into four groups using GMM (color-coded). Insets show the initial 3D structures generated by the decoder from each cluster center. (b) Final 3D reconstructions of the four particle groups after ab-initio reconstruction and non-uniform refinement in CryoSPARC. Structures are colored to match their corresponding clusters in (a). (c) 2D class averages for each cluster, confirming particle purity. All 3D visualizations were created using ChimeraX²⁷.

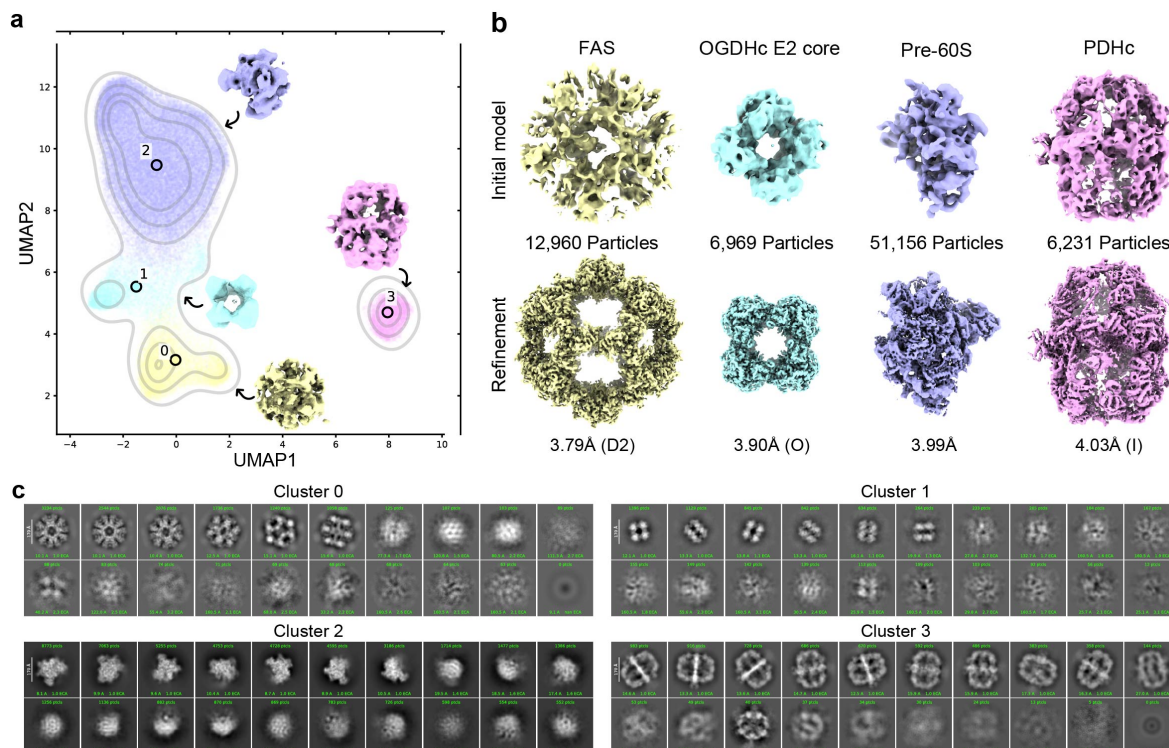


Fig. 4 | CryoDECO classifies endogenous protein communities from a C.

thermophilum native cell extract. (a) UMAP visualization of the particle feature space.

Particles are clustered into four groups using GMM (color-coded). Insets show initial 3D

structures generated by the decoder from each cluster center, colored to match clusters. (b)

Final 3D reconstructions of the four particle groups after ab-initio reconstruction and non-

uniform refinement in CryoSPARC, colored to match (a). (c) Representative 2D class

averages for each cluster. All 3D visualizations were created using ChimeraX.

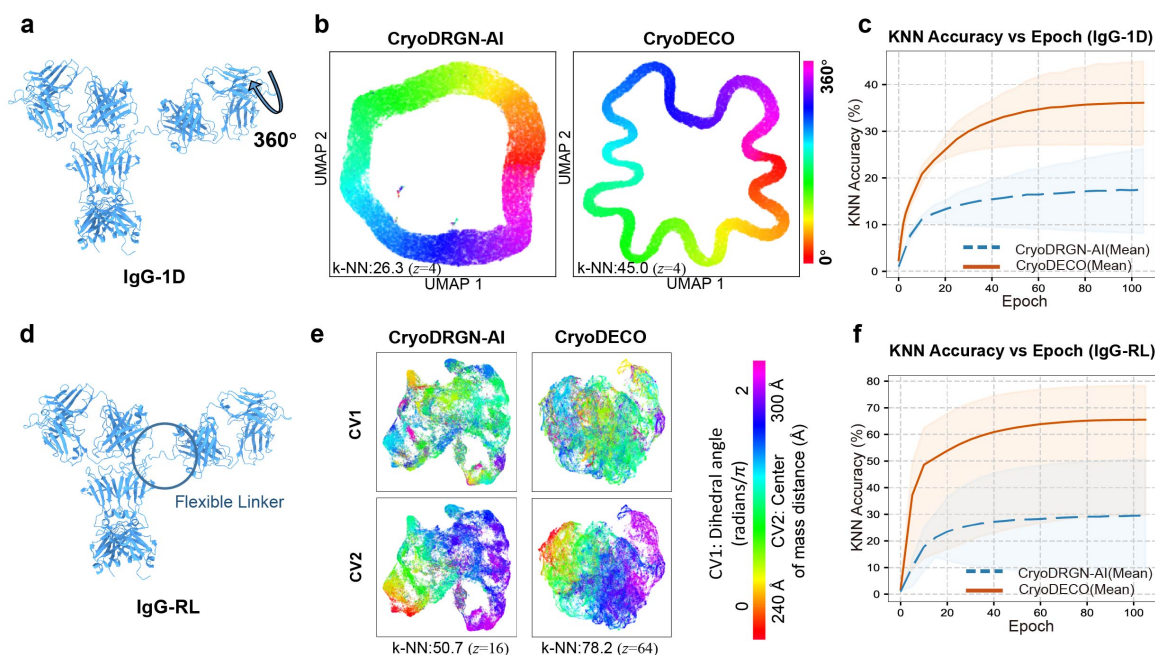


Fig. 5 | CryoDECO performance on CryoBench conformational heterogeneity

datasets. (a) The IgG-1D dataset's heterogeneity is generated by rotating a dihedral angle in one Fab domain. (d) The IgG-RL dataset's heterogeneity is generated by randomizing a linker region. Both datasets contain 100 conformations with 1,000 particles each (SNR = 0.01). (b, e) UMAP visualizations of particle feature spaces from IgG-1D (b) and IgG-RL (e), color-coded by conformation. Features from CryoDRGN- AI (left panels) are shown for comparison. Optimal latent dimensions (z) are shown (IgG-1D: $z = 4$ for both models; IgG-RL: $z = 16$ for CryoDRGN-AI, $z = 64$ for CryoDECO). (c, f) k-NN scores during training for IgG-1D (c) and IgG-RL (f). Scores are averaged over runs with varying z dimensions ($z=4, 16, 32, 64, 128$). CryoDECO (orange line) and CryoDRGN-AI (blue dashed line) are shown. Shaded areas represent the score range across all z dimensions.

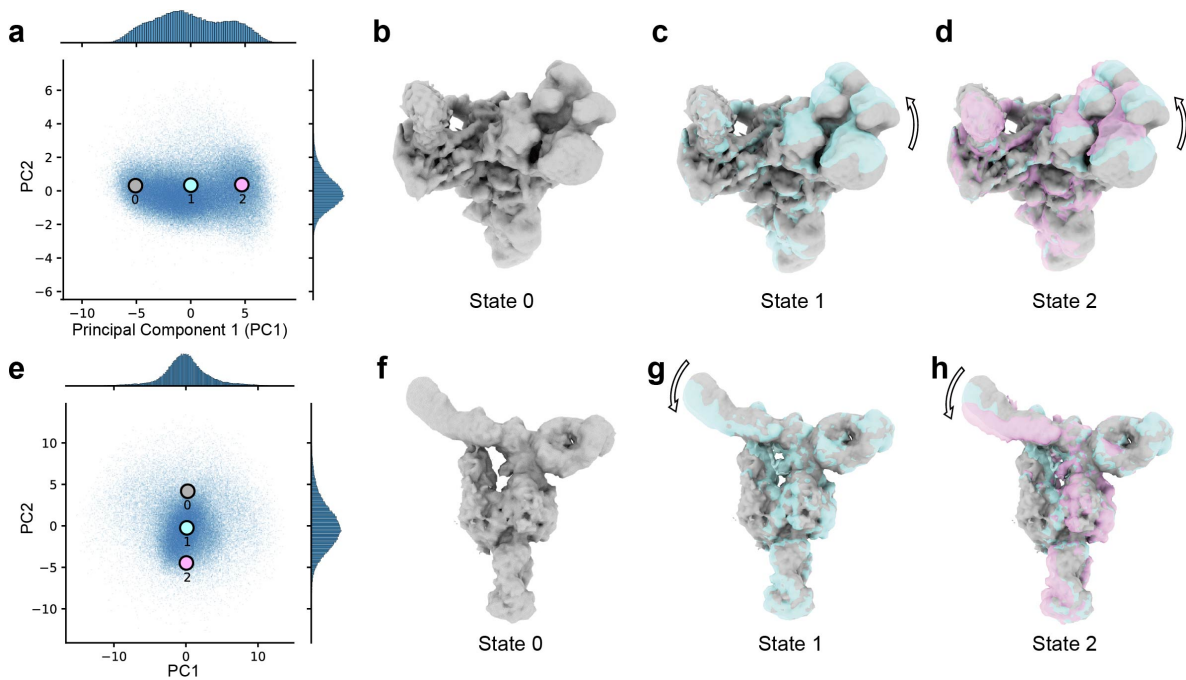


Fig. 6 | Continuous conformational heterogeneity analysis of tri-snRNP and $\alpha V\beta 8$ integrin datasets. (a, e) PCA visualization of the feature space for tri-snRNP (a) and $\alpha V\beta 8$ integrin (e). Three points were selected along PC1 for tri-snRNP (a) and PC2 for $\alpha V\beta 8$ (e) to represent distinct conformations (gray, cyan, violet dots). (b-d) Reconstructions from the selected points for tri-snRNP. (f-h) Reconstructions from the selected points for $\alpha V\beta 8$. (c, g) Superimposition of the first two states (gray, transparent cyan). (d, h) Superimposition of all three states (gray, transparent cyan, transparent violet). All 3D visualizations were created using ChimeraX.

662 **Table 1 Clustering Accuracy**

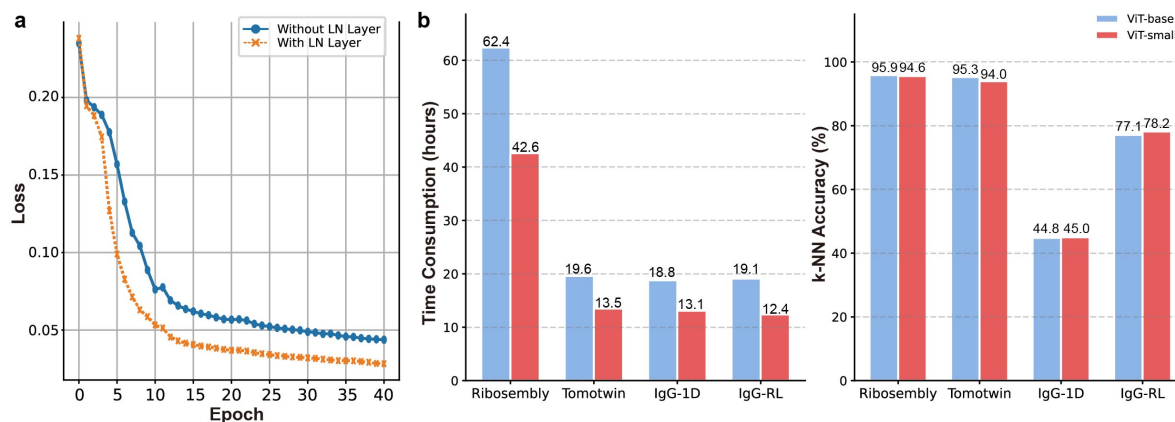
Method		Ribosemblly		Tomotwin-100	
		ARI	AMI	ARI	AMI
Fixed poses†	CryoDRGN*	0.873	0.935	0.956	0.983
	CryoDRGN-AI-fixed*	0.624	0.771	0.791	0.906
	Opus-DSD*	0.891	0.934	0.500	0.781
	3DVA*	0.666	0.823	0.058	0.335
	RECOVAR*	0.968	0.976	0.315	0.649
ab initio‡	CryoDRGN2*	0.529	0.618	0.116	0.374
	CryoDRGN-AI*	0.644	0.729	0.086	0.275
	CryoDECO	0.860	0.908	0.622	0.853

† Reconstruction with fixed (ground truth) poses.

‡ *ab initio* reconstruction where no input poses are provided.

* Result from CryoBench¹⁷.

Extended Data Figures and Tables



Extended Data Fig. 1 | Ablation studies on decoder normalization and encoder scaling

in CryoDECO. (a) Reconstruction loss profiles comparing CryoDECO with (orange

dashed line) and without (blue solid line) layer normalization in the decoder. Training was

performed on the cryoDRGN synthetic toy dataset [12] with 10 epochs for pose searching.

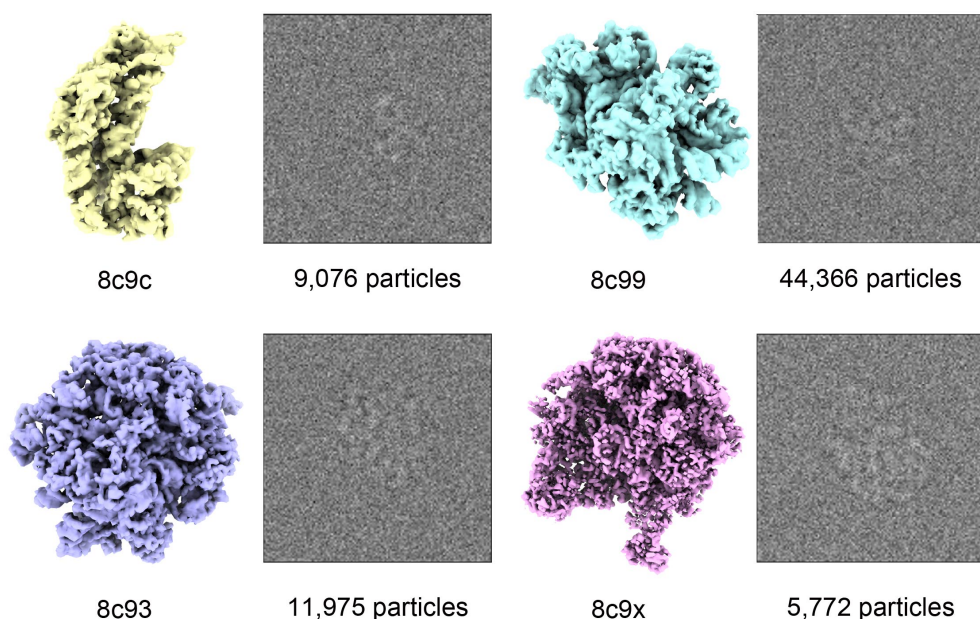
(b) Performance comparison of encoder scaling across four datasets. The bar charts display

the time consumption (left) and latent space quality (k-NN Top-1 accuracy, right) for ViT-

base (blue bars) and ViT-small (red bars) encoders. Note that increasing the encoder size to

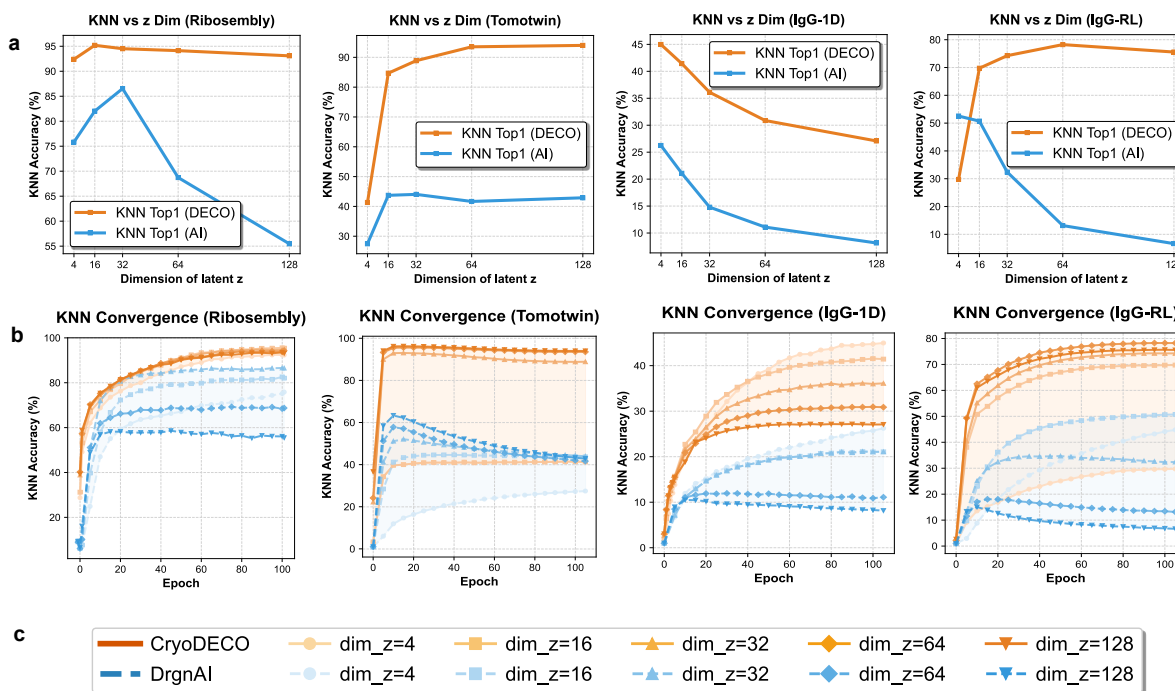
ViT-base significantly raises the computational cost with marginal or negligible gains in

classification accuracy compared to ViT-small.



Extended Data Fig. 2 | The resampled CryoBench Ribosembly dataset. This dataset

contains simulated particles from four distinct structures with SNR=0.01.



Extended Data Fig. 3 | Parameter analysis of the CryoDECO and DrgnAI models on

synthetic datasets. A parameter analysis of the CryoDECO (DECO) and DrgnAI (AI)

models was conducted on the Ribosemblly, Tomotwin, IgG-1D, and IgG-RL synthetic

datasets, specifically examining the effect of latent space dimensionality z on kNN

classification accuracy. (a) Final KNN 10 top-1 classification accuracy as a function of the

latent space dimension z . For each dataset, both models were trained and the final accuracy

was evaluated. (b) KNN 10 top-1 classification accuracy convergence curves over epochs

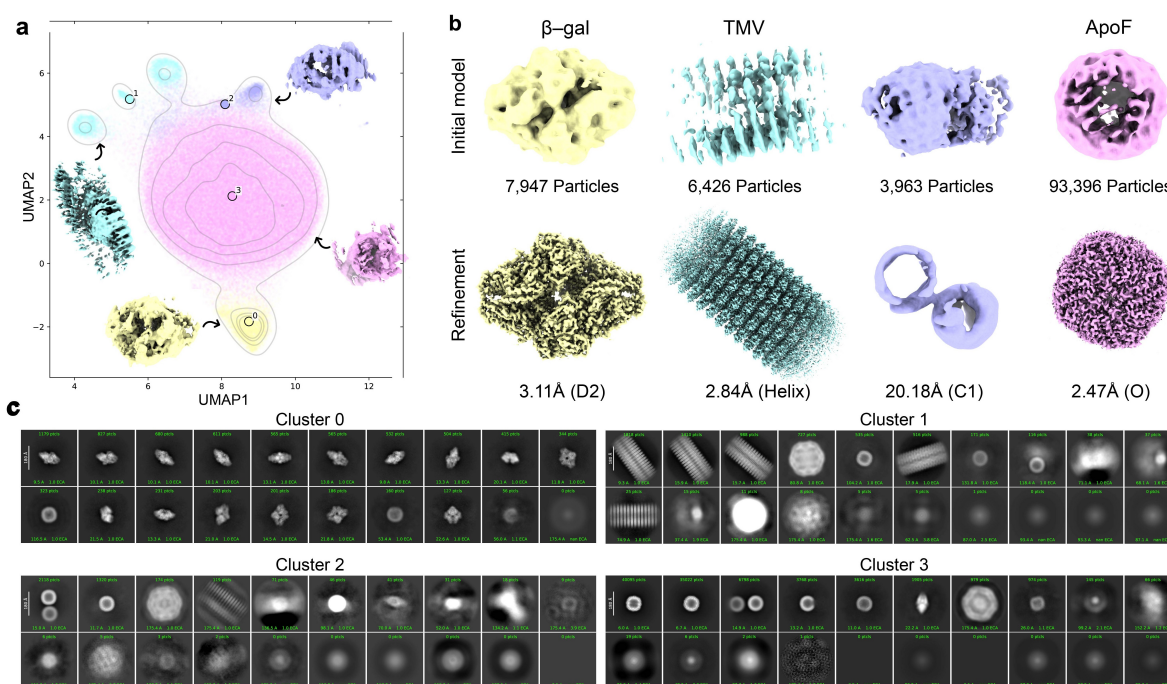
for different latent space dimensions. The orange lines represent CryoDECO (DECO) and

the blue lines represent DrgnAI (AI). Lighter shades correspond to lower dimension, while

lighter shades correspond to lower dimension, as detailed in the legend. (c) Legend

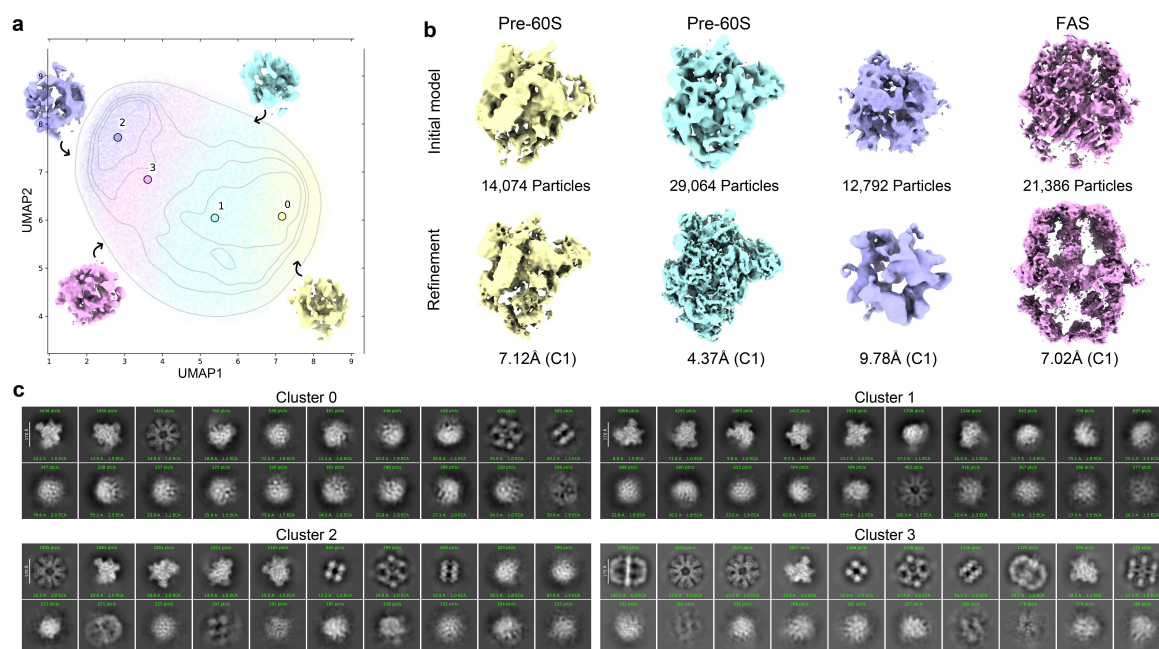
detailing the marker and color scheme used in panel (b) for different latent space

dimensions.



Extended Data Fig. 4 | CryoDRGN-AI benchmark results on the EM ladder. (a)

UMAP visualization of the particle feature space extracted by CryoDRGN-AI. Particles were clustered into four groups (GMM). Insets show 3D structures generated by the decoder from cluster centers. Note the poor separation of clusters compared to Fig. 3a. (b) Final 3D reconstructions of the four clusters after ab-initio reconstruction and non-uniform refinement in CryoSPARC. (c) 2D class averages for each cluster, showing the mixing of particle types. All 3D visualizations were created using ChimeraX.



Extended Data Fig. 5 | CryoDRGN-AI benchmark results on the *C. thermophilum*

native cell extract. (a) UMAP visualization of the particle feature space. Particles were clustered into four groups (GMM). Insets show 3D structures generated by the decoder from cluster centers. Note the significant overlap between clusters compared to Fig. 4a. (b) Final 3D reconstructions of the four clusters. (c) 2D class averages for each cluster, confirming significant structural mixing. All 3D visualizations were created using ChimeraX.

716 **Extended Data Table 1 | Summary of Expanded Data Downloaded from EMPIAR for**

717 **Cryo-IEF Pre-training.**

718

719

EMPIAR ID									
10038	10153	10181	10193	10203	10249	10257	10258	10261	10264
10281	10285	10288	10290	10299	10314	10317	10332	10334	10335
10342	10344	10350	10352	10360	10362	10379	10380	10391	10399
10400	10407	10420	10422	10423	10425	10429	10433	10437	10438
10439	10443	10470	10477	10485	10486	10487	10495	10503	10507
10518	10521	10522	10530	10536	10540	10543	10544	10547	10549
10552	10557	10561	10573	10574	10575	10577	10580	10581	10582
10584	10594	10595	10596	10597	10599	10600	10604	10605	10611
10626	10628	10629	10630	10642	10652	10654	10656	10657	10661
10663	10666	10667	10682	10684	10695	10706	10707	10714	10716
10721	10726	10728	10752	10758	10759	10763	10777	10778	10793
10803	10808	10810	10811	10833	10834	10837	10842	10846	10850
10854	10855	10856	10858	10863	10864	10866	10877	10878	10880
10893	10894	10902	10909	10910	10915	10919	10927	10931	10932
10941	10942	10945	10951	10952	10964	10965	10966	10974	10975
10977	10978	10979	10984	10990	11006	11007	11008	11012	11021
11022	11028	11030	11031	11033	11036	11046	11047	11048	11049
11052	11053	11059	11071	11072	11073	11075	11076	11079	11080
11081	11091	11099	11110	11124	11127	11128	11132	11133	11134
11135	11137	11138	11143	11158	11171	11176	11177	11178	11179
11180	11182	11184	11185	11191	11192	11193	11194	11197	11218
11223	11224	11225	11227	11229	11231	11232	11234	11235	11241
11242	11251	11255	11256	11266	11268	11279	11295	11298	11304
11305	11307	11308	11313	11326	11328	11329	11330	11331	11333
11334	11335	11336	11337	11340	11341	11348	11352	11355	11364
11373	11374	11378	11382	11389	11390	11409	11441	11492	11512
11513	11524	11536	11559	11560	11568	11600	11604	11608	11624
11642	11645	11648	11654	11655	11656	11662	11663	11689	11697
11708	11713	11719	11725	11726	11727	11728	11760	11761	11780
11782	11786	11792	11797	11804	11810	11824	11859	11868	11887
11894	11900	11915	11916	12023	12024	12112	12146		

720