Hidden dynamics of proteasome autoregulation discovered by cryo-EM data-driven deep learning

Zhaolong Wu[1,2,3], Enbo Chen[1,2], Shuwen Zhang[1,2], Wei Li Wang[1,2], Yinping Ma[1], Yuanchen Dong[4], Congcong Liu[5,7], Chang-Cheng Yin[5] & Youdong Mao[1,2,3,6]*

[1]State Key Laboratory for Artificial Microstructure and Mesoscopic Physics, School of Physics, Peking University, Beijing, China. [2]Peking-Tsinghua Joint Center for Life Sciences, Peking University, Beijing, China. [3]Center for Quantitative Biology, Peking University, Beijing, China. [4]Institute of Chemistry, Chinese Academy of Sciences, Beijing, China. [5]Department of Biophysics, School of Basic Medical Sciences, Peking University, Beijing, China. [6]National Biomedical Imaging Center, Peking University, Beijing, China. [7]Current Address: Institute for Hepatology, National Clinical Research Center for Infectious Disease, Shenzhen Third People's Hospital, Shenzhen, Guangdong, China.

*To whom correspondence may be addressed. E-mail: ymao@pku.edu.cn

Abstract

The cellular functions are executed by macromolecular complexes in nonequilibrium dynamic processes, which exhibit a vast diversity of conformational states. Solving conformational continuum of important complexes at atomic level remains a key challenge in the methodology of computational biophysics. Here we introduce a deep-learning framework, named AlphaCryo4D, which enables atomic-level cryogenic electron microscopy reconstructions of nonequilibrium conformational continuum. AlphaCryo4D integrates 3D deep residual learning with manifold embedding of energy landscapes, which simultaneously improves 3D classification accuracy and reconstruction resolution via an energy-based particle-voting algorithm. Applications of this methodology to six experimental datasets demonstrated its generality in breaking resolution barrier of visualizing dynamic components of functional complexes, in choreographing continuous inter-subunit motions and in exploring their 'hidden' conformational space. Importantly, AlphaCryo4D identified 64 conformers of the substrate-bound human 26S proteasome and reconstituted single-nucleotide-exchange dynamics of proteasomal AAA-ATPase motor during translocation initiation, which together illuminates a grand hierarchical allostery for proteasome autoregulation.

Essential cellular functions are executed and regulated by macromolecular complexes comprising many subunits, such as the ribosome, proteasome, replisome and spliceosome. Their structures in functional cycles are often compositionally heterogeneous and conformationally dynamic, involving many reversibly associated components in cells. Visualizing the conformational continuum of highly dynamic megadalton complexes at high resolution is crucial for understanding their functional mechanisms and for guiding structure-based therapeutic development. The current approaches in cryogenic electron microscopy (cryo-EM) structure determination allow atomic-level visualization of major conformational states of these dynamic complexes[1-12]. However, high-resolution structure determination of transient, nonequilibrium intermediates connecting their major states during their functional cycles has been prohibitory at large. To date, there has been lack of appropriate approaches in visualizing conformational space of large biomolecular complexes at high resolution.

The problem of structural heterogeneity in cryo-EM structure determination has been investigated with numerous computational approaches including maximum-likelihood-based classification and multivariate statistical analysis[5,6,11,13-18]. As an extension of multivariate statistical analysis, several machine-learning approaches were recently proposed to estimate a continuous conformational distribution in the latent space[11,12,19-22]. The estimation of continuous conformational distributions does not guarantee an improvement of 3D classification accuracy nor warrant identification of hidden conformational states of biological importance. To date, the most widely used method to counteract structural heterogeneity in cryo-EM for resolution improvement is the hierarchical maximum-likelihood-based 3D classification[1] and its related methods including the multi-body or focused refinement[4,23,24], normal mode analysis[25,26], flexible refinement[27,28], and non-uniform refinement[29]. To improve the resolution of major conformers, low-quality classes are often manually removed during iterations of curated hierarchical classification, inevitably causing an incomplete representation of the true conformational landscape[30,31]. The outcome of refined cryo-EM maps depends on user expertise of decision-making in class selection, which can be biased by user experience and subjectivity. Importantly, a considerable portion of misclassified images can limit the achievable resolution of observed conformational states and result in missing states that could be biologically important.

The energy landscape is a statistical physical representation of the conformational space of a macromolecular complex and is the basis of the transition-state theory of chemical reaction

dynamics[32,33]. The minimum-energy path (MEP) on the energy landscape theoretically represents the most probable trajectory of conformational transitions and can inform the activation energy for chemical reactions[33,34]. Previous studies have demonstrated the benefit of energy landscape estimation in characterizing conformational variation of macromolecules from cryo-EM data[35-41]. The combination of linear dimensionality reduction via 3D principal component analysis (PCA) and energy landscape estimation has been demonstrated to be capable of capturing the overall conformational landscape[36-38]. The benefit of using nonlinear dimensionality reduction and manifold embedding, such as diffusion map[42], to replace PCA for estimating energy landscapes has also been investigated at limited resolution, with the assumption that the conformational changes can be discerned through a narrow angular aperture[35,39,40]. This assumption restricts its potential applications to more complicated conformational dynamics.

Breaking resolution barrier in visualizing the conformational space of highly dynamic complexes represents a major objective of technical advances for leveraging cryo-EM applications in both basic biomedical science and drug discovery. This would require substantial improvement in 3D classification techniques that are optimized for resolution improvement. We hypothesize that a high-quality estimation of energy landscape could be potentially used to improve 3D classification of cryo-EM data. If different conformers or continuous conformational changes can be sufficiently mapped and differentiated on the energy landscape, the energetic visualization of conformational continuum could then be used to discover previously 'invisible' conformers and achieve higher resolution for lowly populated, transient or nonequilibrium states. To explore these ideas, we developed a novel machine learning framework named AlphaCryo4D that can break the existing limitations and enable 4D cryo-EM reconstruction of highly dynamic, lowly populated intermediates or transient states at atomic level. We examined the general applicability of AlphaCryo4D in analyzing several synthetic datasets over a wide range of single-to-noise ratios (SNRs) and six experimental cryo-EM datasets of diverse sample behaviors in complex dynamics. Our approach pushes the envelope of visualizing conformational space of macromolecular complexes beyond the previously achieved scope toward the atomic level.

Results

Design of AlphaCryo4D framework. The conceptual framework of AlphaCryo4D integrates unsupervised deep learning with manifold embedding to learn a free-energy landscape, which

directs cryo-EM reconstructions of conformational continuum or transient states via an energy-based particle-voting algorithm. AlphaCryo4D consists of four major steps (Fig. 1a). First, hundreds to thousands of 3D volumes are bootstrapped with *M*-fold particle shuffling and Bayesian clustering using maximum-likelihood-based 3D (ML3D) classification algorithm[3,7]. To this end, all curated particles are aligned in a common frame of reference through a consensus 3D refinement in RELION. They are then randomly divided into *M*+1 groups of equal data sizes. In the step of '*M*-fold particle shuffling', one group of particles is taken out of the dataset to form a shuffled dataset. This procedure is repeated $M + 1$ times, each time with a different particle group being left out, resulting in $M + 1$ shuffled datasets (Extended Data Fig. 1). Each shuffled dataset is subject to 3D volume bootstrapping separately and is clustered into tens to hundreds of 3D reconstructions through Bayesian clustering in RELION[3,7]. In total, hundreds to thousands of volumes from all shuffled datasets are expected to be bootstrapped through these steps. Due to $M \geq 3$, this step demands a higher computational cost than conventional methods[2-4,6,7].

Next, all bootstrapped volumes are learned by a 3D autoencoder made of a deep residual convolutional neural network in an unsupervised fashion[43,44] (Fig. 1b, Extended Data Tables 1, 2). The 3D deep residual autoencoder are composed of six and five residual learning blocks in the encoder and decoder, respectively. Each residual learning block is derived from the ResNet architecture and consists of two convolutional neural network layers with shortcut connection skipping two adjacent layers, which has been shown to mitigate the degradation problem that learning accuracy gets saturated and then quickly degrades with increasing the network depth[43,44]. A 3D feature map is extracted for each volume by the deep residual autoencoder and is concatenated with the volume data for nonlinear dimensionality reduction by manifold embedding with the t-distributed stochastic neighbor embedding (t-SNE) algorithm[45] (Fig. 1d, e).

Third, the relative free-energy difference between 3D volumes is estimated using the particle number of each volume via the Boltzmann relation[35]. An energy landscape is visualized by mapping the values of energy difference on the learned manifold[35-41]. Note that only relative difference of free energy can be estimated, and the energy values do not represent the total energy of the system, which is not measurable by the Boltzmann relation. A string method can be optionally used to search the MEP on the energy landscape[46,47]. The local energy minima or transition states connecting adjacent minimum-energy states can be defined as the centers of 3D

clustering, with a circular range defined as the cluster boundary for subsequent particle voting (Fig. 1f).

Last, because each particle is used *M* times during volume bootstrapping, it is mapped to *M* locations on the energy landscape. The mapping of each copy of the particle is called a 'vote'. By counting the number of votes of the same particle casted within the same cluster boundary on the energy landscape, the reproducibility of machine learning procedure can be evaluated at single-particle level. Each particle is classified to the 3D cluster that receives more than *M*/2 votes of this particle within its voting boundary (Fig. 1f, Extended Data Fig. 1b). If none of the 3D clusters on the energy landscape receives more than *M*/2 votes of a given particle, the corresponding particle is voted out and excluded for further 3D reconstruction. The resulting 3D classes are expected to be conformationally homogeneous enough for high-resolution cryo-EM refinement.

Since many protein complexes exhibit profound conformational changes in different local regions, we also implemented a focused classification strategy of AlphaCryo4D that applies a local 3D mask[6] throughout the entire procedure, which is executed as an iterative step after initial 3D classification by AlphaCryo4D in the absence of any 3D mask (Fig. 1a).

Solving conformational continuum at atomic level. To assess the numerical performance of AlphaCryo4D, we generated three large synthetic heterogeneous cryo-EM datasets with signal-to-noise ratios (SNRs) of 0.05, 0.01 and 0.005. Each dataset includes 2 million randomly oriented single particles computationally simulated from 20 hypothetical conformer models of the ~130-kDa NLRP3 inflammasome protein[48]. These conformers imitate conformational continuum of the NACHT domain rotating around the LRR domain over an angular range of 90° during inflammasome activation[48] (Fig. 1c). The particles for each conformer are uniformly distributed in each dataset. We conducted blind assessments on 3D classification and heterogeneous reconstructions by AlphaCryo4D, without providing any information of particle orientations, translations and conformational identities (Fig. 1d-f, Extended Data Fig. 2). The 3D classification precision of a retrieved conformer was computed as the ratio of the particle number of correct class assignment (based on the ground truth) versus the total particle number in the class. The results were then compared with several alternative methods, including conventional maximum-likelihood-based 3D (ML3D) classification in RELION[3,4,6], 3D variability analysis

(3DVA) in cryoSPARC[2,11], and deep generative model-based cryoDRGN[12]. In all blind tests by presetting the class number to 20, AlphaCryo4D retrieved all 20 conformers and markedly outperformed the alternative methods, with an average of 3D classification precision at 0.83, 0.82 and 0.65 for datasets with SNRs of 0.05, 0.01 and 0.005, respectively (Fig. 1h, Extended Data Figs. 3a). By contrast, all alternative methods missed two to seven conformers entirely and exhibited 3D classification precisions in the range of 0.2-0.5 in general (Fig. 2i-k, Extended Data Fig. 3c-k). By increasing the preset class number to 25 or 30 that is more than the number of ground-truth conformers, all tested methods appear to be improved in the precision of 3D classification marginally but also reduced in the recall of classification (defined as the percentage of correctly assigned particles versus the number of ground-truth particles for a conformer) (Extended Data Fig. 4). In these cases, AlphaCryo4D still outperformed all alternative methods considerably, with the highest average classification precision reaching 0.91 at the SNR of 0.05.

The 3D classification precision appears to be strongly correlated with the map quality and resolution homogeneity across the density map (Fig. 1h-s, Extended Data Fig. 2h). All density maps from AlphaCryo4D consistently show homogeneous local resolutions (at 2.6-2.9 Å for SNR of 0.01) between the NACHT and LRR domains (Fig. 1l, p, Extended Data Fig. 2h, 3l, r). By contrast, all density maps by the alternative methods show lower average resolutions and notably heterogeneous local resolutions, with the NACHT domain exhibiting substantially lower resolution than that of the LRR domain, causing blurred features, broken loops and invisible sidechains in NACHT (Fig. 1m-o, q-s, Extended Data Fig. 3m-o). Thus, the significantly improved 3D classification accuracy by AlphaCryo4D enables 4D reconstruction of conformational continuum at the atomic level.

Having tested AlphaCryo4D on the simulated datasets with evenly distributed conformers, we further generated a simulated dataset of 0.01 SNR with a non-uniform distribution of the 20 conformers that follows a triangular wave function and presents a more challenging scenario (Extended Data Fig. 5a). Application of AlphaCryo4D on this dataset exhibited only slight reduction in 3D classification accuracy compared to the results on the datasets of evenly distributed conformers, and it still outperformed other methods tested on the uniformly distributed datasets. Moreover, the particle distribution computed by AlphaCryo4D appears to approximately recapitulate the triangular wave function (Extended Data Fig. 5b-f), suggesting

that AlphaCryo4D is potentially capable of reconstituting the overall statistical distribution of underlying conformational continuum.

To understand how the 3D classification accuracy is improved, we further studied the algorithmic mechanism by ablating or replacing certain components of AlphaCryo4D in 24 control tests (Extended Data Figs. 2d-g, 6; see Methods for details). By removing the component of deep manifold learning from AlphaCryo4D, the statistical distributions of high-precision 3D classes after particle voting are reduced by ~60% (Extended Data Fig. 6d-f). Similarly, classifying particles without particle voting or in the absence of deep residual autoencoder reduces the high-precision 3D classes by ~15-30% at a lower SNR (Extended Data Fig. 6m-r). Further, replacing t-SNE with other four dimensionality reduction techniques resulted in the loss of at least half of ground-truth conformers in 3D classification, including PCA, multidimensional scaling (MDS)[49], isometric mapping[50] and locally linear embedding[51] (Extended Data Fig. 2d-g). Throughout AlphaCryo4D, the steps of particle shuffling, defining voting boundaries on energy landscapes and energy-based particle voting contribute to ~16%, ~20% and ~40% improvements of high-precision 3D classes, respectively (Extended Data Fig. 6g-l, s-u). Taken together, these results indicate that all components of AlphaCryo4D contribute to the improved accuracy of 3D classification.

Applications to various experimental datasets. Having conducted the proof-of-principle study of AlphaCryo4D using simulated datasets, we then examined the applicability of AlphaCryo4D on a number of experimental datasets.

***Pf*80S ribosome.** We first evaluated the capability of AlphaCryo4D in detecting residual heterogeneity in a mostly homogeneous dataset of the Plasmodium falciparum 80S (*Pf*80S) ribosome bound to anti-protozoan drug emetine (EMPIAR-10028)[52]. The original cryo-EM analysis of the *Pf*80S ribosome dataset (105,247 particles) determined a major conformational state and detected flexibility in the head region of the 40S small subunit, which was originally unsolved[52]. Application of AlphaCryo4D simultaneously identified two 40S-rotated states (R3 and R4) and another state (R2) with missing 40S head in addition to the original conformational state (R1) located in the deepest energy well (Fig. 2a, b). While the major state was retained at the original resolution of 3.2 Å, three new states R2, R3 and R4 were refined to 3.6, 4.1 and 4.6 Å, respectively, without using more data (Extended Data Fig. 7a-g). Superimposing the density

maps of states R3 and R1 reveals that the rotation of 40S is localized within its head region with concerted movement in the central protuberance of 60S, whereas the rest of the 40S exhibits no obvious movement (Fig. 2d). By contrast, the entire 40S rotates above the 60S in state R4 relative to R1; this coincides with the structural rearrangement within 60S near the central protuberance and L1 stalk, which appears to also drive the 40S translation relative to the 60S (Fig. 2e), indicating a different mode of intersubunit rotation. Moreover, the reconstructions of R3 and R4 also identified differential rearrangements of small structural elements, including the stepwise disappearance of an rRNA helix and of the intersubunit bridge formed by eL8 C-terminal helix, in line with previous studies on the *Pf*80S dynamics during translation[53] (Fig. 2c). Given that the *Pf*80S ribosome was bound to the anti-protozoan drug visible in all new states (Extended Data Fig. 7e), the AlphaCryo4D-enabled analysis of structural dynamics at high resolution opens the door to investigate the drug-modulated allosteric effect relevant to therapeutic improvement[52].

Pre-catalytic spliceosome. We next evaluated AlphaCryo4D on the pre-catalytic spliceosome dataset (EMPIAR-10180)[54]. Previous study using focused classification has provided a multistate-averaged map of the pre-catalytic spliceosome and speculated that the SF3b subcomplex might sample a continuum of conformations[54]. Application of AlphaCryo4D yielded nine intermediate states (S1-S9) characterizing a continuous motion of the SF3b and helicase subcomplexes, as well as several compositionally distinct states (S10-S14) (Fig. 2f-h, Extended Data Fig. 7h-l). The nine states S1-S9 are intuitively located in adjacent clusters, forming a major MEP traversing the long axis of the energy landscape (Fig. 2g). These states show both independent movements of the SF3b and helicase, as well as concerted motions of the two subcomplexes (Extended Data Fig. 7j). Comparison of states S3 and S4 exhibits that the SF3b and helicase move away from each other, whereas states S5 and S6 show co-directional motions of both the SF3b and helicase. The Cus1 subunit in the SF3b subcomplex appears to switch its contact between the N-terminal helicase cassettes of RecA-1 and RecA-2, implicating that part of conformational rearrangement is mediated by the transient interface B between the SF3b and helicase[54]. Notably, comparison of states S4 and S13 exhibits that the motions of U2 snRNP and helicase appear to be allosterically coupled to the disappearance of the Spp381, Prp38 and Snu23 densities in state S13, the three subunits required for spliceosome activation[54] (Fig. 2h). These dynamic features important to understand intermediate steps of spliceosome activation have not

been previously reported from the same dataset[54], demonstrating the applicability of AlphaCryo4D in simultaneously visualizing continuous motion and discrete compositional heterogeneity on a dataset with a moderate size.

RyR1 channel complex. We applied AlphaCryo4D to solve the dynamic domains of the 2.2-MDa ryanodine receptor (RyR), a tetrameric calcium-activated channel protein complex that mediates the release of calcium ions from intracellular stores. The mammalian RyR1 has been determined at near-atomic resolution in its central core in multiple conformations[55,56], which was recently further improved to sub-3 Å resolution[57]. However, approximately half of the entire RyR1 structure located at peripheral domains remain at moderate to low resolution due to their complex dynamics that precludes de novo atomic modeling[55-57]. Application of AlphaCryo4D disentangled the motion of peripheral domains and improved its local resolution to 3.3 Å (Fig. 2i-k, Extended Data Fig. 7m-s). The resulting density map showed high-resolution features including helical pitches and side-chain densities (Fig. 2k). Thus, the application to RyR1 demonstrates that AlphaCryo4D is capable of breaking the resolution barrier of the dynamic domain by resolving hierarchically entangled conformational dynamics of the RyR1 peripheral domains.

Assembly intermediates of bacterial ribosome. We used AlphaCryo4D to analyze a highly heterogeneous dataset of the *Escherichia coli* 50S large ribosomal subunit (LSU) (EMPIAR-10076)[58]. This dataset (131,899 particles) is known to be both compositionally heterogeneous and conformationally dynamic in that the 50S LSU undergoes bL17-depleted intermediate assembly. We first computed the overall energy landscape of the ribosomal assembly using 119 density maps (Fig. 3a). On this energy landscape, the originally reported five major states A-E were straightforwardly reproduced by energy-based particle voting. Obviously, state A containing both 50S and 30S subcomplexes is a rare state distinct from other conformers (Fig. 3b). As the remaining particle numbers of states B, C, D and E were all adequate for deeper classification, we then computed their zoomed-in energy landscapes separately. This allowed us to discover seven new conformational states (designated B1, B2, C5, D5, E6, E7, E8) that have evaded all other methods[11,12,19,58] (Fig. 3), as well as all 13 previously reported sub-class conformers classified from this dataset[12,58] (Extended Data Fig. 8). Notably, state B1 exhibits a low occupancy of domain III and another new state D5 is perhaps a least mature LSU intermediate on the assembly pathway (Fig. 3e, i). The rRNA helix 101 exhibits a small rotation

between states B2 and B3 (Fig. 3e and Extended Data Fig. 8b). The differential occupancy of rRNA helix 63 and the movement of uL1 stalk are observed in states C5 and E6-E8 (Fig. 3g, k). Among the sub-states derived from state E, the new state (E7) with uL1 missing was reconstructed by only 0.3% of the entire dataset. These results suggest that AlphaCryo4D is highly efficient in finding transient states of extremely low populations and in exploring the 'dark' conformers in the conformational space of highly heterogeneous systems.

Visualizing hidden dynamics of the proteasome. Having examined the general applicability of AlphaCryo4D on several experimental datasets, we then turned to two experimental cryo-EM datasets of human proteasome[30,59] to examine whether AlphaCryo4D can make novel biological discoveries. The AlphaCryo4D-reconstituted energy landscape using the substrate-free 26S proteasome dataset (EMPIAR-10090) shows six local energy minima corresponding to the known six conformational states[59], verifying the applicability of AlphaCryo4D on analyzing proteasome dynamics (Fig. 4a). The second dataset (EMPIAR-10669) was collected on the human proteasome mixed with a Lys63-linked polyubiquitinated substrate $Sic1^{PY}$ for 30 seconds before ATP was diluted with ATPγS, which was expected to maximally capture any intermediate states before the degradation reaction was completed[30]. The reconstruction of the energy landscape suggests that the rate-limiting step of substrate degradation lies in the initiation of substrate unfolding after deubiquitylation, reflected in the highest activation energy barrier from state E_B to E_{C1} (Fig. 4b, c), consistent with a recent measurement of proteasome kinetics[60].

Impressively, AlphaCryo4D extracted significantly more proteasome conformers in various compositional forms from the same dataset that previously yielded seven atomic structures of the proteasome[30]. Not only it discovered 20 conformers of regulatory particle (RP)-core particle (CP) subcomplex, but also considerably improved the resolutions of key states, such as E_{A2} and E_{D2}, due to its advantage in keeping more particles without sacrificing conformational homogeneity (Extended Data Figs. 9-12, Extended Data Table 3). It also yielded 8 and 36 conformers of singly capped (SC) and doubly capped (DC) holoenzymes, respectively, with the majority refined to 3.7-4.7 Å (Extended Data Figs. 11, 12a, b, Supplementary Information).

Conformational entanglement. Both the DC and SC proteasomes are abundant in cells[61]. Each CP gate in the DC proteasome appears to be rigorously controlled and is only open when its

proximal RP is in an E_D-compatible state with five RPT C-terminal tails inserted into the surface pockets on the α-ring[30] (Fig. 4e, Extended Data Fig. 11a-c). Consistently, CP gate opening on the RP-controlled side does not remotely open the other CP gate on the RP-free side of the SC proteasome (Extended Data Fig. 11d).

To understand whether the two RP conformations in the same DC proteasome allosterically influence each other, we compare the state distribution matrix obtained by AlphaCryo4D to the prediction of a control model assuming two completely uncoupled RP states in the DC proteasome (Fig. 4f-h). A notable number of DC states exhibit substantial deviation from the model prediction, indicating that the conformational dynamics of the two RPs in the same proteasome are profoundly entangled together (Fig. 4g, h). Intriguingly, the most prominent positive deviations are seen in the state distributions of E_{A1}-E_{A1}, E_B-E_B, E_B-E_{D2}, and E_{D2}-E_{D2} combinations (Fig. 4g), suggesting that the deubiquitylation state E_B and the translocation state E_{D2} of one RP are allosterically coupled to and mutually upregulated by states E_B and E_{D2} of the opposite RP. Coupled actions of deubiquitylation and translocation between two opposite RPs are expected to optimize the efficiency of substrate processing. Moreover, all DC states along the diagonal of the state distribution matrix are more or less upregulated (indicated by red in Fig. 4g, h), suggesting that both RPs in the same proteasome synergistically promote each other to transit into downstream states in a symmetrical fashion.

Notably, the largest discrepancy of the distribution of RP states between the SC and DC proteasomes also lies in state E_B (Fig. 4b-d). The state population of E_B is significantly lower in the SC than in the DC proteasome. This observation verifies the conclusion drawn from the above analysis on the DC proteasome alone that state E_B is upregulated by the conformational entanglement effect in the DC holoenzyme, suggesting that the SC proteasome lacks the functional synergy between two opposite RPs in the DC proteasome (Fig. 4h). Altogether, these results implicate that the DC proteasome could be energetically more efficient in substrate deubiquitylation and degradation than the SC proteasome.

Ubiquitin-proteasome interactions. Ten states, designated $E_{A1.0}$, $E_{A1.1}$, $E_{A1.2}$, $E_{A2.2}$, $E_{A2.3}$, $E_{B.2}$, $E_{B.3}$, $E_{D2.2}$, $E_{D2.3}$ and $E_{D.\alpha5}$, at 3.1-7.5 Å resolution exhibit features resembling ubiquitin that map various putative ubiquitin-binding sites on the RPN1, RPN2, RPN10 and α5 subunits in the proteasome (Fig. 4i, Extended Data Fig. 9p, u). Except states $E_{A1.0}$, $E_{A1.1}$ and $E_{D.\alpha5}$, the

populations of these states are very small and represented in only a few thousand particles or 0.1-0.3% of the entire dataset, indicating their transient nature (Fig. 4i). The local resolutions of the RPN2- and RPN10-bound ubiquitin-like densities are lower than the rest of the reconstructions and other ubiquitin densities on RPN1, RPN11 and α5 subunits. Rigid-body fitting of the ubiquitin structure into these ubiquitin-like densities confirms the agreement of the overall density shape with the ubiquitin fold (Fig. 4i, Extended Data Fig. 10b-g).

RPN1 sites. By focused 3D classification via AlphaCryo4D, we refined the density map of RPN1 in state $E_{A1.0}$ (~5.9% of the dataset) to 3.7 Å, which was previously reconstructed at moderate resolution[30] (Extended Data Figs. 9n, 10d). It reveals that a ubiquitin moiety contacts Glu458, Cys459 and Asp460 at the N-terminal end of an RPN1 helix (homologous to H24 between the T1 and T2 sites of yeast Rpn1) and potentially interacts with the T2 site (Asp423, Leu426 and Asp430) on the RPN1 toroidal domain (Fig. 4k, l)[62]. Because the yeast Rpn1 T2 site was found to preferentially bind Ubp6 ubiquitin-like (UBL) domain[62], this site may be sterically occluded for ubiquitin binding in the presence of USP14/Ubp6. By contrast, state $E_{A1.1}$ exhibits one ubiquitin moiety on the RPN1 T1 site and the other near the T2 site, which together can be modelled as a Lys63-linked diubiquitin (Extended Data Fig. 10e). The much smaller population of state $E_{A1.1}$ (~1.3 % of the dataset) is consistent with the previous finding that Lys63-linked polyubiquitin chains have a lower affinity against the T1 site than those linked via Lys48 and Lys6 residues[62].

RPN2 site. In states $E_{A1.2}$, $E_{A2.2}$, and $E_{B.2}$, a diubiquitin-like density is consistently found on RPN2. The ubiquitin-binding site on RPN2 is centered on three helices harboring an acidic surface around residues Tyr502, Glu530, Asp531 and Ser569 (Fig. 4j-m). Structural comparison suggests that the ubiquitin-binding site on RPN2 is homologous to the T1 site of RPN1 (Fig. 4k, l)[62]. We therefore name it 'the RPN2 T1 site'. Comparison with the NMR structure of Lys48-linked diubiquitin-bound RPN1 T1 site helices[62] further confirms the conservation and similarity of the ubiquitin-binding modes between the RPN1 and RPN2 T1 sites (Fig. 4k, l, Extended Data Fig. 10h-k). Consistently, previous studies have detected interactions of UBL proteins and RPN2 by chemical crosslinking[63]. Furthermore, a Lys63-linked triubiquitin chain can be modelled into the RPN2- and RPN11-bound ubiquitin densities by rotating Lys63 linkages, suggesting that the RPN2 T1 site may function as a coreceptor site for ubiquitin transfer from RPN10 to RPN11 to facilitate deubiquitylation (Fig. 4o). Given that at least one RPN10 ubiquitin-interacting motif

(UIM)-bound ubiquitin invisible in these cryo-EM maps presumably exists, the observation of triubiquitin on RPN2 and RPN11 structurally rationalizes why tetraubiquitin is a preferential degradation signal for the proteasome[64].

RPN10 VWA site. In states $E_{A2.3}$, $E_{B.3}$ and $E_{D2.3}$, a monoubiquitin-like density consistently resides on Lys132, Lys133 and Lys163 of the RPN10 von Willebrand factor type A (VWA) domain (Fig. 4i, Extended Data Fig. 10b-f). By contrast, in state $E_{D2.2}$, a diubiquitin-like density appears to span from this lysine-rich site of RPN10 to the RPN2 T1 site (Extended Data Fig. 10b, g). In line with these observations, previous studies have suggested that the yeast Rpn10 VWA is a ubiquitin-binding domain and ubiquitylation of Rpn10 VWA itself prevents Rpn10 from being incorporated into the proteasome[65-68]. As ubiquitylation of RPN10 appears to regulate substrate recruitment[65-68], it remains to be clarified whether the putative ubiquitin at this lysine-rich site is covalently modified and co-purified from HEK293 cells.

α5-subunit site. In state $E_{D2.2}$, a ubiquitin density is observed on the RPN1 T2 site and another weaker ubiquitin density on the α5-subunit ~6 nm beneath the RPN1 T2 site (Fig. 4i, Extended Data Fig. 10b). To understand if the ubiquitin-binding site on the α5 subunit is used more broadly, we combined all the E_D-compatible states and focused AlphaCryo4D classification on the α5-bound ubiquitin. This yielded a much larger 3D class, named state $E_{D.\alpha5}$ (~4% of the entire dataset) that improved the CP and the α5-bound ubiquitin to 3.1 and 4.9 Å, respectively (Extended Data Figs. 9o, p, r, u, 10f). The ubiquitin-binding site is composed of acidic or hydrophobic residues Glu183, Val184, His186, Ser188 and Glu193 in a loop connecting two helices, resembling the mode of ubiquitin recognition by an RPN13 loop[69,70] (Fig. 4n, Extended Data Fig. 10f), and electrostatically similar to the RPN1 T1/T2 and RPN2 T1 sites (Extended Data Fig. 10h-o). Recent data suggest that the free 20S CP can degrade substrate-conjugated ubiquitin[71], which might be facilitated by ubiquitin binding to the α-ring.

The low-affinity ubiquitin-binding sites all reside in the vicinity of primary receptor sites. The RPN2 T1 and α5 sites are very close to the RPN10 UIMs and RPN1 T2 site, respectively. Modelling of a Lys63-linked tetraubiquitin chain by docking its terminal ubiquitin into the cryo-EM density suggests that a tetraubiquitin can span across the α5 and RPN1 T2 sites (Fig. 4p). Given that RPN1 appears to allosterically regulate the AAA-ATPase motor conformation[30] (Fig. 5a, b), tetraubiquitin binding to both the RPN1 T2 site and α5 subunit could allosterically stabilize the open-CP states (E_{D0}, E_{D1} and E_{D2}). Substrate-conjugated polyubiquitin chains could

be first captured by the high-affinity sites and then moved to the nearby low-affinity sites. Working together, these sites could remodel polyubiquitin chains and deliver the peptide-proximal ubiquitin to the deubiquitylation site[30] (Fig. 4o) as well as assist ubiquitin recycling after isopeptide bond hydrolysis (Fig. 4p). Due to the transient nature of these ubiquitin-bound states, we cannot rule out that some of the ubiquitin-like densities are co-purified endogenous ubiquitin, UBL proteins or ubiquitylation of the proteasome itself[64-68,72]. Future functional studies are required to understand their definitive chemical nature and biological importance in cells.

Single-nucleotide-exchange dynamics. The six intermediate states ($E_{D0.1}$, $E_{D0.2}$, $E_{D0.3}$, $E_{D1.1}$, $E_{D1.2}$ and $E_{D1.3}$) during initiation of substrate unfolding exhibit highly coordinated movements in the RP (Fig. 5a-c, Extended Data Fig. 13), with prominent rotations in the RPN1 N-terminal that detaches from RPT2 in states $E_{D0.1}$, $E_{D0.2}$, and $E_{D0.3}$ and re-binds RPT2 in state $E_{D1.1}$, $E_{D1.2}$ and $E_{D1.3}$. All six conformations share the same architecture in their substrate interactions with four RPT subunits (RPT2, RPT6, RPT3 and RPT4) and in the CP gate (Fig. 5d). The CP gate is open once the pore-1 loop of RPT2 binds the substrate in state $E_{D0.1}$, as five RPT C-tails are inserted in the α-pockets in all six states (Extended Data Fig. 11c). RPT2 re-association with the substrate proceeds that of RPT1 after both RPT1 and RPT2 are dissociated from the substrate in state E_C. They also share the same pattern of nucleotide states, with ATP bound to RPT1, RPT2, RPT6 and RPT3 and ADP bound to RPT4. Poor nucleotide-like densities in the RPT5 nucleotide-binding pocket give rise to the highest B-factor (250-500) when fitted with ADP, suggesting that RPT5 undergoes ADP release (Extended Data Fig. 13k).

Markedly, the pore-1 loop of RPT1 gradually moves toward the substrate from a distance of ~16 Å to ~3 Å, whereas the pore-1 loop of RPT5 moves from ~3 Å to ~16 Å away from the substrate (Fig. 5d, f). The small AAA subdomain of RPT5 moves inward together with the large AAA subdomain of RPT1, as the large AAA subdomain of RPT5 rotates away from the substate (Extended Data Fig. 13f-i). Consistently, the arginine fingers (R-finger) of RPT2 are progressively moved toward ATP bound in RPT1, closing up the nucleotide-binding pocket (Fig. 5e, g). In contrast, the nucleotide-binding pocket of RPT5 is gradually shrunk but is not fully closed by the R-fingers of RPT1, whereas the nucleotide-binding pocket of RPT4 is gradually opened up (Fig. 5g). As a result, the pore-1 loops of RPT5 and RPT1 reach the substate only in

the first and last states $E_{D0.1}$ and $E_{D1.3}$, respectively. Interestingly, the outward flipping of RPT5 appears to complete in state $E_{D0.3}$, much faster than the inward motion of RPT1, which completes in state $E_{D1.3}$, suggesting that substrate re-engagement is the rate-limiting step in the single-nucleotide exchange kinetics of the AAA-ATPase motor. Although sequential hydrolysis around the AAA-ATPase ring has been hypothesized in numerous similar systems[30,64,73,74], our results provide the first direct evidence for sequential intermediates of hand-over-hand actions within a single-nucleotide-exchange cycle (Fig. 5h).

Hierarchical allosteric regulation of proteolytic activity. Both atomic structures of the substrate-bound CPs in the open- and closed-gate states E_A and E_{D2} show a segment of substrate polypeptide binding the catalytically active residue Thr1 in the β2 subunit, whereas no substrate density is observed in the other two catalytic subunits β1 and β5. In both states, the substrate is bound in a β-sheet conformation between residues 20 to 22 and 46 to 48 in the β2 subunit (Fig. 6a, b, Extended Data Fig. 12e, f), reminiscent of the conformations of CP inhibitors binding the β5 subunits in the crystal structures of the CP[75,76]. The hydroxyl oxygen atom in the catalytic Thr1 is ~2-Å closer to the substrate mainchain in state E_{D2} compared to state E_A.

To understand how such a subtle structural change influences the proteolytic activity, we carried out first-principles calculations using the quantum-mechanical density functional theory (DFT)[77]. To initiate the nucleophilic attack, the carbonyl group of the scissile peptide bond must approach the electrophile Thr1-Oγ to facilitate a tetrahedral transition state[76,78,79]. We calculated the charge density difference induced around the interacting pair after proton transfer from Thr1-Oγ to Thr1-NH_2 (Fig. 6c, d, Extended Data Fig. 12g, h). The substrate-interacting pair of atoms in state E_{D2} form a bonding region of slight positive charge difference linking the carbonyl group of the substrate to Thr1-Oγ. By contrast, the pair are separated by a charge depletion non-bonding region in state E_A. Indeed, a recent study has found that the nucleophilic attack is the rate-limiting step in the acylation barrier[78]. Our result suggests that CP gate opening allosterically promotes the nucleophilic attack by lowering the activation energy barrier of acylation and induces hybridization of molecular orbitals, hence upregulating the reaction rate of proteolysis. In support of our structural finding, previous studies have observed indirect evidence for coupling of the CP gating with the proteolytic activity[36,80-82].

Taken together, our analyses point to a grand allostery model for the proteasomal autoregulation through a three-tier hierarchy of inter-molecular interaction networks (Fig. 6e). The first tier is the upregulation of the proteolytic activity of the β-subunits by CP gate opening. The second tier lies in the AAA-ATPase motor bridging the ubiquitin-binding sites and the CP via lid-base and RP-CP interactions. Three modes of coordinated ATP hydrolysis in the AAA-ATPase ring regulate intermediate steps of substrate processing[30,64], in which initiation of substrate unfolding is allosterically coupled with CP gating during E_{C2} to $E_{D0.1}$ transition (Fig. 5). The third tier is via dynamic interactions between ubiquitin signals and a collection of ubiquitin receptor sites and deubiquitinases. This is manifested as the dual roles of ubiquitin-binding sites in recognizing ubiquitin signals and in transmitting such recognitions to allosteric regulation of AAA-ATPase conformations that control the CP gate. Ubiquitin binding destabilizes the resting state of the proteasome and promotes its transition to more active states. On the other hand, a polyubiquitin chain binding multiple proteasomal sites could have profound conformational selection effects (Fig. 4o, p). Such three-tier, hierarchical allosteric regulations may occur mutually with both RPs and underlie the conformational entanglement of the two RPs in the DC proteasome (Fig. 4e), which was not clearly observed in the absence of substrates[59,64,83]. This grand allostery might have been evolved to optimize the proteasome function against extremely complicated intracellular challenges and underpin various phenotypic effects, such as UBL-induced proteasome activation[84,85] and long-range allosteric couplings in the proteasome[36,64,81,86,87].

Discussion

One primary objective of the AlphaCryo4D development is to simultaneously improve reconstruction resolution and 3D classification accuracy without assumption on sample behaviors, such as conformational heterogeneity, continuity or chemical composition. We demonstrate that such a goal is achievable in potentially broad applications. As a highlighted example, AlphaCryo4D-based analysis of a large experimental dataset has enabled unprecedented novel insights into hidden dynamics and allosteric complexity of proteasome autoregulation that are essential for understanding proteasome function and have evaded previous investigations[30]. Our approach expands the realm of structural ensemble determination to the nonequilibrium regime at atomic level, thus potentially transforming biomedical research

and therapeutic development. Despite this advancement, the structural ensemble discovered by AlphaCryo4D should be regarded as discrete approximations visualizing the conformational space of fundamentally unknown continuity.

A key lesson from this study is that the improved accuracy of 3D classification is often positively correlated with the resolution gain in general, as misclassification is presumably a key factor limiting achievable resolution for highly dynamic conformations. Because of extremely low SNR, most objective functions in image similarity measurement and machine learning algorithms tend to poorly perform or entirely fail on cryo-EM data. To date, no methods have included any implicit metrics or strategies for validating 3D classification accuracy in cryo-EM. We address this issue by introducing *M*-fold particle shuffling method and the energy-based particle voting algorithm in AlphaCryo4D that automatically checks the reproducibility and robustness of the 3D classification. This allows users to objectively reject non-reproducible particles and permits a maximal number of particles to be assessed and classified in an integrative procedure based on uniform, objective criteria. Our extensive tests suggest that the particle-voting algorithm synergistically enhances the performance of deep manifold learning and potentially rescues certain particles that are prone to be misclassified, thus boosting the efficiency of particle usage without necessarily sacrificing quality and homogeneity of selected particles. This mechanism simultaneously optimizes the usage of all available particles and their 3D classification accuracy for achieving higher resolution.

Using energy landscape to represent the conformational variation provides several advantages. First, it has theoretical roots in physical chemistry and protein dynamics, in which extensive research tools being built around the energy landscape and transition-state theory allow correlation with other complementary approaches, such as single-molecule florescence microscopy, NMR and molecular dynamics simulation. Second, the energy landscape allows users to intuitively examine kinetic relationships between the adjacent conformers and to potentially discover new conformations. Our approach to use many bootstrapped density maps to estimate the energy landscape is statistically different from other methods proposed to map single particles to continuous distributions in the latent space[11,12,19-22]. The latter lacks energetic information, which can restrict the interpretation of the resulting reconstruction. Third, while our method of mapping 3D volumes to manifolds increases the sampling grain in approximating the conformational space as compared to other methods directly mapping single particles to the

latent space[11,12,19], AlphaCryo4D trades off the fineness of sampling grain in visualizing conformational space for resolution gain. In nearly all cases we examined, AlphaCryo4D apparently still oversampled the conformational space at large in terms of finding those conformational states that can potentially achieve high resolution.

The limitations of AlphaCryo4D lie in several aspects. First, it requires generally larger datasets and more computational costs to fully exploit its advantages and potentials (Extended Data Table 2). The data size must proportionally scale with and match the degree of conformational heterogeneity in order to visualize a full spectrum of conformational states at the atomic level. Further optimization of the code-executing efficiency via software engineering is required to reduce the computational costs and to make AlphaCryo4D more affordable for widespread applications.

Second, its outcomes depend on the success of initial consensus alignment of all particles during particle shuffling and volume bootstrapping. Certain conformational dynamics or heterogeneity can interfere with image alignment, in which case the performance of AlphaCryo4D will be restricted by alignment errors after consensus refinement. This is a common problem for all existing methods developed so far. Our tentative solution is to improve image alignment in a multi-reference refinement procedure. Failure in obtaining accurate alignment parameters can lead to futile, erroneous estimation of conformational heterogeneity in the subsequent steps no matter which methods are used.

Third, the use of t-SNE to visualize energy landscape poses both potential advantages and limitations. The t-SNE algorithm can well preserve the local topology of energy landscape but is not warranted to preserve global topology, which is nonetheless a common limitation for existing manifold learning techniques. The geometry of estimated energy landscape is inevitably influenced by the SNRs of cryo-EM data; lower SNRs cause dispersion of local energy wells and thus the overall geometrical distortion of the landscape (Extended Data Fig. 2). Higher level of noise introduces greater statistical bias in particle clustering into bootstrapped 3D volumes and causes irreproducible particle “votes”. To alleviate this issue, the particle-voting algorithm is employed in final 3D classification at the local areas of energy minima on energy landscape, which rejects particles with irreproducible votes.

Last, all existing approaches for visualizing energy landscape from cryo-EM data suffer from the “curse of dimensionality”, because the dimensions of reaction coordinates related to the

conformational variation are expected to be very high in real space for a megadalton macromolecular complex that includes hundreds of thousands of atoms[35-41]. Thus, reduction of dimensions prior to visualization of energy landscape is necessary. The reduced dimensions then lose intuitive physical meanings that differ from case to case. In all cases of experimental and simulated datasets in this study, the reduced dimensions by t-SNE appear to capture the most prominent conformational changes, analogous to the principal components computed by PCA. Thus, dimensionality reduction to 2D by t-SNE, as well as MEP searching on a 2D energy landscape might limit its capability in disentangling complex dynamics governed by many reaction coordinates. Further investigations are required to understand how to map individual particles to energy landscape without trading off reconstruction resolution, to estimate energy landscape at higher dimensions, to preserve its global topology with high fidelity and to fully automatize MEP solution.

References

1 Scheres, S. H. A Bayesian view on cryo-EM structure determination. *J Mol Biol* **415**, 406-418, doi:10.1016/j.jmb.2011.11.010 (2012).

2 Punjani, A., Rubinstein, J. L., Fleet, D. J. & Brubaker, M. A. cryoSPARC: algorithms for rapid unsupervised cryo-EM structure determination. *Nat Methods* **14**, 290-296, doi:10.1038/nmeth.4169 (2017).

3 Zivanov, J. *et al.* New tools for automated high-resolution cryo-EM structure determination in RELION-3. *Elife* **7**, doi:10.7554/eLife.42166 (2018).

4 Nakane, T., Kimanius, D., Lindahl, E. & Scheres, S. H. Characterisation of molecular motions in cryo-EM single-particle data by multi-body refinement in RELION. *Elife* **7**, doi:10.7554/eLife.36861 (2018).

5 Wu, J. *et al.* Massively parallel unsupervised single-particle cryo-EM data clustering via statistical manifold learning. *PLoS One* **12**, e0182130, doi:10.1371/journal.pone.0182130 (2017).

6 Scheres, S. H. Processing of Structurally Heterogeneous Cryo-EM Data in RELION. *Methods Enzymol* **579**, 125-157, doi:10.1016/bs.mie.2016.04.012 (2016).

7 Scheres, S. H. *et al.* Disentangling conformational states of macromolecules in 3D-EM through likelihood optimization. *Nat Methods* **4**, 27-29, doi:10.1038/nmeth992 (2007).

8 Galves, M., Rathi, R., Prag, G. & Ashkenazi, A. Ubiquitin Signaling and Degradation of Aggregate-Prone Proteins. *Trends Biochem Sci*, doi:10.1016/j.tibs.2019.04.007 (2019).

9 Nakane, T. *et al.* Single-particle cryo-EM at atomic resolution. *Nature* **587**, 152-156, doi:10.1038/s41586-020-2829-0 (2020).

10 Yip, K. M., Fischer, N., Paknia, E., Chari, A. & Stark, H. Atomic-resolution protein structure determination by cryo-EM. *Nature* **587**, 157-161, doi:10.1038/s41586-020-2833-4 (2020).

11 Punjani, A. & Fleet, D. J. 3D Variability Analysis: Resolving continuous flexibility and discrete heterogeneity from single particle cryo-EM. *J Struct Biol*, 107702,

doi:10.1016/j.jsb.2021.107702 (2021).
12 Zhong, E. D., Bepler, T., Berger, B. & Davis, J. H. CryoDRGN: reconstruction of heterogeneous cryo-EM structures using neural networks. *Nat Methods* **18**, 176-185, doi:10.1038/s41592-020-01049-4 (2021).
13 Penczek, P. A., Kimmel, M. & Spahn, C. M. Identifying conformational states of macromolecules by eigen-analysis of resampled cryo-EM images. *Structure* **19**, 1582-1590, doi:10.1016/j.str.2011.10.003 (2011).
14 Tagare, H. D., Kucukelbir, A., Sigworth, F. J., Wang, H. & Rao, M. Directly reconstructing principal components of heterogeneous particles from cryo-EM images. *J Struct Biol* **191**, 245-262, doi:10.1016/j.jsb.2015.05.007 (2015).
15 Sanchez Sorzano, C. O., Alvarez-Cabrera, A. L., Kazemi, M., Carazo, J. M. & Jonic, S. StructMap: Elastic Distance Analysis of Electron Microscopy Maps for Studying Conformational Changes. *Biophys J* **110**, 1753-1765, doi:10.1016/j.bpj.2016.03.019 (2016).
16 Katsevich, E., Katsevich, A. & Singer, A. Covariance Matrix Estimation for the Cryo-EM Heterogeneity Problem. *SIAM J Imaging Sci* **8**, 126-185, doi:10.1137/130935434 (2015).
17 Anden, J., Katsevich, E. & Singer, A. Covariance Estimation Using Conjugate Gradient for 3d Classification in Cryo-Em. *Proc IEEE Int Symp Biomed Imaging* **2015**, 200-204, doi:10.1109/ISBI.2015.7163849 (2015).
18 Lyumkis, D., Brilot, A. F., Theobald, D. L. & Grigorieff, N. Likelihood-based classification of cryo-EM images using FREALIGN. *J Struct Biol* **183**, 377-388, doi:10.1016/j.jsb.2013.07.005 (2013).
19 Chen, M. & Ludtke, S. J. Deep learning-based mixed-dimensional Gaussian mixture model for characterizing variability in cryo-EM. *Nat Methods* **18**, 930-936, doi:10.1038/s41592-021-01220-5 (2021).
20 Gupta, H., Phan, T. H., Yoo, J. & Unser, M. Multi-CryoGAN: Reconstruction of Continuous Conformations in Cryo-EM Using Generative Adversarial Networks. *Computer Vision – ECCV 2020 Workshops* 429-444 (Springer International Publishing).
21 Moscovich, A., Halevi, A., Anden, J. & Singer, A. Cryo-EM reconstruction of continuous heterogeneity by Laplacian spectral volumes. *Inverse Probl* **36**, doi:10.1088/1361-6420/ab4f55 (2020).
22 Lederman, R. R., Andén, J. & Singer, A. Hyper-molecules: on the representation and recovery of dynamical structures for applications in flexible macro-molecules in cryo-EM. *Inverse Problems* **36**, 044005, doi:10.1088/1361-6420/ab5ede (2020).
23 Bai, X. C., Rajendra, E., Yang, G., Shi, Y. & Scheres, S. H. Sampling the conformational space of the catalytic subunit of human gamma-secretase. *Elife* **4**, doi:10.7554/eLife.11182 (2015).
24 Ilca, S. L. *et al.* Localized reconstruction of subunits from electron cryomicroscopy images of macromolecular complexes. *Nat Commun* **6**, 8843, doi:10.1038/ncomms9843 (2015).
25 Jonic, S. & Sorzano, C. O. Versatility of Approximating Single-Particle Electron Microscopy Density Maps Using Pseudoatoms and Approximation-Accuracy Control. *Biomed Res Int* **2016**, 7060348, doi:10.1155/2016/7060348 (2016).
26 Sorzano, C. O. S. *et al.* Survey of the analysis of continuous conformational variability of biological macromolecules by electron microscopy. *Acta Crystallogr F Struct Biol*

Commun **75**, 19-32, doi:10.1107/S2053230X18015108 (2019).
27 Schilbach, S. *et al.* Structures of transcription pre-initiation complex with TFIIH and Mediator. *Nature* **551**, 204-209, doi:10.1038/nature24282 (2017).
28 Punjani, A. & Fleet, D. J. 3D Flexible Refinement: Structure and Motion of Flexible Proteins from Cryo-EM. *bioRxiv*, 2021.2004.2022.440893, doi:10.1101/2021.04.22.440893 (2021).
29 Punjani, A., Zhang, H. & Fleet, D. J. Non-uniform refinement: adaptive regularization improves single-particle cryo-EM reconstruction. *Nature Methods* **17**, 1214-1221, doi:10.1038/s41592-020-00990-8 (2020).
30 Dong, Y. *et al.* Cryo-EM structures and dynamics of substrate-engaged human 26S proteasome. *Nature* **565**, 49-55, doi:10.1038/s41586-018-0736-4 (2019).
31 Wu, Z. *et al.* Deep manifold learning reveals hidden dynamics of proteasome autoregulation. *bioRxiv*, 2020.2012.2022.423932, doi:10.1101/2020.12.22.423932 (2021).
32 Onuchic, J. N., Luthey-Schulten, Z. & Wolynes, P. G. Theory of protein folding: the energy landscape perspective. *Annu Rev Phys Chem* **48**, 545-600, doi:10.1146/annurev.physchem.48.1.545 (1997).
33 Henriksen, N. E. & Hansen, F. Y. *Theories of Molecular Reaction Dynamics: The Microscopic Foundation of Chemical Kinetics*. 2nd edn, (Oxford University Press, Oxford, 2019).
34 Atkins, P., de Paula, J. & Keeler, J. *Physical Chemistry*. 11th edn, (Oxford University Press, Oxford, 2018).
35 Dashti, A. *et al.* Trajectories of the ribosome as a Brownian nanomachine. *Proc Natl Acad Sci U S A* **111**, 17492-17497, doi:10.1073/pnas.1419276111 (2014).
36 Haselbach, D. *et al.* Long-range allosteric regulation of the human 26S proteasome by 20S proteasome-targeting cancer drugs. *Nat Commun* **8**, 15578, doi:10.1038/ncomms15578 (2017).
37 Fischer, N., Konevega, A. L., Wintermeyer, W., Rodnina, M. V. & Stark, H. Ribosome dynamics and tRNA movement by time-resolved electron cryomicroscopy. *Nature* **466**, 329-333, doi:10.1038/nature09206 (2010).
38 Haselbach, D. *et al.* Structure and Conformational Dynamics of the Human Spliceosomal B(act) Complex. *Cell* **172**, 454-464 e411, doi:10.1016/j.cell.2018.01.010 (2018).
39 Dashti, A. *et al.* Retrieving functional pathways of biomolecules from single-particle snapshots. *Nat Commun* **11**, 4734, doi:10.1038/s41467-020-18403-x (2020).
40 Frank, J. & Ourmazd, A. Continuous changes in structure mapped by manifold embedding of single-particle data in cryo-EM. *Methods* **100**, 61-67, doi:10.1016/j.ymeth.2016.02.007 (2016).
41 Maji, S. *et al.* Propagation of Conformational Coordinates Across Angular Space in Mapping the Continuum of States from Cryo-EM Data by Manifold Embedding. *J Chem Inf Model* **60**, 2484-2491, doi:10.1021/acs.jcim.9b01115 (2020).
42 Coifman, R. R. *et al.* Geometric diffusions as a tool for harmonic analysis and structure definition of data: diffusion maps. *Proc Natl Acad Sci U S A* **102**, 7426-7431, doi:10.1073/pnas.0500334102 (2005).
43 He, K., Zhang, X., Ren, S. & Sun, J. Deep Residual Learning for Image Recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV*, 770-778, doi:10.1109/CVPR.2016.90 (2016).

44 Hinton, G. E. & Salakhutdinov, R. R. Reducing the dimensionality of data with neural networks. *Science* **313**, 504-507, doi:10.1126/science.1127647 (2006).
45 van der Maaten, L. & Hinton, G. Visualizing Data using t-SNE. *J Mach Learn Res* **9**, 2579-2605 (2008).
46 E, W., Ren, W. & Vanden-Eijnden, E. String method for the study of rare events. *Phys Rev B* **66**, 052301 (2002).
47 E, W., Ren, W. & Vanden-Eijnden, E. Simplified and improved string method for computing the minimum energy paths in barrier-crossing events. *J Chem Phys* **126**, 164103, doi:10.1063/1.2720838 (2007).
48 Sharif, H. *et al.* Structural mechanism for NEK7-licensed activation of NLRP3 inflammasome. *Nature* **570**, 338-343, doi:10.1038/s41586-019-1295-z (2019).
49 Cox, T. & Cox, M. *Multidimensional Scaling*. (Chapman & Hall, London, 1994).
50 Tenenbaum, J. B., de Silva, V. & Langford, J. C. A global geometric framework for nonlinear dimensionality reduction. *Science* **290**, 2319-2323, doi:10.1126/science.290.5500.2319 (2000).
51 Roweis, S. T. & Saul, L. K. Nonlinear dimensionality reduction by locally linear embedding. *Science* **290**, 2323-2326, doi:10.1126/science.290.5500.2323 (2000).
52 Wong, W. *et al.* Cryo-EM structure of the Plasmodium falciparum 80S ribosome bound to the anti-protozoan drug emetine. *Elife* **3**, doi:10.7554/eLife.03080 (2014).
53 Sun, M. *et al.* Dynamical features of the Plasmodium falciparum ribosome during translation. *Nucleic Acids Res* **43**, 10515-10524, doi:10.1093/nar/gkv991 (2015).
54 Plaschka, C., Lin, P. C. & Nagai, K. Structure of a pre-catalytic spliceosome. *Nature* **546**, 617-621, doi:10.1038/nature22799 (2017).
55 Yan, Z. *et al.* Structure of the rabbit ryanodine receptor RyR1 at near-atomic resolution. *Nature* **517**, 50-55, doi:10.1038/nature14063 (2015).
56 des Georges, A. *et al.* Structural Basis for Gating and Activation of RyR1. *Cell* **167**, 145-157 e117, doi:10.1016/j.cell.2016.08.075 (2016).
57 Wu, Z. *et al.* Chemistry of cation hydration and conduction in a skeletal muscle ryanodine receptor. *bioRxiv*, 732172, doi:10.1101/732172 (2019).
58 Davis, J. H. *et al.* Modular Assembly of the Bacterial Large Ribosomal Subunit. *Cell* **167**, 1610-1622 e1615, doi:10.1016/j.cell.2016.11.020 (2016).
59 Zhu, Y. *et al.* Structural mechanism for nucleotide-driven remodeling of the AAA-ATPase unfoldase in the activated human 26S proteasome. *Nat Commun* **9**, 1360, doi:10.1038/s41467-018-03785-w (2018).
60 Bard, J. A. M., Bashore, C., Dong, K. C. & Martin, A. The 26S Proteasome Utilizes a Kinetic Gateway to Prioritize Substrate Degradation. *Cell* **177**, 286-298 e215, doi:10.1016/j.cell.2019.02.031 (2019).
61 Asano, S. *et al.* Proteasomes. A molecular census of 26S proteasomes in intact neurons. *Science* **347**, 439-442, doi:10.1126/science.1261197 (2015).
62 Shi, Y. *et al.* Rpn1 provides adjacent receptor sites for substrate binding and deubiquitination by the proteasome. *Science* **351**, doi:10.1126/science.aad9421 (2016).
63 Saeki, Y., Sone, T., Toh-e, A. & Yokosawa, H. Identification of ubiquitin-like protein-binding subunits of the 26S proteasome. *Biochem Biophys Res Commun* **296**, 813-819, doi:10.1016/s0006-291x(02)02002-8 (2002).
64 Mao, Y. Structure, Dynamics and Function of the 26S Proteasome. In: Harris J.R., Marles-Wright J. (eds) Macromolecular Protein Complexes III: Structure and Function.

Subcell Biochem **96**, 1-151, doi:10.1007/978-3-030-58971-4_1 (2021).
65 Keren-Kaplan, T. *et al.* Structure of ubiquitylated-Rpn10 provides insight into its autoregulation mechanism. *Nat Commun* **7**, 12960, doi:10.1038/ncomms12960 (2016).
66 Uchiki, T. *et al.* The ubiquitin-interacting motif protein, S5a, is ubiquitinated by all types of ubiquitin ligases by a mechanism different from typical substrate recognition. *J Biol Chem* **284**, 12622-12632, doi:10.1074/jbc.M900556200 (2009).
67 Isasa, M. *et al.* Monoubiquitination of RPN10 regulates substrate recruitment to the proteasome. *Mol Cell* **38**, 733-745, doi:10.1016/j.molcel.2010.05.001 (2010).
68 Lipinszki, Z., Kovacs, L., Deak, P. & Udvardy, A. Ubiquitylation of Drosophila p54/Rpn10/S5a regulates its interaction with the UBA-UBL polyubiquitin receptors. *Biochemistry* **51**, 2461-2470, doi:10.1021/bi3001006 (2012).
69 Husnjak, K. *et al.* Proteasome subunit Rpn13 is a novel ubiquitin receptor. *Nature* **453**, 481-488, doi:10.1038/nature06926 (2008).
70 Schreiner, P. *et al.* Ubiquitin docking at the proteasome through a novel pleckstrin-homology domain interaction. *Nature* **453**, 548-552, doi:10.1038/nature06924 (2008).
71 Sahu, I. *et al.* The 20S as a stand-alone proteasome in cells can degrade the ubiquitin tag. *Nat Commun* **12**, 6173, doi:10.1038/s41467-021-26427-0 (2021).
72 Besche, H. C. *et al.* Autoubiquitination of the 26S proteasome on Rpn13 regulates breakdown of ubiquitin conjugates. *EMBO J* **33**, 1159-1176, doi:10.1002/embj.201386906 (2014).
73 de la Pena, A. H., Goodall, E. A., Gates, S. N., Lander, G. C. & Martin, A. Substrate-engaged 26S proteasome structures reveal mechanisms for ATP-hydrolysis-driven translocation. *Science* **362**, doi:10.1126/science.aav0725 (2018).
74 Zhang, S. & Mao, Y. AAA+ ATPases in Protein Degradation: Structures, Functions and Mechanisms. *Biomolecules* **10**, doi:10.3390/biom10040629 (2020).
75 Schrader, J. *et al.* The inhibition mechanism of human 20S proteasomes enables next-generation inhibitor design. *Science* **353**, 594-598, doi:10.1126/science.aaf8993 (2016).
76 Groll, M. *et al.* Structure of 20S proteasome from yeast at 2.4 A resolution. *Nature* **386**, 463-471, doi:10.1038/386463a0 (1997).
77 Perdew, J. P., Burke, K. & Ernzerhof, M. Generalized Gradient Approximation Made Simple. *Phys Rev Lett* **77**, 3865-3868, doi:10.1103/PhysRevLett.77.3865 (1996).
78 Saha, A., Oanca, G., Mondal, D. & Warshel, A. Exploring the Proteolysis Mechanism of the Proteasomes. *J Phys Chem B* **124**, 5626-5635, doi:10.1021/acs.jpcb.0c04435 (2020).
79 Huber, E. M. *et al.* A unified mechanism for proteolysis and autocatalytic activation in the 20S proteasome. *Nat Commun* **7**, 10900, doi:10.1038/ncomms10900 (2016).
80 Osmulski, P. A., Hochstrasser, M. & Gaczynska, M. A tetrahedral transition state at the active sites of the 20S proteasome is coupled to opening of the alpha-ring channel. *Structure* **17**, 1137-1147, doi:10.1016/j.str.2009.06.011 (2009).
81 Kleijnen, M. F. *et al.* Stability of the proteasome can be regulated allosterically through engagement of its proteolytic active sites. *Nat Struct Mol Biol* **14**, 1180-1188, doi:10.1038/nsmb1335 (2007).
82 Ruschak, A. M. & Kay, L. E. Proteasome allostery as a population shift between interchanging conformers. *Proc Natl Acad Sci U S A* **109**, E3454-3462, doi:10.1073/pnas.1213640109 (2012).
83 Chen, S. *et al.* Structural basis for dynamic regulation of the human 26S proteasome. *Proc Natl Acad Sci U S A* **113**, 12991-12996, doi:10.1073/pnas.1614614113 (2016).

84 Collins, G. A. & Goldberg, A. L. Proteins containing ubiquitin-like (Ubl) domains not only bind to 26S proteasomes but also induce their activation. *Proc Natl Acad Sci U S A* **117**, 4664-4674, doi:10.1073/pnas.1915534117 (2020).
85 Kim, H. T. & Goldberg, A. L. UBL domain of Usp14 and other proteins stimulates proteasome activities and protein degradation in cells. *Proc Natl Acad Sci U S A* **115**, E11642-E11650, doi:10.1073/pnas.1808731115 (2018).
86 Rennella, E., Huang, R., Yu, Z. & Kay, L. E. Exploring long-range cooperativity in the 20S proteasome core particle from Thermoplasma acidophilum using methyl-TROSY-based NMR. *Proc Natl Acad Sci U S A* **117**, 5298-5309, doi:10.1073/pnas.1920770117 (2020).
87 Yu, Z. *et al.* Allosteric coupling between alpha-rings of the 20S proteasome. *Nat Commun* **11**, 4580, doi:10.1038/s41467-020-18415-7 (2020).
88 Zhu, Y., Ouyang, Q. & Mao, Y. A deep convolutional neural network approach to single-particle recognition in cryo-electron microscopy. *BMC Bioinformatics* **18**, 348, doi:10.1186/s12859-017-1757-y (2017).
89 Bepler, T. *et al.* Positive-unlabeled convolutional neural networks for particle picking in cryo-electron micrographs. *Nat Methods* **16**, 1153-1160, doi:10.1038/s41592-019-0575-8 (2019).
90 Ioffe, S. & Szegedy, C. in *Proceedings of the 32nd International Conference on Machine Learning* Vol. 37 (eds Bach Francis & Blei David) 448--456 (PMLR, Proceedings of Machine Learning Research, 2015).
91 Qian, N. On the momentum term in gradient descent learning algorithms. *Neural Netw* **12**, 145-151, doi:10.1016/s0893-6080(98)00116-6 (1999).
92 Penczek, P. A., Yang, C., Frank, J. & Spahn, C. M. Estimation of variance in single-particle reconstruction using the bootstrap technique. *J Struct Biol* **154**, 168-183, doi:10.1016/j.jsb.2006.01.003 (2006).
93 Bonomi, M., Pellarin, R. & Vendruscolo, M. Simultaneous Determination of Protein Structure and Dynamics Using Cryo-Electron Microscopy. *Biophys J* **114**, 1604-1613, doi:10.1016/j.bpj.2018.02.028 (2018).
94 Matsumoto, S. *et al.* Extraction of protein dynamics information from cryo-EM maps using deep learning. *Nature Machine Intelligence* **3**, 153-160, doi:10.1038/s42256-020-00290-y (2021).
95 Unverdorben, P. *et al.* Deep classification of a large cryo-EM dataset defines the conformational landscape of the 26S proteasome. *Proc Natl Acad Sci U S A* **111**, 5544-5549, doi:10.1073/pnas.1403409111 (2014).
96 Wehmer, M. *et al.* Structural insights into the functional cycle of the ATPase module of the 26S proteasome. *Proc Natl Acad Sci U S A* **114**, 1305-1310, doi:10.1073/pnas.1621129114 (2017).
97 Emsley, P. & Cowtan, K. Coot: model-building tools for molecular graphics. *Acta Crystallogr D Biol Crystallogr* **60**, 2126-2132, doi:10.1107/S0907444904019158 (2004).
98 Pettersen, E. F. *et al.* UCSF Chimera--a visualization system for exploratory research and analysis. *J Comput Chem* **25**, 1605-1612, doi:10.1002/jcc.20084 (2004).
99 Adams, P. D. *et al.* PHENIX: a comprehensive Python-based system for macromolecular structure solution. *Acta Crystallogr D Biol Crystallogr* **66**, 213-221, doi:10.1107/S0907444909052925 (2010).
100 Afonine, P. V. *et al.* Real-space refinement in PHENIX for cryo-EM and crystallography.

Acta Crystallogr D Struct Biol **74**, 531-544, doi:10.1107/S2059798318006551 (2018).
101 The PyMOL Molecular Graphics System Version 2.0 (Schrödinger, LLC).
102 Goddard, T. D. *et al.* UCSF ChimeraX: Meeting modern challenges in visualization and analysis. *Protein Sci* **27**, 14-25, doi:10.1002/pro.3235 (2018).
103 Kucukelbir, A., Sigworth, F. J. & Tagare, H. D. Quantifying the local resolution of cryo-EM density maps. *Nat Methods* **11**, 63-65, doi:10.1038/nmeth.2727 (2014).
104 Heymann, J. B. & Belnap, D. M. Bsoft: image processing and molecular modeling for electron microscopy. *J Struct Biol* **157**, 3-18, doi:10.1016/j.jsb.2006.06.006 (2007).
105 Dolinsky, T. J., Nielsen, J. E., McCammon, J. A. & Baker, N. A. PDB2PQR: an automated pipeline for the setup of Poisson-Boltzmann electrostatics calculations. *Nucleic Acids Res* **32**, W665-667, doi:10.1093/nar/gkh381 (2004).
106 Soler, J. M. *et al.* The SIESTA method for ab initio order-N materials simulation. *J Phys Condens Matter* **14**, 2745-2779 (2002).

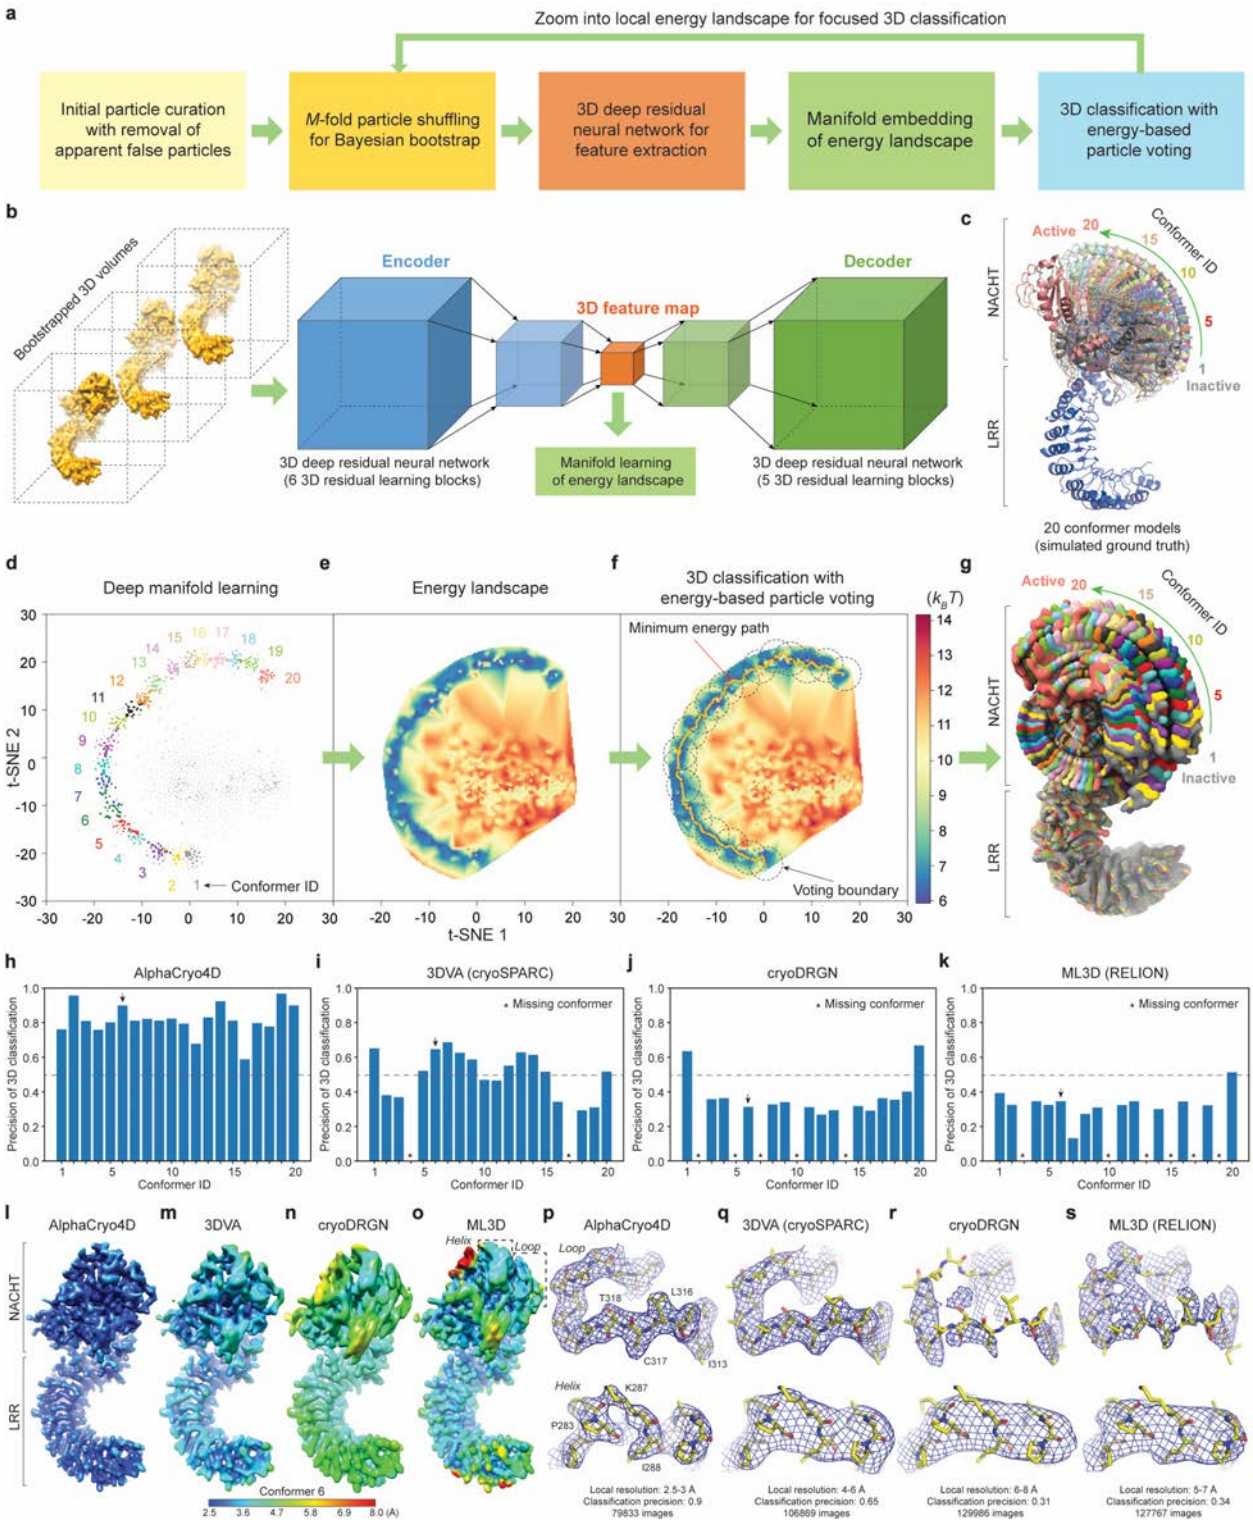

Figure 1. Conceptual framework of AlphaCryo4D for 4D cryo-EM reconstruction. a, Schematic showing the major conceptual steps of single-particle cryo-EM data processing in AlphaCryo4D. **b,** Illustration of deep residual learning of 3D feature maps by an autoencoder

conjugated to a decoder for unsupervised training. **c,** The 20 atomic models of hypothetical conformers of NLRP3 in cartoon representations simulate a 90° rotation of the NACHT domain relative to the LRR domain based on the structure of NLRP3 (PDB ID: 6NPY). **d,** Manifold learning of the bootstrapped 3D volumes and their feature maps learnt by the autoencoder. Each data point corresponds to a 3D volume. The color labels the conformer identity of ground truth for the purpose of verification. **e,** Energy landscape computed from the manifold shown in (**d**) using the Boltzmann distribution. **f,** Minimum energy path (orange line) calculated by the string method is used to find the approximate cluster centers of 20 conformers. The cluster boundaries for energy-based particle voting are shown as dashed circles. **g,** The 20 reconstructions resulting from the blind assessment of AlphaCryo4D using simulated dataset with SNR of 0.01 are superimposed together, shown in surface representation, and aligned against its LRR domain. **h-k,** Plots of 3D classification precisions of the 20 NLRP3 conformers from blind assessments on the simulated NLRP3 dataset with SNR of 0.01, using AlphaCyo4D (**h**), 3DVA algorithm in cryoSPARC[11] (**i**), cryoDRGN[12] (**j**), and conventional ML3D in RELION[3] (**k**). All 3D conformers in panel (**h**) were reconstructed to 2.6-2.9 Å resolution (Extended Data Fig. 3l). Asterisks mark the missing conformers that were completely lost due to misclassification by 3DVA, cryoDRGN and ML3D. **l-o**, Typical side-by-side comparison of density map quality and local resolution of the same conformer (ID 6) reconstructed by AlphaCryo4D (**l**), 3DVA (**m**), cryoDRGN (**n**) and ML3D (**o**). The maps are colored by their local resolutions calculated by Bsoft blocres program. **p-s,** Closeup side-by-side comparison of the same two secondary structures, including a loop (upper row) and a helix (lower row), in the NACHT domain illustrates considerable improvements in local density quality and resolution by AlphaCryo4D (**p**) as opposed to 3DVA/cryoSPARC (**q**), cryoDRGN (**r**) and ML3D/RELION (**s**). The locations of the loop and helix in the NLRP3 structure are marked by dashed boxes in panel (**o**). The same ground-truth atomic model of Conformer 6 shown in stick representations is superimposed with the density maps shown in blue mesh representations, all from the same perspective. The atomic model was not further refined against each map for visual validation of the map accuracy.

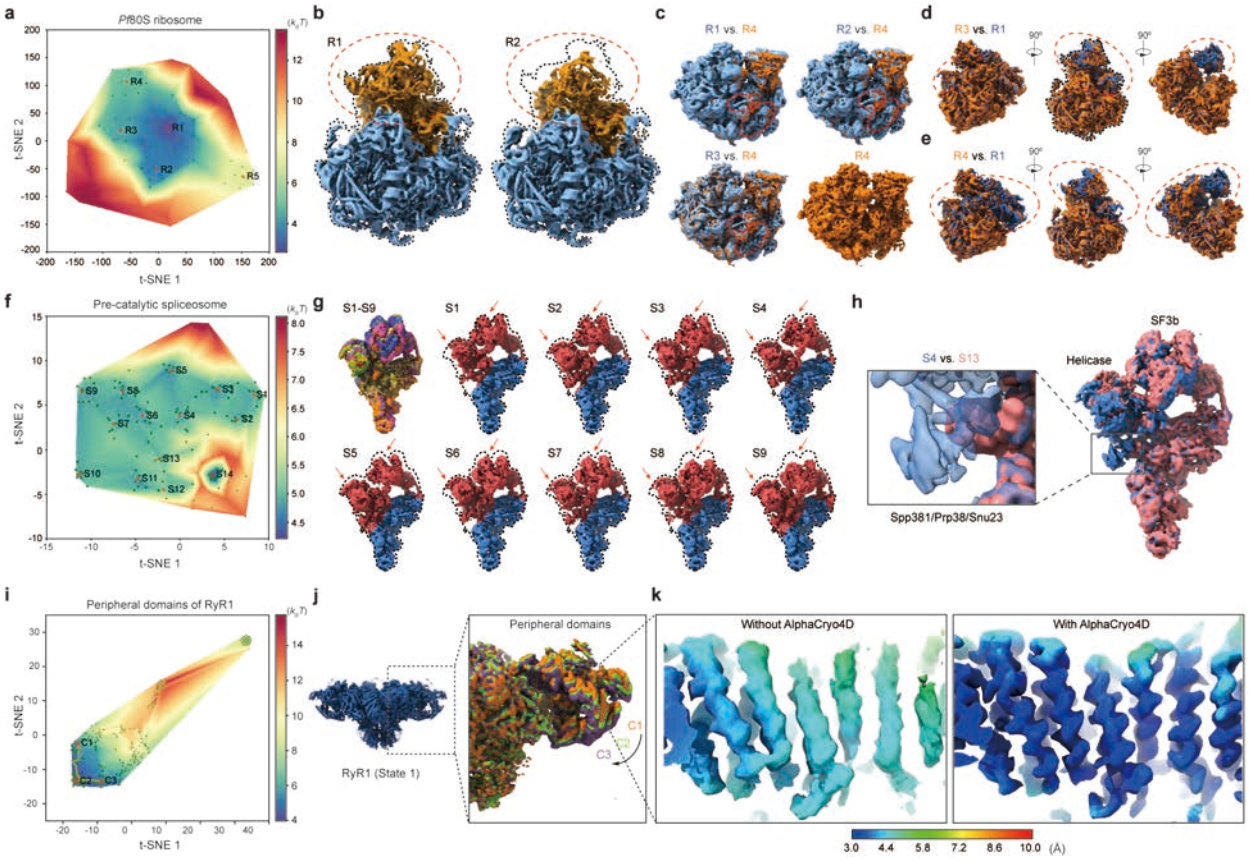

Figure 2. **Applications of AlphaCryo4D to three experimental datasets**. **a-e,** Visualizing conformational dynamics of the *Pf*80S ribosome bound to anti-protozoan drug at high resolution by AlphaCryo4D on the dataset EMPIAR-10028. **a,** Clusters R1-R5 on the energy landscape by AlphaCryo4D using 79 bootstrapped density maps. Energy-based particle voting was conducted within each cluster to keep particles of high reproducibility. **b,** Comparison of the density maps of clusters R1 and R2. Cluster R1 is the major state in the *Pf*80S ribosome dataset, while cluster R2 accounts for the state missing the head of 40S (red circles). The dynamic region of the density map is colored orange and the stable region is colored blue. **c,** Comparisons of states R1 and R4 (upper left), states R2 and R4 (upper right), and states R3 and R4 (lower left) reveal the differential appearance of the C-terminal helix of eL8 and an rRNA helix. Note that the rRNA helix is only partially presented in state R3. **d,** State R3 superimposed with state R1 from three different viewing angles. The density of the 40S head region (red circles) rotated relative to the 60S. **e,** State R4 superimposed with state R1 from three different viewing angles. The density of the whole 40S (red circles) is rotated relative to the 60S. **f-h**, Choreographing continuous inter-subunit motions of the pre-catalytic spliceosome by AlphaCryo4D on the dataset EMPIAR-

10180. **f,** Energy landscape by AlphaCryo4D, which bootstrapped 160 density maps using 327,490 particles, reveals 14 clusters designated as S1-S14. Clusters S1-S9 are related to the continuous rotational motion of the SF3b and helicase subcomplexes of pre-catalytic spliceosome, and the density maps of clusters S10-S14 are shown in Extended Data Fig. 7i. **g,** Cryo-EM density maps of clusters S1-S9. The dynamic component of the density map is colored red and the static component is colored blue. The red arrows suggest the rotational motion of the SF3b and helicase subcomplexes within the dashed outline. **h,** Density map of the Spp381/Prp38/Snu23 missing state S13 superimposed with the state S4. Closeup comparison shows the density difference between states S13 and S4. **i-k**, Resolution improvement of the dynamic components of RyR1 and the 26S proteasome by AlphaCryo4D. **i,** Energy landscape of the peripheral domains of the RyR1 using 310 bootstrapped density maps. Clusters C1, C2 and C3 were selected for particle voting by this energy landscape. **j,** Cryo-EM density maps of clusters C1-C3 classified by AlphaCryo4D. **k,** Comparison of the helical domain density map of the peripheral domains without and with AlphaCryo4D.

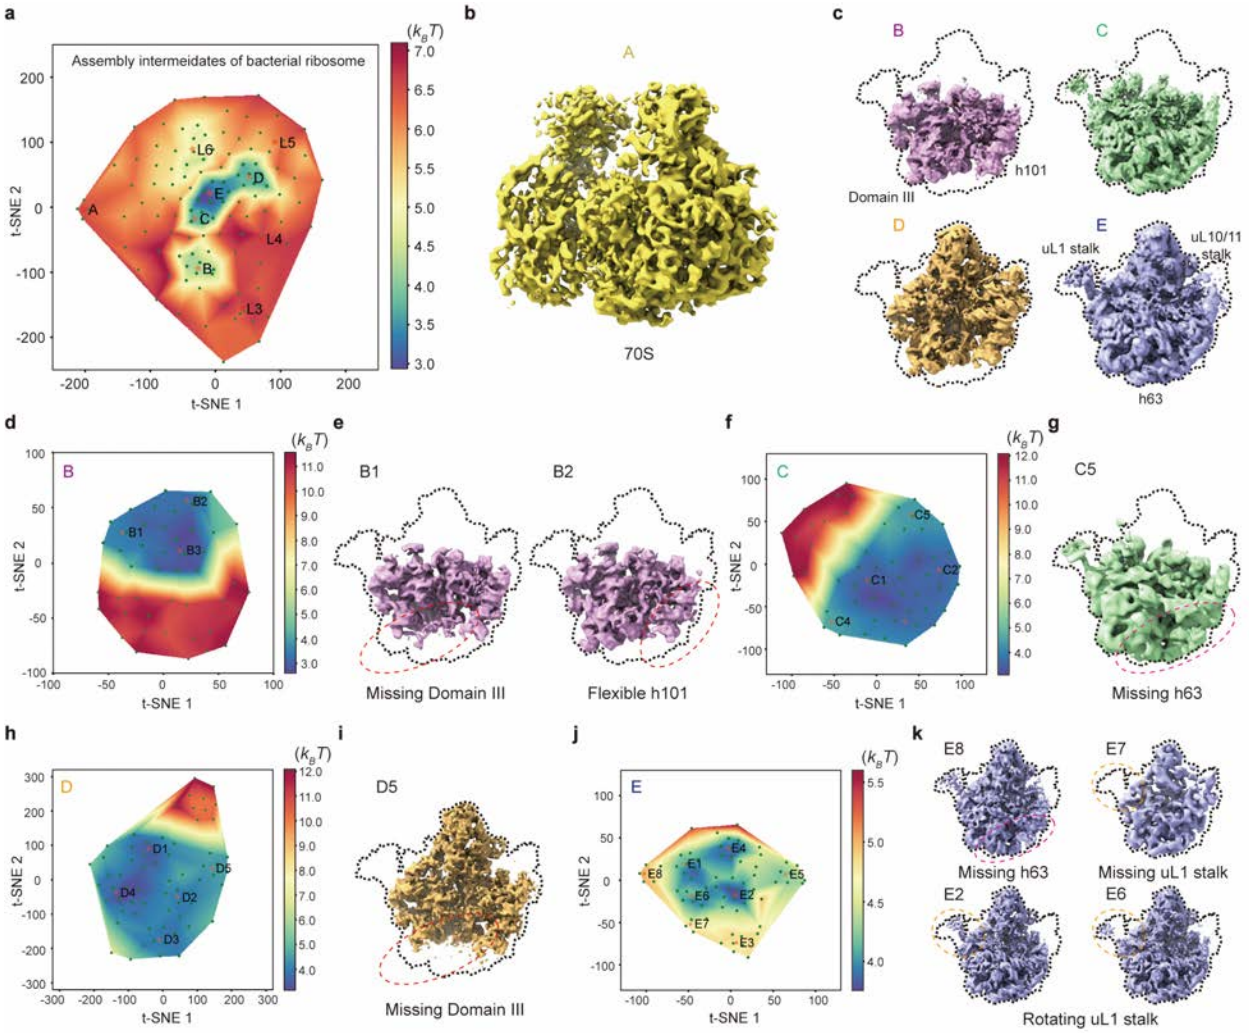

Figure 3. Exploring hidden conformational space of the bacterial ribosomal assembly intermediates by AlphaCryo4D. **a,** Global energy landscape of bacterial ribosomal assembly intermediates determined by AlphaCryo4D on the dataset EMPIAR-10076. Clusters A-E and L3-L6 were reconstructed for further analysis. **b,** Cryo-EM density map of state A reconstructed by only 766 particles still retains enough details of the complete 70S assembly. This state is located far from other states on the energy landscape shown in (**a**). **c,** Density maps of the major states B, C, D and E reconstructed by AlphaCryo4D. The density maps of these four states were refined using the particles of clusters B, C, D and E respectively from the energy landscape. **d,** Zoomed-in energy landscape of the major state B. **e,** Density maps of new conformers B1 and B2 that are subclasses of state B. Comparing to the original state B (Extended Data Fig. 8b), the new states B1 and B2 appear to miss Domain III and show flexible rRNA helix 101, respectively. **f,** Zoomed-in energy landscape of the major state C generating clusters C1-C5. **g,** A new state (C5)

of C determined by AlphaCryo4D. The rRNA helix 63 was not observed in the density map of C5. **h,** Clusters D1-D5 from the zoom-in energy landscape of the major state D. **i,** A newly discovered minor state D5. Like the minor state B1 of the major state B, Domain III was not seen in the density map of D5. **j,** Clusters E1-E8 on the zoom-in energy landscape of the major state E by AlphaCryo4D. **k,** Discovery of new conformers in the major state E. Like C5, the new state E8 lacks the rRNA helix 63. Comparing to the previously reported state E3 (Extended Data Fig. 8e), the uL1 stalk is not found in the subclass E7. Two new states, E2 and E6, with different conformational changes of the uL1 stalk were discovered by AlphaCryo4D.

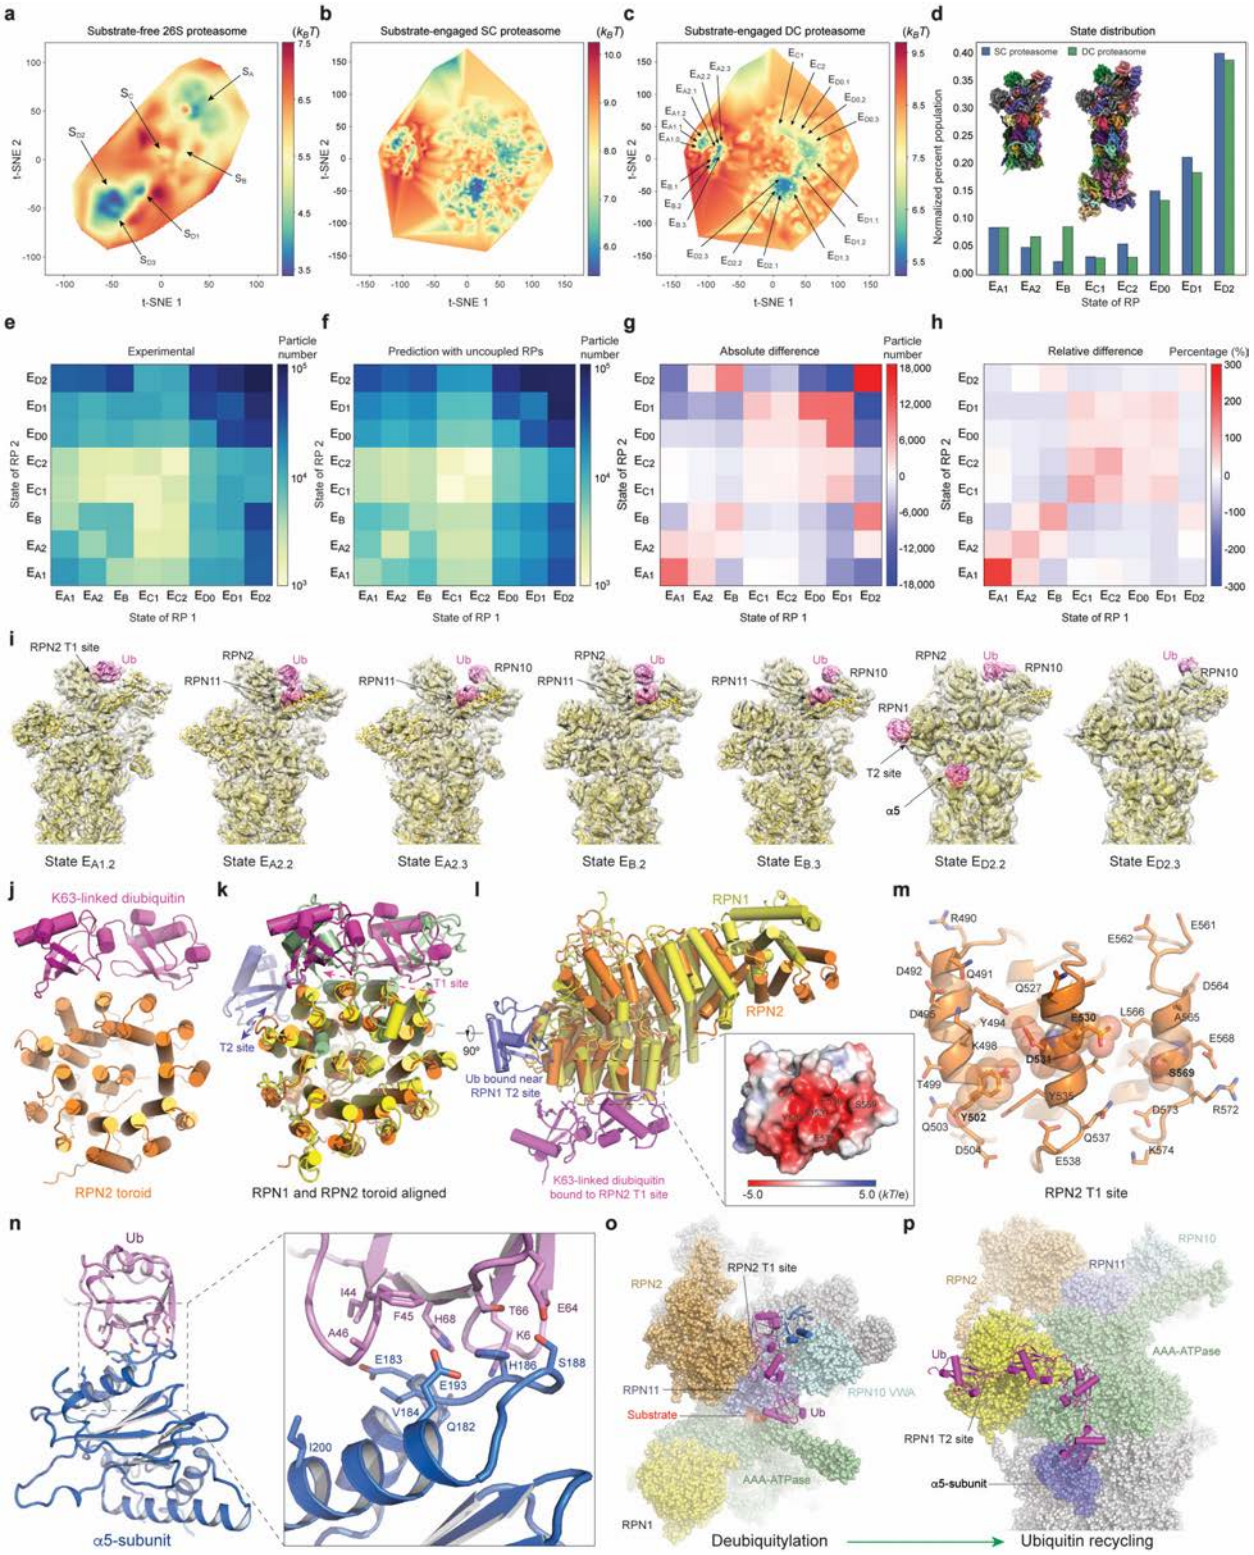

Figure 4. **Discovery of hidden dynamics of human proteasome autoregulation by AlpahCryo4D**. **a-c**, Energy landscape of the substrate-free ATPγS-bound 26S proteasome from the dataset EMPIAR-10090 (**a**), the substrate-engaged SC (**b**) and DC (**c**) proteasome from the

dataset EMPIAR-10669, computed by AlphaCryo4D. 20 states of the RP-CP subcomplex are marked on the energy landscape. **d**, Comparison of the state distribution of RP in the SC and DC proteasomes. **e**, State distribution matrix of the DC proteasome obtained using AlphaCryo4D on experimental data, colored by the particle numbers in the DC states, with the horizonal and vertical axes representing the states of two RPs bound to the same CP. **f**, State distribution matrix of the DC proteasome predicted by a control model assuming the two completely uncoupled RPs in the same DC proteasome. The model assumes the same total number of DC particles and the same probability of observing each RP state as experimentally measured. **g**, State distribution difference matrix of the DC proteasome by experimental results minus the model predictions, colored by the absolute particle number difference. **h**, Relative differences of the DC proteasome computed as the ratio of absolute difference divided by the particle numbers from the model prediction in each DC state. **i**, Seven lowly populated states visualize various ubiquitin-binding sites on the RPN1, RPN2, RPN10 and α5 subunits. For clarity, the cryo-EM maps are low-pass filtered at 8 Å and shown in transparent surface representations superimposed with their corresponding atomic models in cartoon representations. The ubiquitin (Ub) densities are colored magenta. **j**, Structure of the RPN2 toroidal domain (orange) in complex with Lys63-linked diubiquitin (magenta) modelled from the cryo-EM maps of states $E_{A1.2}$, $E_{A2.2}$ and $E_{B.2}$, in cartoon representation. **k**, Structure-based alignment of the human RPN1 (yellow) and RPN2 (orange) toroidal domains. Ubiquitin bound near the RPN1 T2 site is shown as light blue. The NMR structure (PDB ID 2N3V) of the Lys48-linked dibuiquitin-bound yeast Rpn1 T1 site segment[62], shown in light green, is aligned and superimposed with the human RPN1 and RPN2. **l**, The aligned human RPN1 and RPN2 are viewed from a perspective orthogonal to that in panel (**k**). Insert, electrostatic surface of the RPN2 T1 site showing its acidic nature. **m**, Closeup view of the RPN2 T1 site composed of three helices in carton representation with the side chains shown in stick representation. The key residues contacting ubiquitin are highlighted with transparent sphere representation. **n**, Structure of ubiquitin-bound α5 subunit shown in cartoon representation. The side chains of residues mediating intermolecular interactions are shown in stick representation. **o**, A Lys63-linked triubiquitin chain modelled from the cryo-EM maps of states $E_{A2.2}$ and $E_{B.2}$ is shown in purple cartoon representation and spans from the RPN2 T1 site to the RPN11 deubiquitylation site. The RPN10 VWA-bound monoubiquitin fitted in states $E_{A2.3}$ and $E_{B.3}$ is superimposed in blue cartoon representation. The proteasome is shown as space-

filling sphere representations. **p**, A Lys63-linked tetraubiquitin chain model derived from the cryo-EM map of state $E_{D2.2}$. The two terminal ubiquitin molecules are fitted from their highest-resolution density maps available at the RPN1 T2 and α5-subunit sites, whereas the middle two molecules are hypothetically modelled.

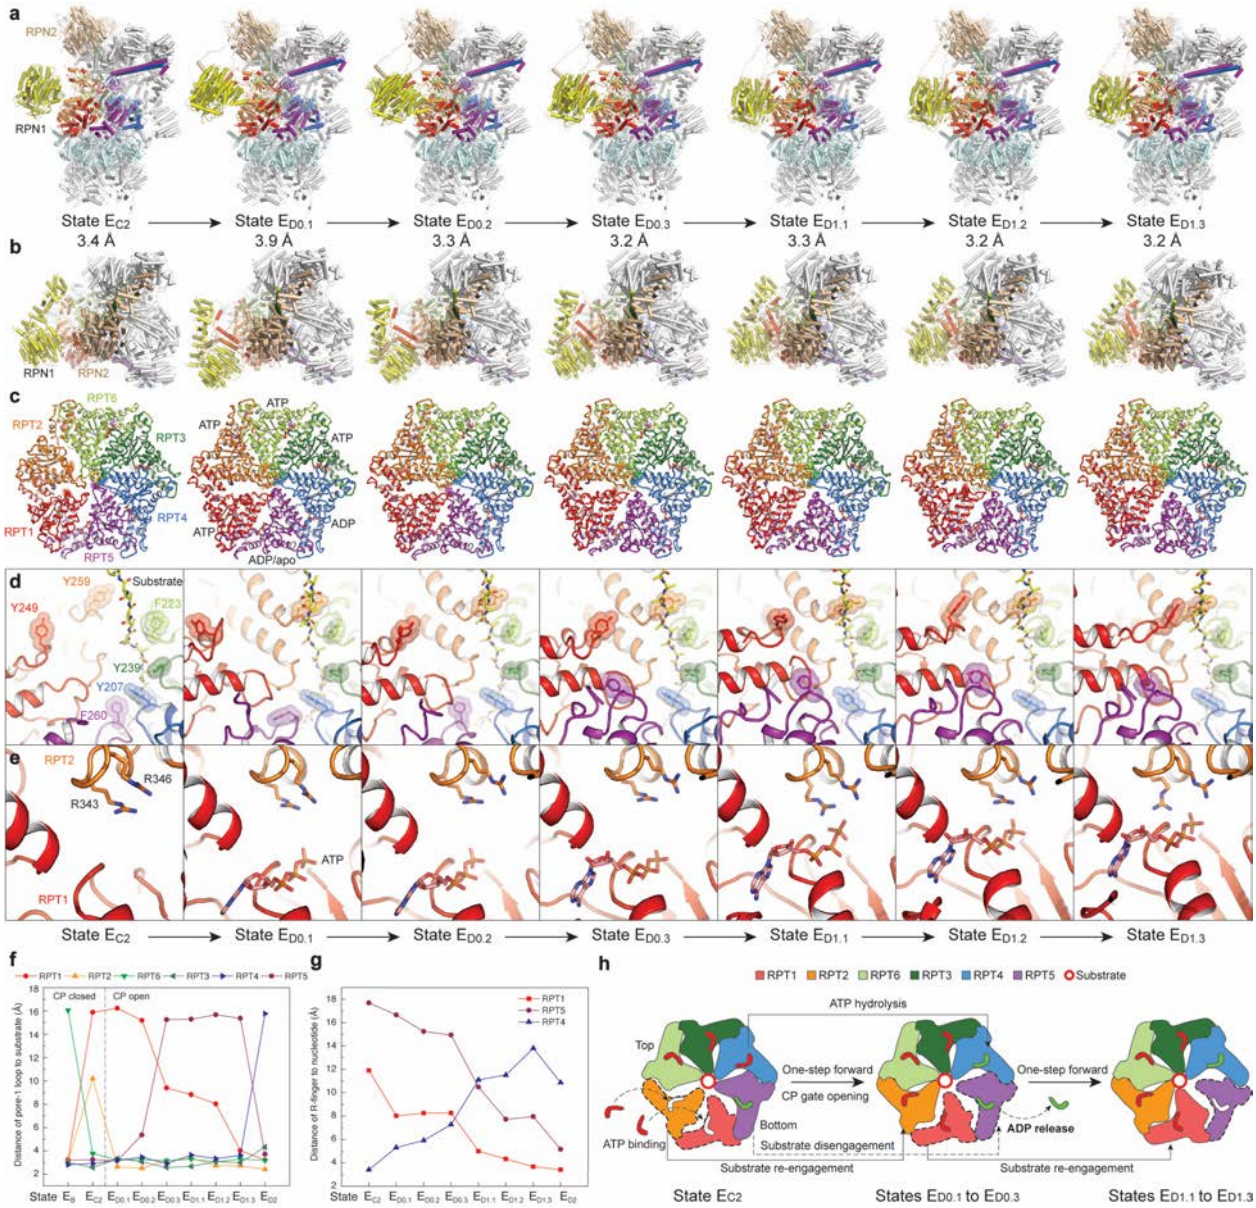

Figure 5. Single-nucleotide-exchange dynamics of the proteasomal AAA-ATPase motor during translocation initiation. **a**, Side views of the substate-engaged RP-CP subcomplex of seven sequential conformers in cartoon representation. **b**, Top views of the seven sequential conformers, showing the overall lid rotation and the relative rotation between RPN1 and the rest of the RP structure. **c**, The AAA-ATPase structures of the seven conformers in cartoon representations. The substrates and nucleotides are shown in stick representations. **d**, Close-up views of the pore-1 loop interactions with the substrate (in yellow stick representation), with the aromatic residues in the pore-1 loops highlighted by transparent sphere representations. **e**, Close-up views of the RPT1 nucleotide-binding sites in different states. **f**, Plots of distances of the pore-

1 loops of all RPT subunits to the substrate measured in the atomic structures of 9 sequential conformational states. **g**, Plots of distances of R-finger to the nucleotides bound to three RPT subunits measured in the atomic structures of 8 sequential conformational states. **h**, Schematic illustrating the ATP hydrolysis reaction and nucleotide exchange associated with the intermediates during the transition from state E_{C2} to $E_{D1.3}$.

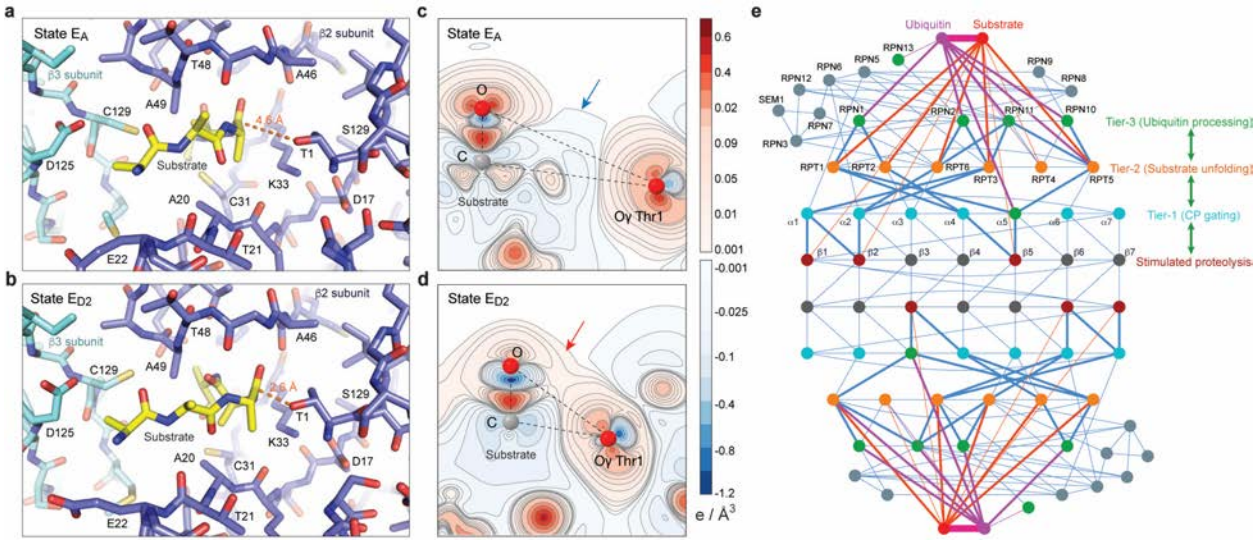

Figure 6. **Hierarchical allosteric regulation of proteasome activity**. **a** and **b**, Structure of the catalytic site with a substrate polypeptide bound at residue Thr1 in the β2 subunit of the CP in state E_A at 2.7-Å resolution (**a**) and in state E_{D2} at 2.5-Å resolution (**b**). **c** and **d**, Charge density difference contour maps at the proteolytic active site of the β2 subunit computed by the quantum mechanical DFT based on the atomic structures of states E_A (**c**) and E_{D2} (**d**). The contours are plotted in the plane passing through the Thr1-Oγ atom and the nearest carbonyl carbon and oxygen atoms of the substrate. Blue and red arrows highlight a non-bonding charge depletion region in state E_A and a bonding interaction indicated by the slight positive charge difference region in state E_{D2}, respectively. **e**, A network model of hierarchical allosteric regulation of proteasome activity. The interactome network diagram illustrating the grand allostery of the proteasome. Steel blue lines represent inter-subunit interactions; purple lines represent the interactions with ubiquitin; and red lines represent substrate polypeptide interactions. Bold lines highlight the shortest pathways connecting the three tiers of regulatory subunits that are expected to propagate the allosteric regulation from ubiquitin interactions down to the six proteolytic active sites in the CP. The network diagram is constructed according to all 64 structural models of the substrate-bound proteasome in this study. Notably, the subnetwork architecture of RPN1 connecting to the proteolytic subunits is identical to that of RPN2 connecting to the proteolytic subunits, consistent with the structural homology of RPN1 and RPN2 as well as their potential functional similarity in ubiquitin interactions.

Methods

***M*-fold particle shuffling and Bayesian bootstrapping**. A key philosophy of the AlphaCryo4D design is to avoid subjective judgement on the particle quality and usability if it is not apparent false particles like ice contaminants, carbon edges and other obvious impurities. Deep-learning-based particle picking in DeepEM[88], Torpaz[89] or other similarly performed software is favored for data preprocessing prior to AlphaCryo4D. To prepare particle datasets for AlphaCryo4D, an initial unsupervised 2D image classification and particle selection, preferentially conducted by the statistical manifold-learning-based algorithm in ROME[5], is necessary to ensure that no apparent false particles are selected for further analysis and that the data have been collected under an optimal microscope alignment condition, such as optimized coma-free alignment. No particles should be discarded based on their structural appearance during this step if they are not apparent false particles. Any additional 3D classification should be avoided to pre-maturely reject particles prior to particle shuffling and volume bootstrapping in the first step of AlphaCryo4D processing. Pre-maturely rejecting true particles via any 2D and 3D classification is expected to introduce subjective bias and to impair the native conformational continuity and statistical integrity intrinsically existing in the dataset.

In raw cryo-EM data, 2D transmission images of biological macromolecules suffer from extremely heavy background noise, due to the use of low electron dose to avoid radiation damage. To tackle the conformational heterogeneity of the macromolecule sample of interest in the presence of heavy image noise, the particle shuffling and volume bootstrapping procedure was devised to incorporate the Bayesian or maximum-likelihood-based 3D clustering in RELION[3]. To enable the particle-voting algorithm in the late stage of AlphaCryo4D, each particle is reused *M* times during particle shuffling to bootstrap a large number of 3D volumes (Extended Data Fig. 1a). First, all particle images are aligned to the same frame of reference in a consensus 3D reconstruction and refinement in RELION[3] or ROME[5] to obtain the initial alignment parameters of three Euler angles and two translational shifts. Optimization for alignment accuracy should be pursued to avoid the error propagation to the subsequent steps in AlphaCryo4D. For a dataset with both compositional and conformational heterogeneity, coarsely classifying the dataset to a few 3D classes during initial image alignment, limiting the 3D alignment to a moderate resolution like 10 Å or 15 Å during global orientational search, progressing to small enough angular steps in the final stage of consensus refinement, may be

optionally practiced to optimize the initial 3D alignment. Failure of initial alignment of particles would lead to failure of all subsequent AlphaCryo4D analysis.

Next, based on the results of consensus alignment, in the particle-shuffling step, all particles were divided into $M + 1$ groups, and M was set to an odd number of at least 3. Then the whole dataset was shuffled $M + 1$ times by removing a different group out of the dataset each time. Each shuffled dataset is classified into a large number (B) of 3D volumes, often tens to hundreds, by the maximum-likelihood 3D classification algorithm without further image alignment in RELION (with 'skip-align' option turned on). This is necessary because the alignment accuracy often degrades when the sizes of 3D classes decrease considerably. This step is repeated $M + 1$ times, each time on a shuffled dataset missing a different group among the $M + 1$ groups. Due to the effect of M-fold particle shuffling, the outcome of this entire process is expected to bootstrap up to thousands of 3D volumes in total (i.e., $B(M+1) > 1000$). Each particle is used and contributed to the 3D reconstructions of M volumes, which prepare it for the particle-voting algorithm to evaluate the robustness and reproducibility of each particle with respect to the eventual 3D classification.

For processing a large dataset including millions of single-particle images, it becomes infeasible even for a modern high-performance computing system to do the consensus alignment by including all particles once in a single run due to limitation of supercomputer memory and the scalability of the alignment software. To tackle this issue, the whole dataset is randomly split into a number (D) of sub-datasets for batch processing, with each sub-dataset including about one to two hundred thousand particles, depending on the scale of available supercomputing system. In this case, the initial reference should be used for the consensus alignment of different sub-datasets to minimize volume alignment errors in the later step. The total number of resulting bootstrapped volumes becomes $BD(M+1) > 1000$. This strategy can substantially reduce the supercomputer memory pressure and requirement. In each sub-dataset, all particles were divided into $M + 1$ groups and subject to the particle shuffling and volume bootstrapping procedure as described above.

Deep residual autoencoder for 3D feature extraction. The bootstrapped volumes may still suffer from reconstruction noises and errors due to variation of particle number, misclassification of conformers and limited alignment accuracy. Thus, we hypothesize that unsupervised deep

learning could help capture the key features of structural variations in the bootstrapped volumes and potentially enhance the quality of subsequently reconstituted energy landscapes for improved 3D classification. To this end, a 3D autoencoder was constructed using a deep Fully Convolutional Network (FCN) composed of residual learning blocks[43,44]. The architecture of the 3D autoencoder consists of the encoder and the decoder, which are denoted as $\mathcal{E}$ and $\mathcal{D}$, respectively (Fig. 1b). The relation between the output $\boldsymbol{y}$ and the input $\boldsymbol{x}$ of the network can be expressed as $\boldsymbol{y} = \mathcal{D}(\mathcal{E}(\boldsymbol{x}))$, in which $\boldsymbol{x}$ is the input 3D density volume with the size of N^3, where N is the box size of the density map in pixel units. For reconstruction of the 3D volumes and further optimization, the decoding output $\boldsymbol{y}$ should be in the same size and range with the input data $\boldsymbol{x}$. In this way, the framework of FCN is established to restore the input volume, using the sigmoid function $S(x) = \frac{1}{1+\exp(-x)}$ as the activation function of the decoding layer to normalize the value of $\boldsymbol{y}$ into the range (0, 1). Meanwhile, all 3D density maps $\boldsymbol{x}$ are normalized with the following equation before input to the deep neural network:

$$x_{ijk} := \frac{x_{ijk} - x_{min}}{x_{max} - x_{min}}, \quad i, j, k = 1,2, \dots, N. \qquad (1)$$

where x_{min} and x_{max} are, respectively, the minimum and minimax value in all x_{ijk}.

The distance between the decoded maps and the input volumes can be used for constructing the loss function to train the 3D kernels and bias of the networks. The value distribution of the encoded 3D feature maps $\boldsymbol{z} = \mathcal{E}(\boldsymbol{x})$ is expected to be an abstract, numerical representation of the underlying structures in the volume data, which may not necessarily have any intuitive real-space physical meanings. The neural network can extract such abstract information in the prediction step, with no restriction on the feature maps $\boldsymbol{z}$ in the expression of training loss. The loss function is then formulated as:

$$L(\boldsymbol{\theta}; \boldsymbol{x}, \boldsymbol{y}) = \frac{1}{N^3} \sum_{i,j,k=1}^{N} \left\| x_{ijk} - y_{ijk} \right\|^2 + \lambda \|\boldsymbol{\theta}\|^2 \qquad (2)$$

where $\boldsymbol{\theta}$ denotes the weights and bias of the network, and λ is L2 norm regularization coefficient. As the feature of the complex structure is difficult to be learned from the 3D volume data, the value of λ is set to 0.001 in default to prevent overfitting in the applications to experimental datasets. But it is set to 0 in the tests on the simulated datasets (see below).

To improve the learning capacity of the 3D autoencoder, residual learning blocks containing 3D convolutional and transposed convolutional layers are employed in the encoder and decoder, respectively. In each residual learning block, a convolutional layer followed by a Batch Normalization[90] (BN) layer and activation layer appears twice as a basic mapping, which is added with the input to generate the output. The mathematical expression of the lth block can be shown as:

$$\begin{cases} \boldsymbol{y_l} = \mathcal{F}(\boldsymbol{x_l}; \boldsymbol{\theta}_l) + \boldsymbol{x_l} \\ \mathcal{F}(\boldsymbol{x_l}; \boldsymbol{\theta}_l) = \mathcal{C}\big(\mathcal{C}(\boldsymbol{x_l})\big) \\ \mathcal{C}(\boldsymbol{x_l}) = a(b(c(\boldsymbol{x_l}))) \end{cases} \tag{3}$$

where $\boldsymbol{x_l}$ and $\boldsymbol{y_l}$ represent the input and the output of the block, respectively. $\mathcal{F}(x_l; \boldsymbol{\theta}_l)$ denotes the basic mapping of the lth block parameterized by $\boldsymbol{\theta}_l$, and $\mathcal{C}(\boldsymbol{x_l})$ is the sequential operation of convolution or transposed convolution c, BN b and activation a. The rectified linear unit (ReLU) function is used in all the activation layers but the last one. In addition, the output of mapping function $\mathcal{F}(\boldsymbol{x_l}; \boldsymbol{\theta}_l)$ must have the same dimension as the input $\boldsymbol{x_l}$. If this is not the case, the input $\boldsymbol{x_l}$ must be rescaled along with $\mathcal{F}(\boldsymbol{x_l}; \boldsymbol{\theta}_l)$ using a convolutional or transposed convolutional transformation $c'(\boldsymbol{x_l})$ with an appropriate stride value, the parameters of which can be updated in the training step.

Autoencoder training. To analyze a large number of volumes, we carefully tuned the training of the 3D deep residual autoencoder to obtain suitable kernels and bias. First, parallel computation with multiple GPUs has been implemented to reduce the training time. Then the parameters of the network are optimized by the stochastic gradient descent Adam (Adaptive moment estimation) algorithm, in which the gradients of the objective function $L(\boldsymbol{\theta}; \boldsymbol{x}, \boldsymbol{y})$ with respect to the parameters $\boldsymbol{\theta}$ can be calculated by the chain rule. Moreover, the learning rate is reduced to one tenth when the loss function does not decrease in three epochs based on the initial value of 0.01. After trained about 50 epochs, the best model is picked to execute the task of structural feature extraction. Using the unsupervised 3D autoencoder, the feature maps $\boldsymbol{z} = \mathcal{E}(\boldsymbol{x})$ encoding the structural discrepancy among the 3D volume data can be extracted automatically without any human intervention.

Autoencoder hyperparameters. The recommended hyperparameters for the deep residual

autoencoder architecture are provided in Supplementary Table 1. Although we expand the residual network (ResNet) into 3D, we keep the original design rules of ResNet for constructing the 3D autoencoder[43,44]. The first and last convolutional layers have 5 × 5 × 5 filters (kernels). The remaining convolutional layers have 3 × 3 × 3 filters. Because the cubic filter in convolutional layer is computationally expensive, only one 3D filter is used in each of the last three convolutional layers in the encoder and of the first two and last transposed convolutional layers in the decoder. To accommodate a large volume input, two filters are used in the first three convolutional layers and in the third and fourth transposed convolutional layers. Downsampling is directly performed by convolutional layers with a stride of 2. The output dimension of encoding layers is set to be a quarter of the input dimension to avoid over-compressing the feature maps. We employed six 3D convolutional layers in the encoder and five transposed convolutional layers in the decoder based on the expected tradeoff between the learning accuracy and training cost, which is roughly comparable to a 2D ResNet with more than 200 layers with respect to the training cost. Lastly, one should randomly inspect some feature maps to ensure that all hyperparameter setting works as expected. If a feature map clearly shows no structural correspondence to its input volume and exhibits strong noise or abnormal features, it might signal the existence of overfitting. In this case, a non-zero value of λ in the L2 norm regularization should be empirically examined to avoid overfitting.

Manifold embedding of energy landscape. To prepare for the energy landscape reconstitution, each bootstrapped 3D volume was juxtaposed with its 3D feature map learned by the 3D autoencoder to form an expanded higher-dimensional data point. Each expanded data point is either low-pass filtered at a given resolution (5 Å in default setting) or standardized by the z-score method prior to t-SNE processing:

$$\begin{cases} x_{ijk} := \dfrac{x_{ijk} - \mu_x}{\sigma_x}, \quad i,j,k = 1,2,\dots,N. \\ z_{ijk} := \dfrac{z_{ijk} - \mu_z}{\sigma_z}, \quad i,j,k = 1,2,\dots,N'. \end{cases} \tag{4}$$

in which μ_x and σ_x denote the mean and standard deviation of the 3D volume voxels x_{ijk} with the box size of N, respectively. Similarly, μ_z and σ_z are the mean and standard deviation of the 3D feature map voxels z_{ijk} with the box size of N', respectively. All the input data points were then embedded onto a low-dimensional manifold via the t-SNE algorithm by preserving the

geodesic relationships among all high-dimensional data[45]. During manifold embedding, it is assumed that the pairwise similarities in the high dimensional data space and low dimensional latent space follow a Gaussian distribution and a student's t-distribution, respectively. To find the low-dimensional latent variable of each data point, the relative entropy, also called the Kullback-Leibler (KL) divergence, is used as an objective function measuring the distance between the similarity distribution p_{ij} in the data space and q_{ij} in the latent space:

$$KL(P \parallel Q) = \sum_{i,j} p_{ij} \log \frac{p_{ij}}{q_{ij}} \qquad (5)$$

which is minimized by the gradient descent algorithm with the momentum method[91]. The hyperparameter of t-SNE may be tuned to better preserve both local and global geodesic distances (https://distill.pub/2016/misread-tsne/).

The idea of using concatenated data composed of both volumes and their feature maps for manifold embedding is similar to the design philosophy of deep residual learning, in which the input and feature output of a residual learning block are added together to be used as input of the next residual learning block. Such a design has been demonstrated to improve learning accuracy and reduce the performance degradation issues when the neural network goes much deeper[43,44]. Although the 3D feature maps or the bootstrapped volumes alone can be used for manifold embedding, both appear to be inferior in 3D classification accuracy (Extended Data Fig. 6). The juxtaposed format of input data for manifold learning is potentially beneficial to the applications in those challenging scenarios, such as visualizing a reversibly bound small protein like ubiquitin (~8.6 kDa) of low occupancy by the focused AlphaCryo4D classification (see below).

After the manifold embedding by t-SNR, each 3D volume is mapped to a low-dimensional data point in the learned manifold. The coordinate system, in which the low-dimensional representation of the manifold is embedded, is used for reconstructing energy landscape. The difference in the Gibbs free energy ΔG between two states with classified particle numbers of N_i and N_j is defined by the Boltzmann relation $N_i/N_j = \exp(-\Delta G/k_B T)$. Thus, the free energy of each volume can be estimated using its corresponding particle number:

$$\Delta\Delta G_i = -k_B T \ln \frac{N_i}{\sum_k N_k}, \qquad (6)$$

where $\Delta\Delta G_i$ denotes the free energy difference of the data point with the classified particle number of N_i against a common reference energy level, k_B is the Boltzmann constant and T is

the temperature in Kelvin. The energy landscape was plotted by interpolation of the free energy difference in areas with sparse data. We suggest that linear interpolation is used for the energy landscape with loosely sampled areas to avoid overfitting by polynomial or quadratic interpolation. For densely sampled energy landscape, polynomial interpolation could give rise to smoother energy landscape that is easier to be tackled by the string method for MEP solution (see below).

The coordinate system of the embedded manifold output by t-SNE is inherited by the reconstitution of energy landscape as reaction coordinates. They do not have intuitive physical meaning of length scale in the real space. However, they can be viewed as transformed, rescaled, reprojected coordinates from the real-space reaction coordinates along which the most prominent structural changes can be observed. Alternatively, they can be intuitively understood as being similar to transformed, reprojected principal components (PCs) in principle component analysis (PCA). In the case of the 26S proteasome, the two most prominent motions are the rotation of the lid relative to the base and the intrinsic motion of AAA-ATPase motor[30], which can be approximately projected to two reaction coordinates. In real space, both motions are notably complex and in fact are characterized by a considerable number of degrees of freedom, which are partially defined by the solved conformers.

String method for finding minimum energy path (MEP). The string method is an effective algorithm to find the MEP on the potential energy surface[46]. To extract the dynamic information implicated in the experimental energy landscape, an improved and simplified version of the string method has been previously developed[47]. Along the MEP on the energy landscape, the local minima of interest could be defined as 3D clustering centers to guide the particle-voting algorithm for 3D classification to generate high-resolution cryo-EM density maps (Extended Data Fig. 1b). The objective of the MEP identification in energy barrier-crossing events lies in finding a curve γ having the same tangent direction as the gradient of energy surface ∇G. It can be expressed as:

$$(\nabla G)^{\perp}(\gamma) = 0, \qquad (7)$$

where $(\nabla G)^{\perp}$ denotes the component of ∇G perpendicular to the path γ. To optimize this objective function, two computational steps, referred to as 'evolution of the transition path' and 'reparameterization of the string', are iterated until convergence within a given precision

threshold.

After initialization with the starting and ending points, in the step of evolution of the transition path, the positions of interval points along the transition path were updated according to gradient of the free energy at the tth iteration:

$$\varphi_i^{*(t)} = \varphi_i^{(t-1)} - h\nabla G\left(\varphi_i^{(t-1)}\right), \qquad (8)$$

with $\varphi_i^{(t)}$ $(i = 0,1,\dots,N)$ being the ith intermediate point along the transition path at the tth iteration ($*$ denoting the temporary values), and h the learning rate.

In the step of reparameterization of the string, the values of positions $\varphi_i^{(t)}$ $(i = 0,1,\dots,N)$ were interpolated onto a uniform mesh with the constant number of points. Prior to interpolation, the normalized length $\alpha_i^{*(t)}$ $(i = 0,1,\dots,N)$ along the path was calculated as:

$$\alpha_0^{*(t)} = 0,\ \ \alpha_i^{*(t)} = \alpha_{i-1}^{*(t)} + \frac{\left\|\varphi_i^{*(t)} - \varphi_{i-1}^{*(t)}\right\|}{\sum_{i=1}^{N}\left\|\varphi_i^{*(t)} - \varphi_{i-1}^{*(t)}\right\|},\ \ i = 1,2,\dots,N. \qquad (9)$$

Given a set of data points $(\alpha_i^{*(t)}, \varphi_i^{*(t)})$, the linear interpolation function was next used to generate the new values of positions $\varphi_i^{(t)}$ $(i = 0,1,\dots,N)$ at the uniform grid points $\alpha_i^{(t)}$ $(i = 0,1,\dots,N)$. The iteration was terminated when the relative difference $\sum_{i=0}^{N}\left\|\varphi_i^{(t)} - \varphi_i^{(t-1)}\right\|^2/N$ became small enough.

String method parameters. The string method of searching a rational MEP on the energy landscape can only guarantee the solution of local optimum and is not capable of ensuring a solution of global optimum[47]. The outcome of the string method depends on the initialization of the starting and ending points of the MEP on the energy landscape. If there are too many local minima to sample along a long MEP, the string method could retrieve a MEP solution that partly misses some local energy minima by going off the pathway. In this case, the search of an expected MEP solution can be divided into several segments of shorter MEPs connecting one another, with each MEP being defined by a closer pair of starting and ending points, which travels through a smaller number of energy minima. Another parameter affecting the MEP solution is the step size being set to explore the MEP on energy landscapes. Although a default value of 0.1 empirically recommended may work in many cases, it may be helpful to tune the step size according to the quality of the energy landscape. The value of step size can be

decreased if the computed path runs out of the energy landscape and can be increased if the path updates too slowly during iterative searching by the string method.

Energy-based particle voting algorithm. The particle-voting algorithm was designed to conduct 3D classification, particle quality control, reproducibility test and particle selection in an integrative manner. The particle-voting algorithm mainly involves two steps (Extended Data Fig. 1b). First, we count the number of votes for each particle mapped within each voting boundary, the particles within which define a new 3D cluster approximately centered on a local energy minimum of the energy landscape. One vote is rigorously mapped to one copy of the particle used in reconstructing a 3D volume and to no more than one 3D cluster on the energy landscape where the corresponding volume is located. Thus, each particle can have M votes casted for no more than M 3D clusters. If the vote is mapped outside of any 3D cluster boundary, it becomes an 'empty vote' with no cluster label. Each non-empty vote is thus labeled for both its particle identify and corresponding cluster identity. For each pair of particle and cluster, we compute the total number (V) of votes that the cluster receives from the same particle. Each particle is then assigned and classified to the 3D cluster that receives $V > M/2$ votes from this particle (Extended Data Fig. 1b). Note that after particle voting, each particle is assigned no more than once to a 3D class, with its redundant particle copies removed from this class. This strategy only retains the particles that can reproducibly vote for a 3D cluster corresponding to a homogeneous conformation, while abandoning those non-reproducible particles with divergent, inconsistent votes.

Alternative strategy of 3D clustering without particle voting. Because the particle-voting algorithm imposes strong constraints on the reproducibility of particle classification by deep manifold learning, some 3D classes might be assigned with a small number of particles that are insufficient to support high-resolution reconstruction. To remedy this limitation, an alternative, distance-based classification algorithm was devised to replace the particle-voting algorithm when there are not enough particles to gain the advantage of particle voting (Extended Data Fig. 1c). In this method, the distances of all M copies of each particle to all 3D cluster centers on the energy landscape are measured and ranked. Then, the particle is classified to the 3D cluster of the shortest distance. A threshold could also be manually preset to remove particles that are too far

away from any of the cluster centers. The distance-based classification method can keep more particles, but it ignores the potential issue of irreproducibility of low-quality particles. Thus, it is proven to be less accurate in 3D classification (Extended Data Fig. 6m-o). In other words, it trades off the classification accuracy and class homogeneity to gain more particles, which is expected to be potentially useful for small datasets or small classes. By contrast, the energy-based particle-voting algorithm imposes a more stringent constraint to select particles of high reproducibility during classification, resulting in higher quality and homogeneity in the classified particles, which is superior to the distance-based classification method (Extended Data Fig. 6s-u).

Practical consideration of particle shuffling and voting parameters. The parameter M determines the degree of implicit cross-validation of classification reproducibility, as well as the sampling densities of energy landscape. To establish reproducibility for 3D classification, M should be no less than 3. The variation of M is not supposed to change the energy landscape because it multiplies on both denominator and enumerator in the Boltzmann relation and is cancelled in the ratio of particle densities in computing the free-energy differences. Increasing M value will allow more volumes to be bootstrapped, which then leads to a higher sampling density in computing the manifold of energy landscape and potentially enhances its reconstitution quality. The default classification threshold $M/2$ is the minimal number of votes for verifying the reproducibility of 3D classification. However, a higher threshold, such as $2M/3$, will give rise to more stringent criteria in cross-validation, with a tradeoff of voting out more particles. The particle voting algorithm does not entirely eradicate misclassified particles. However, increasing either M or the classification threshold could theoretically have similar impact and improve the conformational homogeneity, because it gets less probable for a particle being misclassified to the same cluster for more times.

Several considerations may be applied to the choice of M and to set up particle shuffling and voting. First, to obtain a high-quality energy landscape, a thousand or more bootstrapped volumes are expected for datasets of more than 1-2 million particles. Second, the average particle number per volume is expected to be no less than 5000 or more to ensure that majority of volumes include sufficient particles for quality reconstructions. The expected particle number per volume must increase if the average SNR per particle is decreased. It may need to reach 10000-

20000 or more for small proteins or lower SNR datasets. Third, for a dataset of a moderate size, M can be adjusted to a higher value to mitigate the lack of image data for bootstrap.

Data-processing workflow of AlphaCryo4D. The following summarizes the complete procedure of AlphaCryo4D, with the detailed algorithm rationale explained in the previous subsections.

Input: Single-particle cryo-EM dataset after initial particle rejection of apparent false particles.

Output: Energy landscape, MEP, 3D class assignment of each particle.

Step 1. Bootstrap many 3D volumes through particle shuffling, consensus alignment and Bayesian clustering.

1. Split the particle dataset randomly to many sub-datasets, if necessary, in case of a large dataset (e.g., >150,000 particle images), for batch processing of particle shuffling and volume bootstrapping. Otherwise, skip this step if the dataset is small enough (e.g., <150,000).
2. For each sub-dataset, conduct a consensus alignment to generate initial parameters of Euler angles and translations in RELION with one or more consensus models.
3. Divide each sub-dataset into $M + 1$ groups, shuffle the sub-dataset $M + 1$ times and each time take a different group out of the shuffled sub-dataset, giving rise to $M + 1$ shuffled sub-datasets all with different collection of particles.
4. Conduct 3D Bayesian classification on all the $M + 1$ shuffled sub-datasets to generate hundreds of 3D volumes, making each particle to contribute to M different volumes.
5. Execute steps (2) to (4) using the same initial model (low-pass filtered at 60-Å) for all sub-datasets.

Step 2. Extract deep features of all volume data with the 3D deep residual autoencoder.

6. Align all 3D volumes and adjust them to share a common frame of reference.
7. Initialize the hyper-parameters of the 3D autoencoder (Supplementary Table 1).
8. Train the neural network with the 3D volume data to minimize mean square error between the decoding layer and the input by the Adam algorithm of initial learning rate 0.01.
9. Extract the 3D feature maps of all volumes from the encoding layer.

Step 3. Embed the volume data to two-dimensional manifolds through the t-SNE algorithm and

compute the energy landscape and find the MEP.

10. Calculate the pairwise similarities between volumes using their feature-map-expanded volume vectors, and randomly initialize the low-dimensional points.
11. Minimize the Kullback-Leibler divergence by the Momentum algorithm to generate 2D manifold embeddings with t-SNE.
12. Compute the energy landscape from the manifold using the Boltzmann relation.
13. Initialize searching of the MEP with a straight line between given starting and ending points.
14. Find the optimal MEP solution using the string method.

Step 4. Classify all particles through the energy-based particle-voting algorithm.

15. Sample the clustering centers along the MEP and calculate the recommended clustering radius.
16. Define the local energy minima as the 3D clustering centers and their corresponding cluster boundary for particle voting.
17. For each particle, cast a labelled vote for a 3D class when a volume containing one of the M particle copies is located within the voting boundary.
18. Count the number of votes of each particle with respect to each 3D class and assign the particle to the 3D class that receives more than $M/2$ votes from this particle.
19. Refine each 3D density map separately to high resolution in RELION or cryoSPARC using particles classified into the same 3D classes.

Blind assessments with simulated datasets. Three simulated large datasets with the SNRs of 0.05, 0.01 and 0.005, each including 2-million particle images, were employed to benchmark AlphaCryo4D and to compare its performance with alternative methods. For each synthetic dataset, the particles were computationally simulated by projecting the 20 3D density maps calculated from 20 hypothetical atomic models emulating continuous rotation of the NLRP3 inflammasome protein. The 20 atomic models were interpolated between the inactive NLRP3 structure and its hypothetical active state, which was generated through homology modeling using the activated NLRC4 structure[48] (Fig. 1c). The 20 atomic models represent sequential intermediate conformations during a continuous rotation in its NATCH domain against its LRR domain over an angular range of 90°. Each conformation is thus rotated 4.5° over its immediate

predecessor in the conformational continuum sequence. 100,000 simulated particle images per conformational state were generated with random defocus values in the range of -0.5 to -3.0 μm, resulting in 2 million particles for each dataset of a given SNR. The pixel size of the simulated image was set to the same as the pixel size (0.84 Å) of the real experimental dataset of NLRP3-NEK7 complex[48]. To emulate realistic circumstances in cryo-EM imaging, Gaussian noises, random Euler angles covering half a sphere and random in-plane translational shifts from -5.0 to 5.0 pixels were then applied to every particle image.

Each simulated heterogeneous NLRP3 dataset was analyzed separately by AlphaCryo4D and used to characterize the performance and robustness of AlphaCryo4D against the variation of SNRs. In the step of particle shuffling and volume bootstrapping, 2,000,000 particles in the dataset of any given SNR were divided randomly into 10 sub-datasets for batch processing. The orientation of each particle was determined in the initial 3D consensus alignment in RELION, which did not change in the subsequent 3D classification. In this step, the maximum number of iterations of the 3D alignment was set up as 30, with the initial reference low-pass filtered to 60 Å. 3-fold particle shuffling (indicated as × 3 below) was conducted on each sub-dataset for volume bootstrapping. The first round of maximum-likelihood 3D classification divided the input shuffled particle sub-dataset into 5 classes, each of these classes were then further classified into 8 classes. This procedure was separately executed on all shuffled particle sub-datasets. The particle shuffling and volume bootstrapping generated 1,372, 1,489, and 1,587 volumes by the datasets with SNRs of 0.05, 0.01 and 0.005, respectively. These volume data were used as inputs for deep residual autoencoder to compute low-dimensional manifolds and energy landscapes (Fig. 1d-f, Extended Data Fig. 1b, 2a, b). After searching the MEP on the energy landscapes by the string method, 20 cluster centers along the MEP were defined by the local energy minima along the MEP by approximately equal geodesic distance between adjacent minima, which represent potentially different conformations of the molecule (Extended Data Fig. 2a-c). The particle-voting algorithm was applied in every cluster to determine the final particle sets for all 3D classes. For validation of the methodology and investigation of 3D classification improvement, we labeled each bootstrapped 3D volume with the conformational state that held the maximum proportion of particles in the class and computed its 3D classification precision as the ratio of the particle number belonging to the labelled class versus the total particle number in the volume (Fig. 1h, Extended Data Fig. 3a, b).

Analysis of algorithm mechanism. To understand how the 3D classification accuracy is improved by AlphaCryo4D, we analyzed the algorithmic mechanism by a number of control experiments (Extended Data Fig. 2d-g, 7). In total, we conducted 24 conditional tests using 10 variations of algorithmic design by removing or replacing certain components of AlphaCryo4D. First, by removing the entire component of deep manifold learning, the distributions of high-precision 3D classes obtained by *M*-fold particle shuffling and voting alone drop by ~60% relative to that from complete AlphaCryo4D processing (Extended Data Fig. 6d-f). In these control experiments, the particle voting was achieved through counting the votes of a particle against the same conformers classified by RELION via computing the Intersection-over-Union (IoU) metric after the step of volume bootstrapping. Similarly, by removing the 3D deep residual autoencoder but keeping the manifold learning by t-SNE for energy landscape reconstitution, the distributions of high-precision 3D classes is reduced by ~15% at the SNR of 0.005 (Extended Data Fig. 6p-r). The effect of accuracy degradation is less prominent at a higher SNR (0.05 or 0.01), indicating that the unsupervised feature learning using deep residual autoencoder promotes the tolerance of the algorithm against higher noise level in the data. Together, these control tests suggest that deep manifold learning plays a crucial role in improving 3D classification accuracy with low SNR data.

Next, we replaced t-SNE with four other algorithms in manifold learning step, including two classic linear dimensionality reduction techniques, principal component analysis (PCA)[13] and multidimensional scaling (MDS)[49], and two nonlinear dimensionality reduction algorithms, isometric mapping (Isomap)[50] and locally linear embedding (LLE)[51]. We applied the four algorithms to reduce the dimensionality of the same sets of bootstrapped volume data. Although being capable of differentiating a portion of ground-truth conformers, the resulting 2D mappings by the four techniques exhibit considerable overlap between adjacent conformers and are incapable of distinguishing all 20 conformers of ground truth, thus inevitably missing many conformers (Extended Data Fig. 2d-g). The PCA and Isomap approximately missed half of the ground-truth conformers, whereas MDS and LLE both missed around 70% of the ground-truth conformers. The inferior performance of these techniques is consistent with the original control experiments conducted by the t-SNE developers[45].

Further, we conducted control experiments to evaluate the impact of particle voting on the 3D

classification accuracy by replacing the particle voting component with a clustering strategy that directly classifies each particle to the cluster of the nearest clustering center (Extended Data Fig. 1c). In this case, the distributions of high-precision 3D classes are reduced by ~15-30% (Extended Data Fig. 6m-o). The reduction is more prominent at the lower SNR. This indicates that the particle voting considerably improves the 3D classification accuracy but is less impactful than the component of deep manifold learning.

Last, by tracking the statistics of classification precisions step by step, we evaluated how the 3D classification accuracy is improved over the intermediate steps of AlphaCryo4D (Extended Data Fig. 6g-l). We found that the steps of particle shuffling, defining cluster boundaries on the energy landscapes and energy-based particle voting contribute to ~16%, ~20% and ~40% improvements in the distributions of high-precision 3D classes, respectively. Taken together, these results indicate that all components contribute to the improved performance of AlphaCryo4D in 3D classification toward the atomic level. Neither deep manifold learning nor particle voting alone is sufficient to achieve the present level of 3D classification accuracy.

Comparisons with alternative methods. Using 3D classification precision as a benchmark indicator, the performance of AlphaCryo4D preprocessed by standardization or 5 Å low-pass filtering was compared with several other methods: (1) the conventional maximum-likelihood-based 3D (ML3D) classification in RELION[3,4,6] with and without a hierarchical strategy, (2) the conventional heterogeneous reconstruction in cryoSPARC[2], (3) the 3DVA algorithm with 2 and 3 principal components (PCs) in cryoSPARC[11] and (4) the deep generative model-based cryoDRGN[12]. A total of 24 comparative tests by these alternative methods have been conducted blindly using the three simulated datasets, which includes 6-million images in total. In all our comparative tests on AlphaCryo4D and the alternative methods, the ground-truth information of particle orientations, in-plane translations and conformational identities were completely removed from the methods being tested and were not used for any steps of data processing or training. The ground-truth conformational identities of particles were only used as validation references to compute the 3D classification precisions of the blind testing results.

For the tests using ML3D in RELION, we classified all particles directly into 20 classes and hierarchically into 4 × 5 classes, which first divided the dataset into 4 classes, with each class further classified into five sub-classes (Extended Data Fig. 3g, h). For testing the conventional

discrete ab initio heterogeneous reconstructions in cryoSPARC, each synthetic dataset was directly classified into 20 classes without providing any low-resolution reference model (Extended Data Fig. 3f). For comparison with the 3DVA algorithm in cryoSPARC, we first did the blind consensus alignment of the entire dataset to find the orientation of each particle. Then the alignment and the mask generated from the consensus reconstruction were used as inputs into the 3DVA calculation, with the number of orthogonal principal modes being set to 2 or 3 in 3D classification. The 3D variability display module in the cluster mode was used to analyze the results of 3D classification. For blind tests with cryoDRGN, a default 8D latent variable model was trained for 25 epochs. The encoder and decoder architectures were 256 × 3, as recommend by the cryoDRGN developers[12]. The particle alignment parameters prior to cryoDRGN training were obtained by the same blind consensus refinement in RELION used for other parallel tests. The metadata of 3D classification precisions as well as the 3D density maps from all the algorithms applied on the three simulated datasets were collected to conduct the statistical analysis.

Computational costs. Although the computational cost of AlphaCryo4D is generally higher than the conventional approach, it does not appear to increase drastically and likely falls in an affordable range, while reducing the average cost of computation per conformational state. In a nutshell, we can have a brief comparison of the computational efficiency on the simulated 2-million image dataset with an SNR of 0.01. In the step of 3D data bootstrapping, we split the dataset into 20 subsets, which contained 100,000 particles each. The 3D consensus alignment of all 2,000,000 particles cost about 75 hours using 8 Tesla V100 GPUs interconnected with the high-speed NVLink data bridge in a NVIDIA DGX-1 supercomputing system. Within each subset, the 3D Bayesian classification for one leave-one-group-out dataset cost about 2.5 hours using 320 CPU cores (Intel Xeon Gold 6142, 2.6 GHz, 16-core chip), so the total time spent in one subset was about 10 hours using 320 CPU cores in an Intel processor-based HPC cluster. In addition, it spent about 3 hours to extract feature via deep neural network using 8 Tesla V100 GPUs of the NVIDIA DGX-1 system. In contrast to about 160 hours cost in traditional classification methods, this approach cost a little more than 200 hours using 8 Tesla V100 GPUs and 320 CPU cores. Taken together, these observations suggest that the computational cost and efficiency of AlphaCryo4D are within the acceptable range considering the output of more high-

resolution conformers yielded by the procedures (Extended Data Table 2). The benefits likely outweigh the tradeoff of computational costs, particularly for large datasets that equally slow down any methods.

Relation to other techniques proposed to analyze structural heterogeneity. Our particle-shuffling and Bayesian bootstrap strategy is conceptually distinct from the stochastic bootstrap method previously proposed for estimating 3D variance map[92] or performing 3D PCA[13] to detect conformational variability. One essential difference lies in that our bootstrap requires that only particles of similar or identical conformations are grouped together via likelihood-based similarity estimation instead of being combined or resampled in a stochastic manner[1,18]. Another difference is that the goal of *M*-fold particle shuffling is to improve the quality of reconstituted energy landscape and to enable energy-based particle voting for intrinsic cross-validation of 3D classification. Furthermore, AlphaCryo4D does not assume any prior knowledge regarding the conformational landscape, its continuity and topology, as well as chemical composition of the macromolecular system. Our approach allows multiple consensus models being used for optimizing the particle alignment during Bayesian bootstrap. This can theoretically deal with more complicated conformational dynamics when no stable core structure is available to guide the consensus alignment over a single model.

The recently proposed 3DVA algorithm[11] implementing probabilistic PCA-based 3D classification in cryoSPARC, deep generative model-based cryoDRGN[12] and Laplacian spectral volumes method[21] were designed to tackle the problem of recovering continuous conformational changes without reconstituting energy landscapes. cryoDRGN uses a variational autoencoder (VAE) to map raw particle images to a feature space. The Laplacian spectral volume method is based on graph Laplacian representations of conformational heterogeneity and may also be considered as a generalization of PCA[21]. This algorithm was only examined with synthetic datasets. In addition, methods of integrating traditional molecular dynamics (MD) simulations with cryo-EM data, with or without use of deep learning, have been recently proposed to extract protein dynamics information[93,94]. These methods cannot be used to recover structures and dynamics hidden in the particle sub-dataset not being accurately classified and not used in final cryo-EM reconstructions and are limited by the size of biomolecules that can afford MD simulations.

The focused classification strategy implemented in our approach is analogously derived from the multi-body or focused refinement[4,23,24], and can be combined with flexible refinement[27,28], non-uniform refinement[29] and normal mode analysis[25,26]. These approaches may be limited to the analysis of local conformational dynamics. A comprehensive review on the methods previously proposed to study continuous conformational changes is available[26].

Applications to experimental cryo-EM datasets. Six experimental cryo-EM datasets were processed using AlphaCryo4D to examine its applicability in processing experimental cryo-EM data, as described in detail in the following. Five of the six datasets are publicly available in the Electron Microscopy Public Image Archive (EMPIAR) database at EMDataResource.org under accession codes EMPIAR-10028, EMPIAR-10180, EMPIAR-10076, EMPIAR-10090, and EMPIAR-10669.

Application to the experimental cryo-EM dataset of the *Pf*80S ribosome. The *Pf*80S ribosome dataset (EMPIAR-10028) contains 105,247 particle images. First, the alignment of all particles was refined with a global mask in RELION. To bootstrap enough 3D volumes for the energy landscape, we set $M = 7$ in the particle shuffling step, resulting in 79 volumes in total with a global mask and 15 Å resolution limit in the expectation step. Moreover, the regularization coefficient was set to 0.001 when training the 3D residual network with this small dataset. All volumes were then low-pass filtered at 5 Å prior to manifold embedding. Five clusters were obtained from the energy landscape, which had 66,035, 21,482, 9,922, 6,424, 657 particle images respectively after voting. All these clusters were refined independently using RELION.

Application to the experimental cryo-EM dataset of the pre-catalytic spliceosome. The pre-catalytic spliceosome dataset (EMPIAR-10180) with the particle number of 327,490 shows high conformational dynamics. The consensus alignment of these particles in the original dataset was used in the first step. $M = 7$ particle shuffling was utilized to bootstrap 160 volumes with a soft global mask and 15 Å resolution limit in the expectation step. The regularization coefficient of 0.001 was set to train the 3D Autoencoder with this dataset. All volumes were then low-pass filtered at 5 Å prior to manifold embedding. Based on the energy landscape, we obtained 14

classes with the particle number of 15,749, 17,703, 27,522, 13,820, 20,841, 14,989, 16,377, 22,047, 25,049, 21,327, 18,949, 6,157, 9,790 and 36,878 after voting. Each class was refined independently using RELION.

Application to the experimental cryo-EM dataset of the membrane channel protein complex RyR1. To determine the dynamics of the peripheral domains of the RyR1, density subtraction and focused energy landscape were applied in the further analysis of state 1 conformation[57]. Due to the low SNR of the particles, a large $M = 30$ was used in the particle shuffling in AlphaCryo4D. After symmetry expansion of particles of state 1, a total of 1,557,168 particles with the core domain subtracted were utilized to compute the energy landscape with 310 volumes. Three 3D classes with the particle numbers of 196,669, 848,655 and 189,796 were obtained by the particle-voting procedure on the energy landscape. All these 3 classes were merged together for high resolution refinement with a mask of the peripheral domains using RELION. To further improve the quality of the density map, a zoom-in energy landscape with $M = 22$ was computed to result in a final particle set with the number of 1,198,946. Furthermore, as the P2 domain showed high flexibility, we subtracted the non-P2 high-resolution components from the particles and computed the P2 domain-focused energy landscape using $M = 23$. This energy landscape resulted in three classes with the particle numbers of 14,089, 1,108,282 and 57,381, which was combined to refine the final density map of the P2 domain with a specific mask using RELION.

Application to the experimental cryo-EM dataset of the bacterial 50S ribosomal intermediates. The 131,899 particle images of the 50S ribosomal large subunit (EMPIAR-10076) are highly heterogeneous. We refined the alignment of all particles with a global mask of 50S using RELION. In the first step, 119 bootstrapped volumes were utilized to calculate the energy landscape. $M = 7$ particle shuffling was conducted to generate these volumes without any mask. The regularization coefficient of 0.001 was set when training the deep neural network. All volumes were then low-pass filtered at 5 Å prior to manifold embedding. Then the first energy landscape resulted in nine clusters with the particle number of 766 (A), 15,129 (B), 1,236, 322, 670, 22,037, 25,445 (D), 33,976 (E) and 25,115 (C) respectively. For the clusters of states B-E, four zoom-in energy landscapes were then computed to discover more sub-states with 46, 56, 66

and 72 bootstrapped volumes respectively. In the zoom-in energy landscape calculation, $M = 7$ was set in particle shuffling for bootstrapping 3D volumes with a global mask and 20 Å resolution limit in the expectation step. Then the regularization coefficient of 0.001 was used to train the deep residual network. All volumes were then low-pass filtered at 5 Å prior to manifold embedding. After particle voting, the energy landscape of state B resulted in three sub-clusters with the particle number of 3,494, 3,600 and 7,803. The state C energy landscape resulted in five sub-clusters with the particle number of 2,444, 7,442, 711, 7,633 and 5,557. The state D energy landscape resulted in five sub-clusters with the particle number of 6,231, 8,785, 2,123, 3,718 and 4,092. And the energy landscape of state E resulted in eight sub-clusters with the particle number of 4,783, 1,588, 5,891, 3,620, 7,352, 5,265, 357 and 2,241. All these clusters were refined using RELION.

Application to the substrate-free proteasome dataset. The dataset of the ATPγS-bound human 26S proteasome[59] (EMPIAR-10090) contains 455,680 particles and was previously used to determine six coexisting conformational states of the substrate-free proteasome. The cryo-EM data were recorded with the super-resolution pixel size of 0.75 Å. The size of ATPγS-bound proteasome dataset was enhanced to 455,680 × 3 particles and was divided into 4 sub-datasets for volume bootstrapping. 160 volumes were bootstrapped to compute the energy landscape the substrate-free ATPγS-bound proteasome (Fig. 4a).

Applications to the substrate-engaged human 26S proteasome dataset. The substrate-engaged human proteasome dataset[30] (EMPIAR-10669) includes 3,254,352 RP-CP particles (combined with particle images from both the DC and SC proteasomes) in total, with the super-resolution counting mode pixel size of 0.685 Å and the undecimated box size of 600 × 600 pixels. Sample preparation, cryo-EM imaging and data collection condition as well as data pre-processing prior to 3D classification by AlphaCryo4D have been previously described in detail[30]. All particles were binned twice down to 300 × 300 pixels prior to data processing. To proceed with AlphaCryo4D, the whole dataset was randomly split into 32 sub-datasets, each with 100,000 particles. Particle shuffling and volume bootstrapping procedures were conducted on each sub-dataset, which enhanced the total data size to 3,254,352 × 3 particles and yielded 32 × 4 × 10 = 1280 volumes in total that were used to compute the energy landscape of the substrate-

engaged proteasome (Fig. 4b, c). Particle-voting on this energy landscape detected previously missing states (such as $E_{D0.2}$, $E_{D0.3}$, $E_{D1.1}$ and $E_{D1.2}$, etc.) and yielded previously determined states with improved map quality, such as E_{A2} and E_{D2}. Guided by this initial discovery of new conformational states, extensive focused 3D classifications in AlphaCryo4D were performed, as described in detail in the next section, to visualize novel ubiquitin-binding sites and hidden dynamics of proteasomal AAA-ATPase motor during single-nucleotide exchange.

Since many protein complexes exhibit profound conformational changes in different local regions, we implemented a focused classification strategy that applies a local 3D mask[6] throughout the entire AlphaCryo4D procedure, often being executed as an iterative step after initial AlphaCryo4D analysis without any 3D mask (Fig. 1a). To apply AlphaCryo4D to process the experimental dataset of the substrate-engaged human 26S proteasome[30], a focused 3D classification protocol using AlphaCryo4D was implemented to enhance its capability in discovering transient states of extremely low particle populations and in solving highly dynamic components that may be too small to be analyzed by conventional methods[2-4,6]. To study the local features of ubiquitin binding and substrate interactions, 3D masks corresponding to the local densities of interest were applied to all bootstrapped volume data, and the subset of interest on the energy landscape was extracted to conduct a zoomed-in analysis. In this case, the energy landscape calculated via masked volumes is expected to reflect the local structural variations of interest (Extended Data Fig. 9a-k).

To detect novel ubiquitin-binding sites in the proteasome, local masks were applied to focus on the features around the putative locations of ubiquitin chains. After particle shuffling for M = 3 folds, 732,666 × 3 particle images of the state E_A and E_B were utilized for computing the zoomed-in energy landscape with 280 volumes. These particles were first aligned in a consensus alignment without any mask or resolution limit to restrict alignment parameter calculations. On the focused energy landscape, only the region within the soft mask of the ubiquitin and its binding site was applied in the volume bootstrapping step, with the signals of the CP of the proteasome subtracted from each raw particle when dealing with the state E_A and E_B. Accordingly, the particle diameter applied in the CP subtracted proteasome was shrunk to 274 Å, instead of 411 Å in the complete 26S proteasome. When the region around RPN2 masked for focused 3D classification to search for potential ubiquitin-like features, the total particle numbers were enhanced to 208,700 × 3 (E_{A1}), 240,475 × 3 (E_{A2}), 260,021 × 3 (E_B) and 129,492 × 3 (E_{D2}).

The energy landscapes were computed with 80, 80, 80 and 40 masked volumes bootstrapped by upper resolution bounds of 15, 15, 15 and 20 Å, respectively. Several 3D classes with the particle numbers of 192,219, 2,539, 3,982 and 609 were obtained by the particle-voting procedure on the zoomed-in energy landscape of state E_{A1}, while the case of state E_{A2} gave four classes containing 147,108, 5,842, 5,500 and 33,083 particles (Extended Data Fig. 9e, f). Likewise, the particle numbers of 3D classes on the zoomed-in energy landscape of state E_B were 173,931, 5,117 and 9,754 (Extended Data Fig. 5g); and the zoomed-in energy landscape of state E_{D2} resulted in four classes having 61,580, 6,192, 24,674 and 1,673 particles, respectively (Extended Data Fig. 9h).

To detect and improve RPN1-bound ubiquitin densities, a local mask around RPN1 was used, with 15 Å resolution limit applied in the volume-bootstrapping step. The zoomed-in energy landscape around state E_{A1} computed from 120 volumes resulted in five classes with the particle numbers of 128,161, 7,754, 6,557, 6,826 and 3,937. Five 3D classes on the zoomed-in energy landscape of E_{A2}, also computed with 120 volumes, including 81,612, 27,483, 12,247, 26,029 and 12,777 particles, were analyzed (Extended Data Fig. 9b, c). Based on the subclass of E_{A1} that contains 128,161 particles, we continued to refine the alignment locally to the component of RPN1. Then a 3D mask around RPN1 was applied to bootstrap 40 volumes with 20 Å resolution limit in the expectation step. After voting on the energy landscape, two classes with the particles number of 43,957 and 43,620 were generated, one of which exhibited a ubiquitin-like density on the T1 site of RPN1 and was designated state $E_{A1.1}$ (Extended Data Fig. 9d). In the case of the ubiquitin-binding site on the α5 subunit of the CP, five 3D classes with the particle numbers of 136,071, 19,351, 4,100, 131,570 and 4,932 were generated by particle voting on the zoomed-in energy landscape of E_D computed with 120 volumes, which were bootstrapped from 326,409 × 3 particles via the mask containing this ubiquitin-binding site and no high-resolution restriction (Extended Data Fig. 9i). During these analyses, all masked volume-bootstrapping steps were performed for 50 iterations, whereas unmasked 3D classifications were performed for 30 iterations with the initial reference low-pass filtered to 60 Å. At the last step, 3D refinement and reconstruction of each class was done to finalize the high-resolution structure determination.

To exploit intermediate conformations during the substrate translocation initiation, a zoomed-in energy landscape was computed for the focused 3D classification procedure. First, the energy landscape of all 3,254,352 × 3 particles was computed and analyzed. Then the zoomed-in energy landscape of the states E_C and E_D containing 2,521,686 × 3 particle images were computed from

1,000 volumes that were bootstrapped without masking or resolution restriction (Extended Data Fig. 9j). For further zoomed-in processes, the particles around the translocation initiation states (from E_C to E_{D0}) on this energy landscape were extracted by the distance-based 3D classification strategy. Subsequently, these particles were shuffled 3 times to the number of 410,098 × 3 to generate 160 volumes for the computation of the zoomed-in energy landscape with the RP mask and no resolution restriction imposed in the volume bootstrapping, which led to the 3D classes E_{C1}, E_{C2}, and $E_{D0.1}$ with the particle number of 68,506, 77,545, and 25,755, respectively (Extended Data Fig. 9k). Moreover, state $E_{D0.3}$ came from two clusters with 86,415 and 83,107 particles merged together on the unmasked energy landscape, and the classes $E_{D0.2}$ and $E_{D1.1}$ including 105,081 and 51,808 particles, respectively, were also obtained after particle voting on this energy landscape. In the local areas of the state E_{D1} on the energy landscape, two clusters with 80,841 and 94,273 particles were merged to reconstruct the high-resolution structure of state $E_{D1.2}$, while states $E_{D1.1}$ and $E_{D1.3}$ with 40,696 and 145,506 particles, respectively, were separated out from the clusters on this unmasked energy landscape. In these studies, the maximum number of iterations in the 3D classifications of volume bootstrapping was 30, with the initial reference low-pass filtered to 60 Å routinely.

Beyond resolving significantly more conformers from the same dataset, AlphaCryo4D also allows us to push the envelope of the achievable resolution of dynamic components and key metastable states due to its advantage in keeping more particles without sacrificing conformational homogeneity. To improve the resolution of state E_{D2}, a rigorous particle voting procedure was performed in each cluster on the unmasked energy landscape (Fig. 4b, c), which improved the overall quality of particle dataset while classifying significantly more particles to this E_{D2} conformational class. Next, we conducted a focused 3D classification by using AlphaCryo4D to calculate the zoomed-in energy landscape of the entire E_{D2} dataset to detect any potential conformational changes within this class. Five clusters near the energy wells of state E_{D2}, whose particle numbers were 154,661, 243,799, 198,221, 179,257 and 106,047, were investigated by separate high-resolution 3D refinement and reconstruction. The resulting refined density maps at 3.0-3.5 Å were found to be all in an identical conformation as previously reported[30]. These five clusters were then merged together to refine a final 3D density map of state E_{D2}. Thus, AlphaCryo4D allowed us to obtain a new E_{D2} dataset that appears to retain high conformational homogeneity while including three-fold more particles, which improves its 3D

reconstruction from the previously published resolution at 3.2 Å to 2.5 Å, measured by gold-standard Fourier shell correlation (FSC) at 0.143-cutoff. By contrast, in our previous work[30], we had attempted to include more E_{D2}-compatible 3D classes from 3D maximum-likelihood classification in RELION, which resulted in lower resolution at 3.3 Å by doubling the E_{D2} dataset, indicating that including more particles by conventional methods gave rise to higher heterogeneity in the class presumably due to increased misclassification.

Analysis of doubly and singly capped proteasomes. Previous cryo-EM analyses of the proteasomal states were nearly all focused on the RP-CP subcomplex, by converting double-cap (DC) proteasome particles to pseudo single-cap (SC) proteasome, which was always mixed with true SC particle images, in order to improve the resolution with conventional methods[30,59,73,83,95,96]. As a result, the cryo-EM density of the CP gate in the other side distal to the RP in the RP-CP subcomplex has been an average of closed and open states with unknown stoichiometry. This has made it impossible to appreciate the difference between the SC and DC proteasomes in the previous studies. To remove this ambiguity and local conformational heterogeneity, we reconstructed 8 density maps of the pure SC proteasome and 36 density maps of the pure DC proteasome including all possible combination of the 8 major RP-CP states by AlphaCryo4D classification. Both CP gates in the 44 states of the pure SC and DC reconstructions are expected to be in a homogeneous conformation. To reconstruct the density maps of the DC proteasome, pseudo-SC particles corresponding to 8 different RP states were voted out from the energy landscape, with the sub-states belonging to the same RP state combined together. Then the particles with one RP-CP were extracted from the particle stacks to refine the RP and CP density maps with RP and CP contour masks applied in the local search step in RELION, respectively. For each reconstruction of the DC proteasome, two RP density maps on the opposite sides of the CP were expanded from 600 × 600 × 600 pixels to 800 × 800 × 800 pixels using RELION and aligned by the density of CP in Chimera before being merged together. All 8 SC and 29 DC states were refined to 3.7-4.7 Å, with 7 DC states limited to 6-9 Å due to their transient nature and extremely low population (Extended Data Fig. 11b, e, f).

Counting the particle number of the 36 states of the DC proteasome, we could derive the distribution of the DC states, p_{ij}, and that of RP p_i only, where i and j denote the states of two RPs in the same DC proteasome. Based on the value of p_{ij} and total particle number, the

experimental state distribution of the DC proteasome was plotted for quantitative analysis in two dimensions with respect to the states of two RPs in the same DC proteasome (Fig. 4e). If the states of two RPs of the same DC proteasome were independent of each other, the predicted state distribution $\tilde{p}_{ij}$ of the DC proteasome can be calculated as:

$$\tilde{p}_{ij} = p_i \cdot p_j$$

By using the total particle number and experimentally measured p_i, the 2D state distribution of the DC proteasome with two uncoupled RPs was calculated for further comparison to investigate the conformational entanglement effect of the two RP states (Fig. 4f-h).

Atomic model building and refinement. To build the initial atomic model of the newly discovered states, we use previously published proteasome structures as a starting model and then manually improved in Coot[97]. For the conformational states at resolutions lower than 5 Å, pseudo-atomic modelling was conducted in the following steps. First, each subunit was fitted as a rigid body, often in UCSF Chimera[98] as well as in Coot[97]. Then, the local structural domains and secondary structure elements were fitted as a rigid body, with linking loops flexibly fitted and modelled. After the mainchain structures were well fitted, no further fitting of the sidechain rotamers were pursued due to insufficient resolution except for energy minimization to remove unrealistic side-chain clashes. For the conformational states at resolutions higher than 5 Å, the mainchain traces were first fit in Coot. Then, extensive fitting, adjustment and optimization of sidechain rotamers were conducted through local real-space refinement and manual rectification in Coot to the degree commensurate to the map resolution and quality. Atomic model refinement was conducted in Phenix[99] with its real-space refinement program[100]. We used both simulated annealing and global energy minimization with secondary structure, rotamer and Ramachandran constraints. Partial rebuilding, model correction and density-fitting improvement in Coot[97] were iterated after each round of atomic model refinement in Phenix[99]. The improved atomic models were then refined again in Phenix, followed by rebuilding in Coot[97]. The refinement and rebuilding cycle were often repeated for three rounds or until the model quality reached expectation (Extended Data Table 3, Supplementary Information). All figures of structures were plotted in Chimera[98], PyMOL[101], or ChimeraX[102]. Local resolutions of cryo-EM density maps were evaluated using ResMap[103] or Bsoft Blocres program[104]. Structural alignment and comparison were performed in both PyMOL and Chimera. Electrostatic surfaces were calculated

by APBS plugin[105] in PyMOL.

Given a large range of variations of local resolution in an expanded number of proteasome states and substates, as well as reconstructions with different stoichiometric ratios of RP versus CP, additional caveats and cautions were practiced in order to avoid misinterpretation and overall-fitting of atomic and pseudo-atomic models. First, when fitting the ubiquitin structure to the low-resolution local density, we consider both the existing homologous structural models as a modelling reference. Specifically, the fitting of diubiquitin to the density on RPN2 in states $E_{A2.1}$, $E_{A2.2}$, and $E_{B.2}$, we took the NMR structure of the yeast diubiquitin-bound Rpn1 T1 site as a reference, because of the high structural homology between RPN1 and RPN2. For the fitting of RPN10-bound density, we also took into account the electrostatic complementarity (Extended Data Fig. 10h-o). Second, even within an overall high-resolution cryo-EM map, it often presents certain local densities at a lower local resolution due to poor occupancy of ligands or flexibility of less well-folded segments. For example, for the nucleotide in RPT5 in states $E_{D0.1}$ to $E_{D1.3}$, the high-resolution atomic model of ADP refined from other states like E_{C1} were used to fit these lower-local-resolution densities as a rigid body. However, the B-factors of the fitting atoms of ADP in RPT5 were reported to be more than twice higher than those of other nucleotides, indicating its partial occupancy or unstable binding. Thus, ADP fitting in RPT5 should be regarded tentative for the most or for the purpose of evaluation of nucleotide states rather than a reliable high-resolution atomic modeling. Similarly, the atomic models of these locally low-resolution features were only further adjusted with strict stereochemical constraints when the map resolution and features permit. Third, for the CP gate at the other side opposite to the RP in the RP-CP subcomplex reconstructed from a mixture of true and pseudo-SC proteasomes (converted from the DC proteasomes), their local densities in states $E_{D0.1}$ to E_{D2} appear to be an average of the open and closed states, and thus are weaker in amplitude than the rest of the CP. In the atomic model building of these reconstructions, we chose to model them with the closed CP gate state as long as their density resolution allows atomic modeling. This issue was completely solved when we separated the 3D classes of the SC and DC proteasomes and reconstructed their states respectively, where the atomic modelling of any CP gate no longer suffers from such ambiguities.

Quantum mechanical calculation. First-principles calculations on the electronic structures of the substrate-proteolytic site complex were conducted as described. In brevity, local coordinates of the substrate-bound CP near the catalytic site at residue Thr1 were taken out from the experimentally fitted atomic models of CP and the substrate, in both E_A and E_{D2} conformations. The proteolytic site is the Thr1 residue at the β2 subunit but extended along the chain for another two residues to make the system less finite. Hydrogen atoms are added to complete the residues, and to saturate the system boundaries. The structural relaxations and electronic structure calculation were carried out using the density functional theory (DFT) method with norm-conserving pseudopotentials as implemented in the SIESTA code[106]. To set up the initial conditions of the finite system, the hydrogen atom in the hydroxy group of Thr1 is manually relocated to the nearby nitrogen atom in the same residue, corresponding to the initial proton transfer in the catalytic mechanism. All hydrogen atoms were subsequently optimized without symmetric restrictions using the conjugate gradient algorithm and a 0.04 eV/Å maximum force convergence criterion while keeping the rest of the system fixed. The generalized gradient approximation (GGA) exchange-correlation density functional PBE[77] was employed together with a double-zeta plus polarization basis set, and a mesh cutoff of 200 Ry (corresponds to 0.23 Å smallest grid size). The charge density and charge density difference contour maps were plotted with the Siesta Utility programs denchar and Python, respectively.

Data availability. The experimental data used in this study are based on the public datasets available from the raw cryo-EM data online repository EMPIAR (https://www.ebi.ac.uk/empiar/) through the accession numbers EMPIAR-10028, EMPIAR-10180, EMPIAR-10076, EMPIAR-10090, and EMPIAR-10669.

Code availability. Source code of AlphaCryo4D is freely available at http://github.com/alphacryo4d/alphacryo4d/.

Acknowledgments. The authors thank Q. Ouyang, A. Goldberg, D. Finley, S. Elsasser, M. Kirschner, Y. Lu, and for constructive discussions. This work was funded in part by the Natural Science Foundation of Beijing Municipality (grant No. Z180016/Z18J008), the National Natural Science Foundation of China (grant No. 11774012) and a gift academic grant from Intel

Corporation. The cryo-EM data were collected at the Cryo-EM Core Facility Platform and Laboratory of Electron Microscopy at Peking University. The data processing was performed in the High-Performance Computing Platform at Peking University.

Author contributions. Y.Mao and Z.W. conceived this study and devised the methodology. Z.W. developed the prototype source code of the AlphaCryo4D system, conducted numerical studies of the system using the synthetic datasets, and built the DC proteasome models. S.Z., Y.D., Z.W. and W.L.W. contributed to the experimental cryo-EM data acquisition. Z.W., E.C., and S.Z. analyzed the experimental cryo-EM datasets and refined the density maps. C.L. and C.C.Yin contributed to the sample preparation and data collection on the RyR1 dataset. W.L.W. conducted the quantum mechanical simulation and contributed to the improvement of the AlphaCryo4D source code. Y.Ma contributed to early studies in algorithm design and testing. Y.Mao supervised this study, verified the density maps, built and refined the atomic models, interpreted the data and wrote the manuscript with inputs from all authors.

Competing interests. The authors declare no competing financial interests.

Additional information

Supplementary information is available for this paper online.

Correspondence and requests for materials should be addressed to Y.Mao.

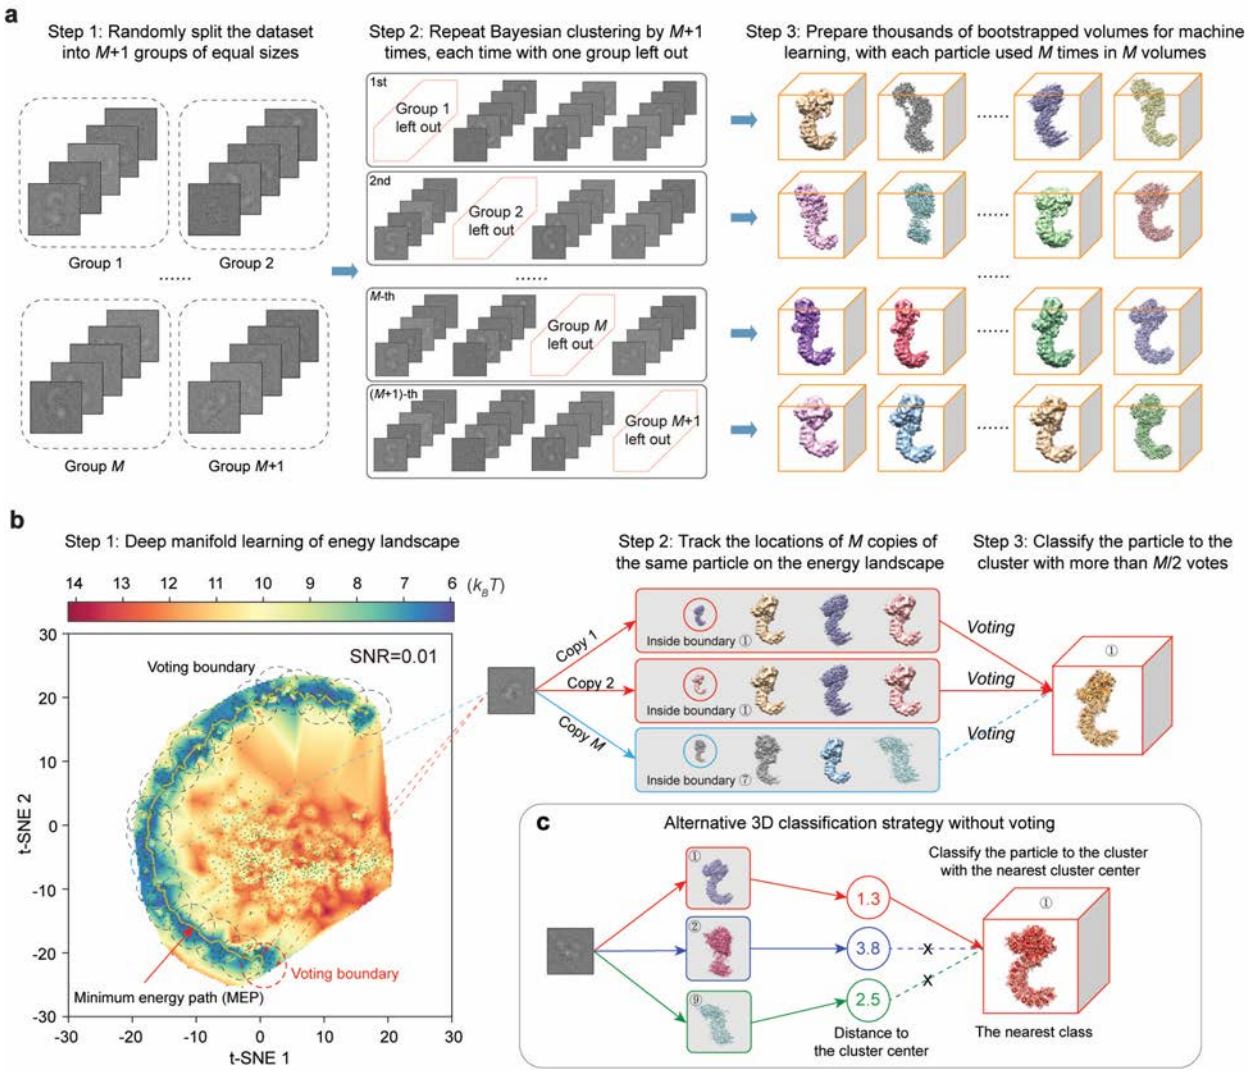

Extended Data Fig. 1. Detailed algorithmic design of particle shuffling and voting in AlphaCryo4D. a, Schematic showing the method of *M*-fold particle shuffling and bootstrapping of 3D volumes. In the step 1, all particles are split randomly into *M* + 1 groups equally. Then the step 2 carries out the Bayesian clustering for reconstructions of 3D density maps within each of the *M* + 1 particle sets that are shuffled via the leave-one-group-out approach. After these two steps, thousands of 3D volumes are generated for the subsequent 3D deep learning, with each particle contributing to *M* volumes. **b,** Schematic showing the algorithmic concept of energy-based particle voting for 3D classification. The left panel shows the energy landscape obtained by deep manifold learning. After clustering along the minimum energy path, all *M* locations of each particle on the energy landscape can be tracked to cast *M* votes. A vote of the particle is only counted for the cluster when it is mapped within the voting boundary of the cluster, as

indicated by the circles marked on the energy landscape. Eventually, this particle is classified to the 3D cluster with over $M/2$ votes from this particle. **c,** Alternative distance-based 3D classification method as a control in the analysis of algorithmic performance of particle voting. Instead of particle voting, each particle is directly classified to the cluster with the nearest clustering center among the M volume data points in the distance-based 3D classification strategy.

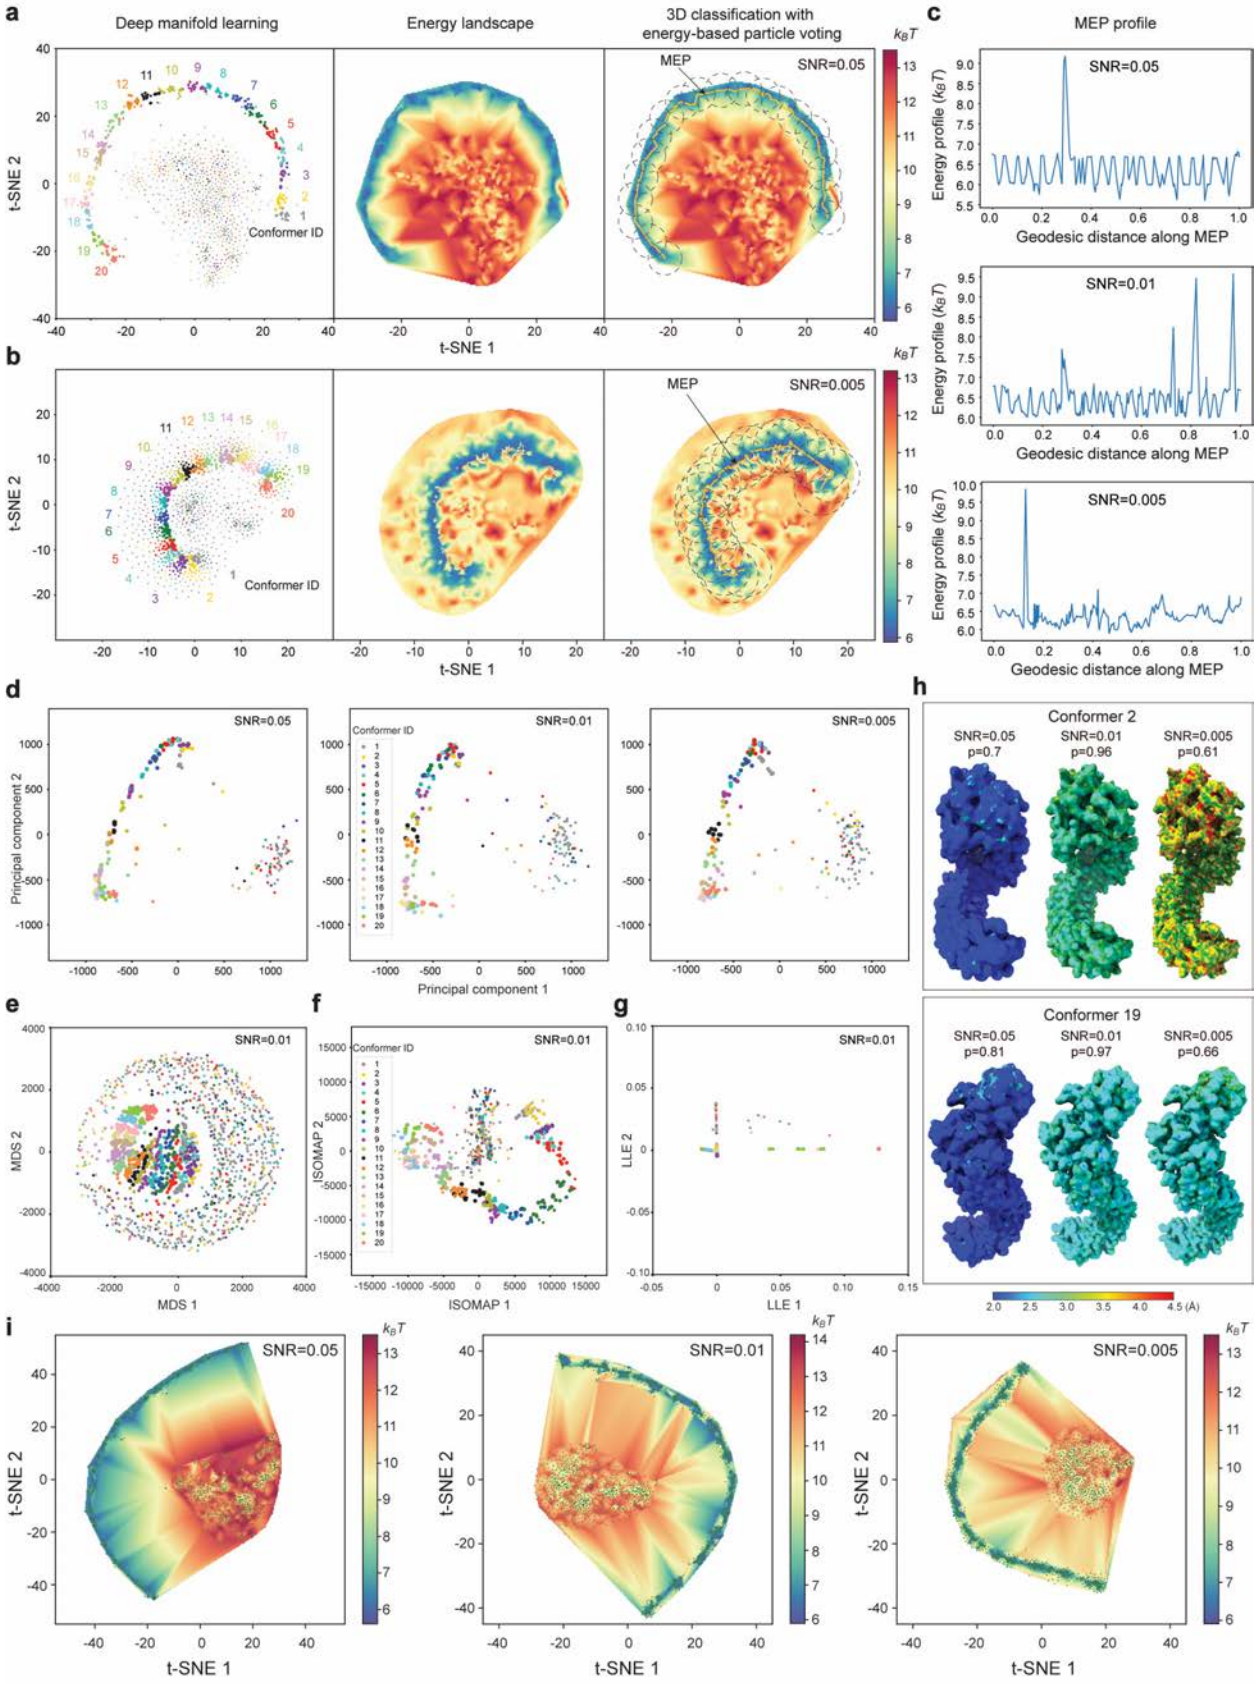
a
Deep manifold learning
Energy landscape
3D classification with energy-based particle voting
MEP
SNR=0.05
Conformer ID
t-SNE 1
t-SNE 2
b
SNR=0.005
c
MEP profile
SNR=0.05
SNR=0.01
SNR=0.005
Energy profile (k_BT)
Geodesic distance along MEP
d
Principal component 1
Principal component 2
e
MDS 1
MDS 2
f
ISOMAP 1
ISOMAP 2
g
LLE 1
LLE 2
h
Conformer 2
SNR=0.05 p=0.7
SNR=0.01 p=0.96
SNR=0.005 p=0.61
Conformer 19
SNR=0.05 p=0.81
SNR=0.01 p=0.97
SNR=0.005 p=0.66
i

Extended Data Fig. 2. Blind assessments of AlphaCryo4D and its comparison with the 3D PCA method using the simulated heterogeneous NLRP3 datasets of different SNRs. a and **b,** Reconstruction of the energy landscape of the simulated NLRP3 datasets at SNRs of 0.05 (**a**) and 0.005 (**b**) by the t-SNE algorithm using the bootstrapped volumes and their corresponding feature maps. Colors in the left panels indicate the ground truth of 3D volume data points. **c,** Free energy profiles along the MEP calculated by the string method in the 2D energy landscapes of the simulated NLRP3 datasets at SNRs of 0.05 (top), 0.01 (middle) and 0.005 (bottom). **d,** Linear dimensionality reduction of bootstrapped 3D volumes from the simulated datasets at three distinct SNRs by the 3D PCA method. **e-g,** Dimensionality reduction of bootstrapped 3D volumes from the simulated dataset at SNR of 0.01 by the multidimensional scaling (MDS)[49] (panel **e**), isometric mapping (Isomap)[50] (panel **f**) and locally linear embedding (LLE)[51] algorithms (panel **g**). Colors of data points indicate the ground truth of their corresponding 3D volumes. **h,** Comparison of local resolution assessment of AlphaCryo4D-classified NLRP3 reconstructions of conformers 2 (upper inset) and 19 (lower inset) from the simulated datasets of three distinct SNRs. The local resolutions were computed by ResMap[103]. Conformers 2 and 19 have been missed 10 and 8 times in 18 control tests using several other methods (see Extended Data Fig. 3c-h), which are the most and second-most frequently missed conformers, respectively. The 3D classification precision (P) is labelled above each density map. The color bar of local resolution is shown in the lower insert. **i,** Energy landscape calculated by 3D volumes with preprocessing of 5 Å low-pass filtering instead of standardization. The gap between the majority and minority data points was widened at different SNRs.

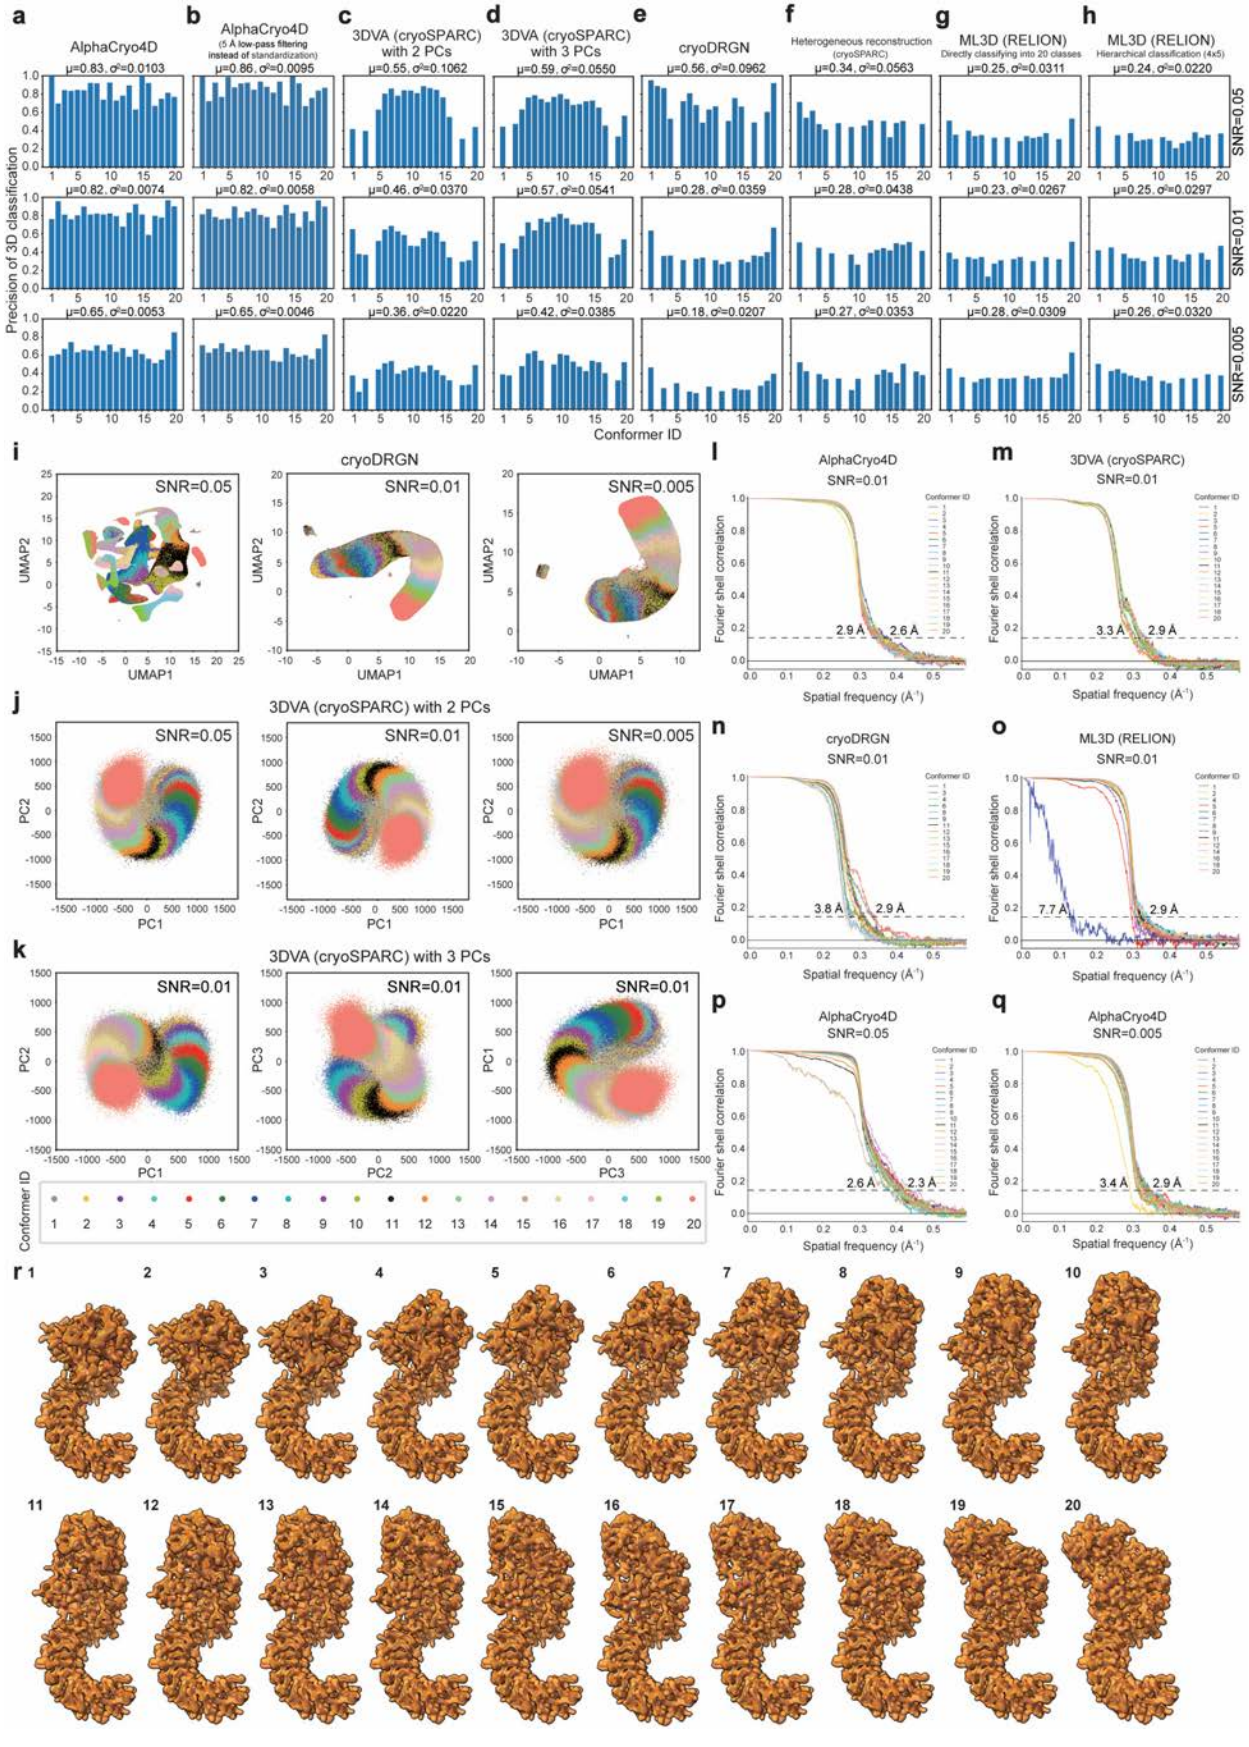
a
AlphaCryo4D
b
AlphaCryo4D
(5 Å low-pass filtering instead of standardization)
c
3DVA (cryoSPARC) with 2 PCs
d
3DVA (cryoSPARC) with 3 PCs
e
cryoDRGN
f
Heterogeneous reconstruction (cryoSPARC)
g
ML3D (RELION)
Directly classifying into 20 classes
h
ML3D (RELION)
Hierarchical classification (4x5)
Precision of 3D classification
SNR=0.05
SNR=0.01
SNR=0.005
Conformer ID
i
cryoDRGN
UMAP1
UMAP2
j
3DVA (cryoSPARC) with 2 PCs
PC1
PC2
k
3DVA (cryoSPARC) with 3 PCs
PC3
Conformer ID
l
AlphaCryo4D
SNR=0.01
m
3DVA (cryoSPARC)
SNR=0.01
n
cryoDRGN
SNR=0.01
o
ML3D (RELION)
SNR=0.01
p
AlphaCryo4D
SNR=0.05
q
AlphaCryo4D
SNR=0.005
Fourier shell correlation
Spatial frequency (Å⁻¹)
r

Extended Data Fig. 3. Performance comparison of AlphaCryo4D with alternative methods using the simulated heterogeneous NLRP3 datasets of different SNRs. a-h, 3D classification precision of the simulated datasets by AlphaCryo4D (**a**), AlphaCryo4D preprocessed by 5 Å low-pass filtering instead of standardization (**b**), 3DVA with two principal components (PCs) in cryoSPARC[11] (**c**), 3DVA with three principal components (PCs) in cryoSPARC[11] (**d**), cryoDRGN[11] (**e**), heterogeneous reconstruction in cryoSPARC (**f**) and maximum-likelihood classification in RELION (**g**, **h**). The results of SNRs of 0.05 (the first row), 0.01 (the second row) and 0.005 (the third row) are shown on three rows for side-by-side comparison. On the top of each panel, the symbols of μ and σ^2 denote the mean and variance of precision, respectively, with the values of missing classes treated as zeros. In the maximum-likelihood classification of RELION, both direct and hierarchical strategies are compared in the study. **i-k,** Visualization of 3D classification by cryoDRGN in autoencoder-learned feature space (**i**), and by 3DVA with 2 PCs (**j**) and 3 PCs (**k**). **l-o,** The gold-standard FSC plots of the 20 maps resulting from 3D classification by AlphaCryo4D (**l**), of the 18 maps resulting from the 3D classification by 3DVA in cryoSPARC (**m**), of the 15 maps resulting from cryoDRGN (**n**) and of the 14 maps resulting from the maximum-likelihood 3D classification in RELION (**o**) on the simulated data of 0.01 SNR. They correspond to the precision results presented in the second row of panels (**a**), (**c**), (**e**) and (**g**), respectively. **p** and **q,** Additional gold-standard FSC plots of the refined density maps resulting from AlphaCryo4D on the simulated datasets of SNRs of 0.05 and 0.005. **r,** The 20 maps of distinct NLRP3 conformers resulting from the 3D classification by AlphaCryo4D. All maps are shown in transparent surface representations superimposed with their corresponding atomic models of the ground truth in cartoon representations, which are fitted to the maps as rigid bodies without further atomic modelling. The conformer ID numbers are marked on the upper left of each map panel.

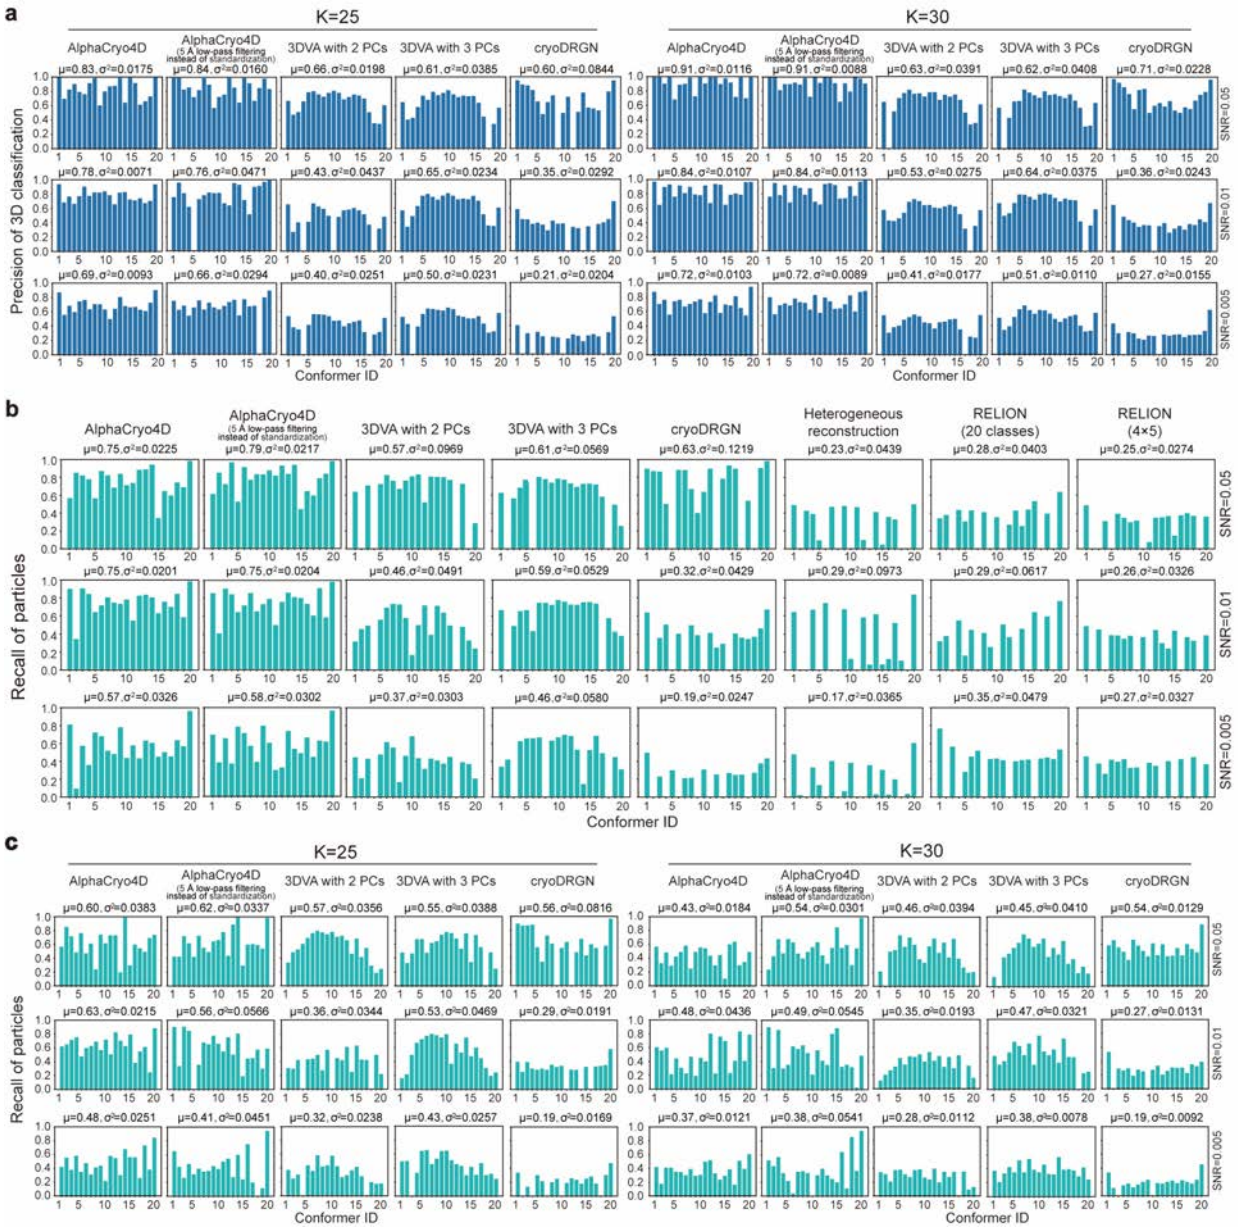

Extended Data Fig. 4. Performance evaluation of 3D classification with alternative parameters and comparison of the recall of 3D classification using the simulated datasets of three distinct SNRs. **a,** 3D classification precision of AlphaCryo4D preprocessed by standardization or 5 Å low-pass filtering, 3DVA using two PCs, 3DVA using three PCs and cryoDRGN when the class number K was set to 25 or 30, which are more than the class number 20 of ground truths. **b,** 3D classification recall of AlphaCryo4D preprocessed by standardization or 5 Å low-pass filtering, 3DVA using two PCs, 3DVA using three PCs, cryoDRGN, heterogeneous reconstruction of cryoSPARC and maximum-likelihood classification of RELION at SNRs of 0.05 (the first row), 0.01 (the second row) and 0.005 (the third row). Particles recall

of each state was calculated as the ratio of true positive particles in all ground truth particles. In all methods, the class number was set to 20. **c,** 3D classification recall of AlphaCryo4D preprocessed by standardization or 5 Å low-pass filtering, 3DVA using two PCs, 3DVA using three PCs and cryoDRGN when the class number was set to 25 or 30 (more than the class number 20 of ground truths).

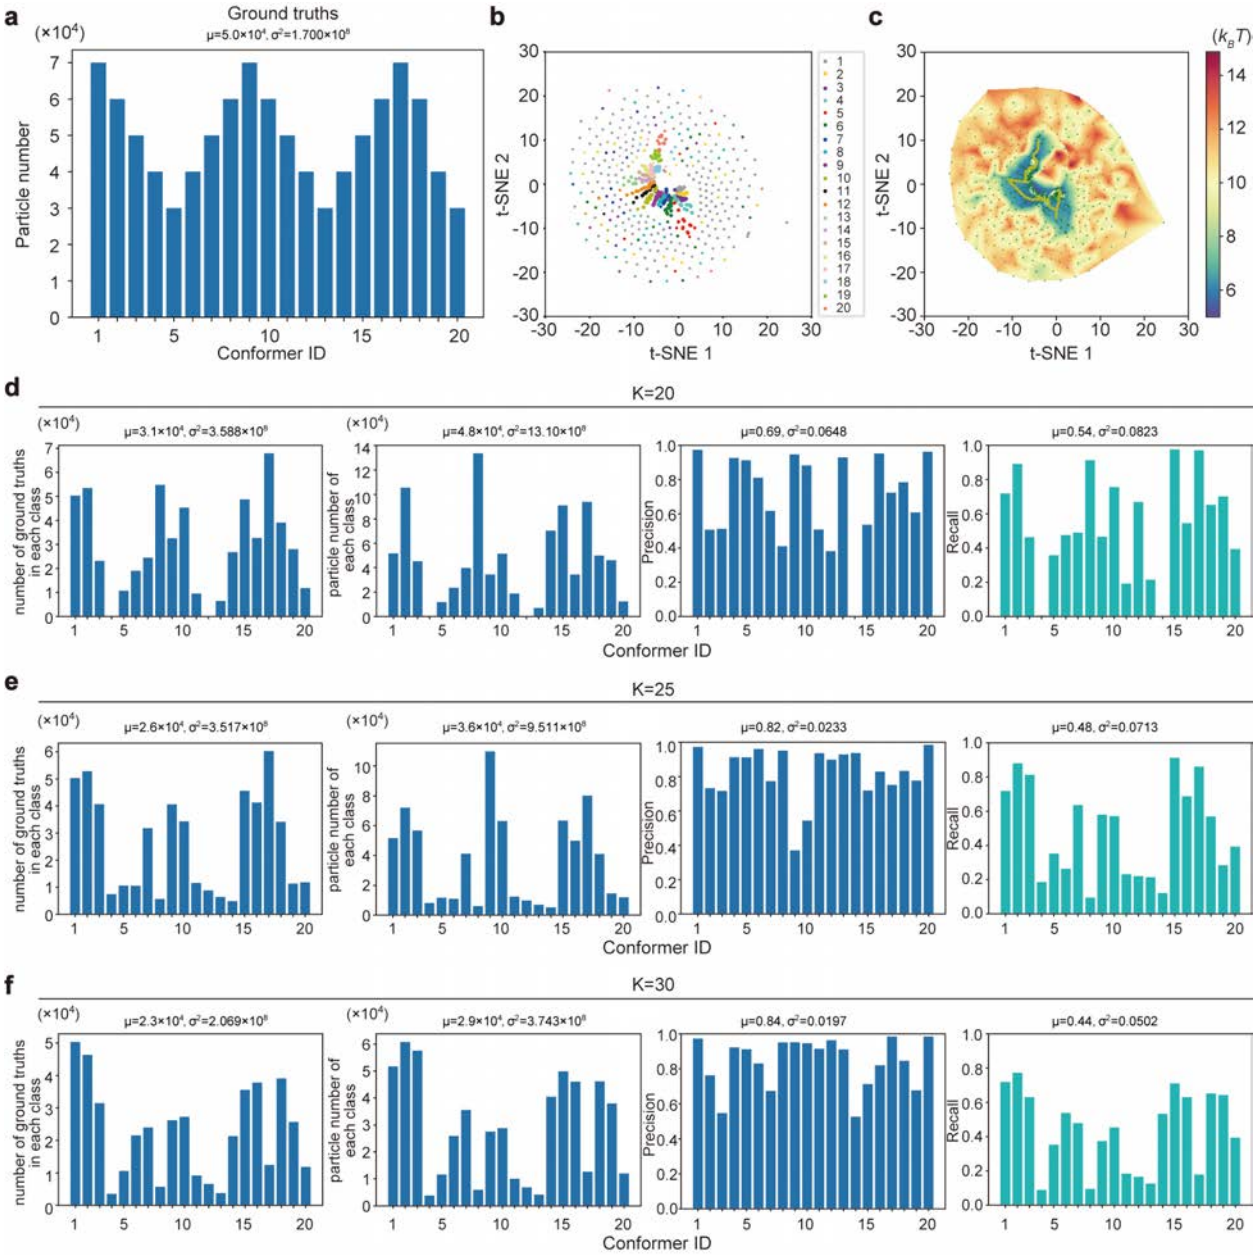

Extended Data Fig. 5. Performance evaluation of 3D classification by AlphaCryo4D on the synthetic dataset with non-uniform distribution of conformational continuum. **a**, The non-uniform distribution of 20 conformers of another simulated NLRP3 dataset at SNR of 0.01 that was used to examine the performance of AlphaCryo4D. This dataset was utilized to calculate the energy landscape shown in panels (**b, c**) and its corresponding outcome of 3D classification by AlphaCryo4D shown in panels (**d-f**), which tests the robustness of AlphaCryo4D in the case of non-uniform distributions of the underlying conformational states. **b**, Dimensionality reduction of bootstrapped 3D volumes and their corresponding feature maps from the simulated dataset by

t-SNE. Colors of data points indicate the ground truth of their corresponding 3D volumes. **c**, Reconstruction of the energy landscape of the simulated NLRP3 dataset. **d-f**, The number of correctly classified particles in each class, the particle number of each class, the precision of 3D classification and the corresponding recall of AlphaCryo4D blind tests using this simulated dataset when the class number is set to 20 (**d**), 25 (**e**) and 30 (**f**).

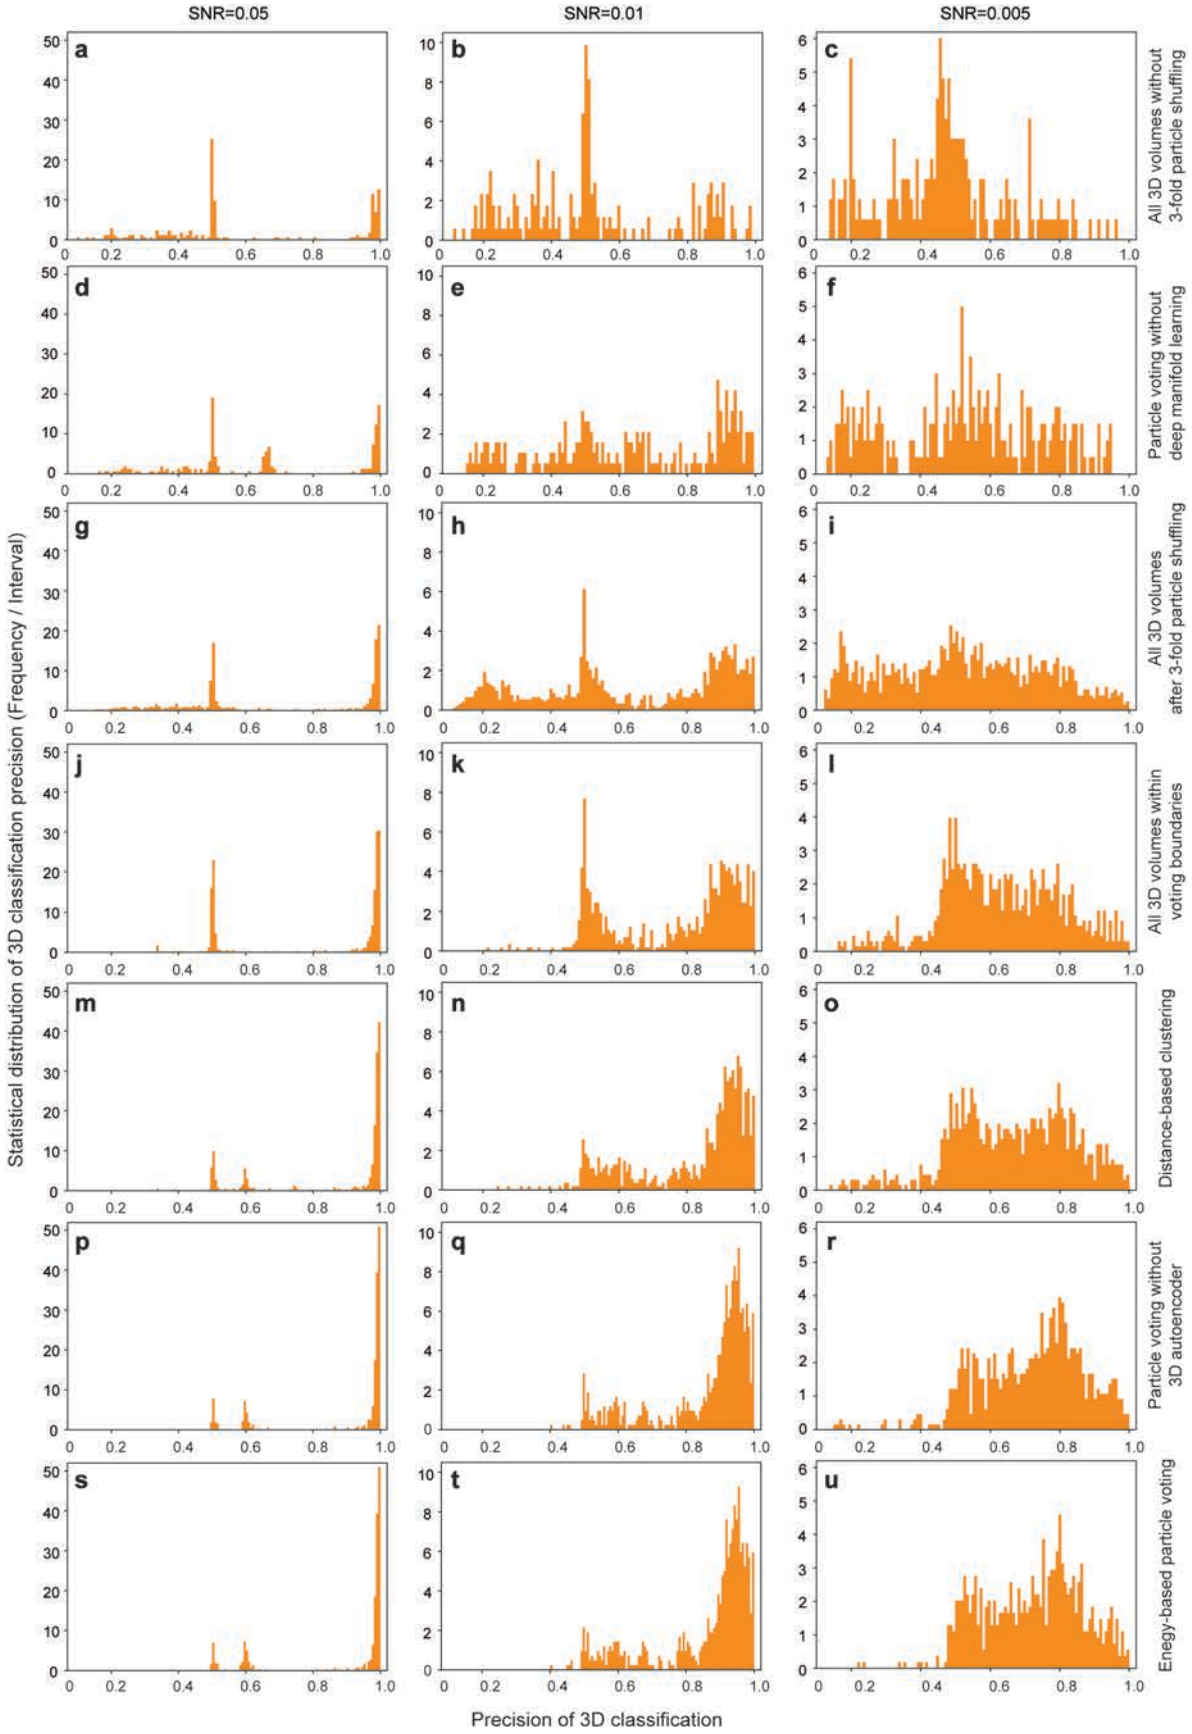
SNR=0.05
SNR=0.01
SNR=0.005
a
b
c
d
e
f
g
h
i
j
k
l
m
n
o
p
q
r
s
t
u
All 3D volumes without 3-fold particle shuffling
Particle voting without deep manifold learning
All 3D volumes after 3-fold particle shuffling
All 3D volumes within voting boundaries
Distance-based clustering
Particle voting without 3D autoencoder
Enegy-based particle voting
Statistical distribution of 3D classification precision (Frequency / Interval)
Precision of 3D classification

Extended Data Fig. 6. Mechanistic characterizations of the improvement of 3D classification accuracy by AlphaCryo4D using the simulated NLRP3 datasets of three typical SNRs. Left, middle and right vertical columns show the precision analysis on the simulated datasets with SNRs of 0.05, 0.01 and 0.005, respectively. In total, there are 21 conditional controls under six different algorithmic design variations being analyzed in panel (**a**)-(**r**). **a-c,** Statistical distribution of 3D classification precision in bootstrapped 3D volumes without particle shuffling. **d-f,** Statistical distribution of 3D classification precision by implementing particle voting directly on the bootstrapped 3D volumes without using deep manifold learning. The procedure of 3-fold particle shuffling and volume bootstrapping is identical to AlphaCryo4D. **g-i,** Distribution of 3D classification precision in bootstrapped 3D volumes after 3-fold particle shuffling in the intermediate step of AlphaCryo4D. **j-l,** Distribution of 3D classification precision in bootstrapped 3D volumes screened by voting boundary on the energy landscape in the intermediate step of AlphaCryo4D. **m-o,** Distribution of 3D classification precision after distance-based 3D clustering in the absence of energy-based particle voting with all prior steps identical to AlphaCryo4D. **p-r,** Distribution of 3D classification precision after particle voting by a modified AlphaCryo4D variation only without using deep residual autoencoder in the manifold embedding step. **s-u,** Distribution of 3D classification precision after energy-based particle voting via a complete AlphaCryo4D procedure.

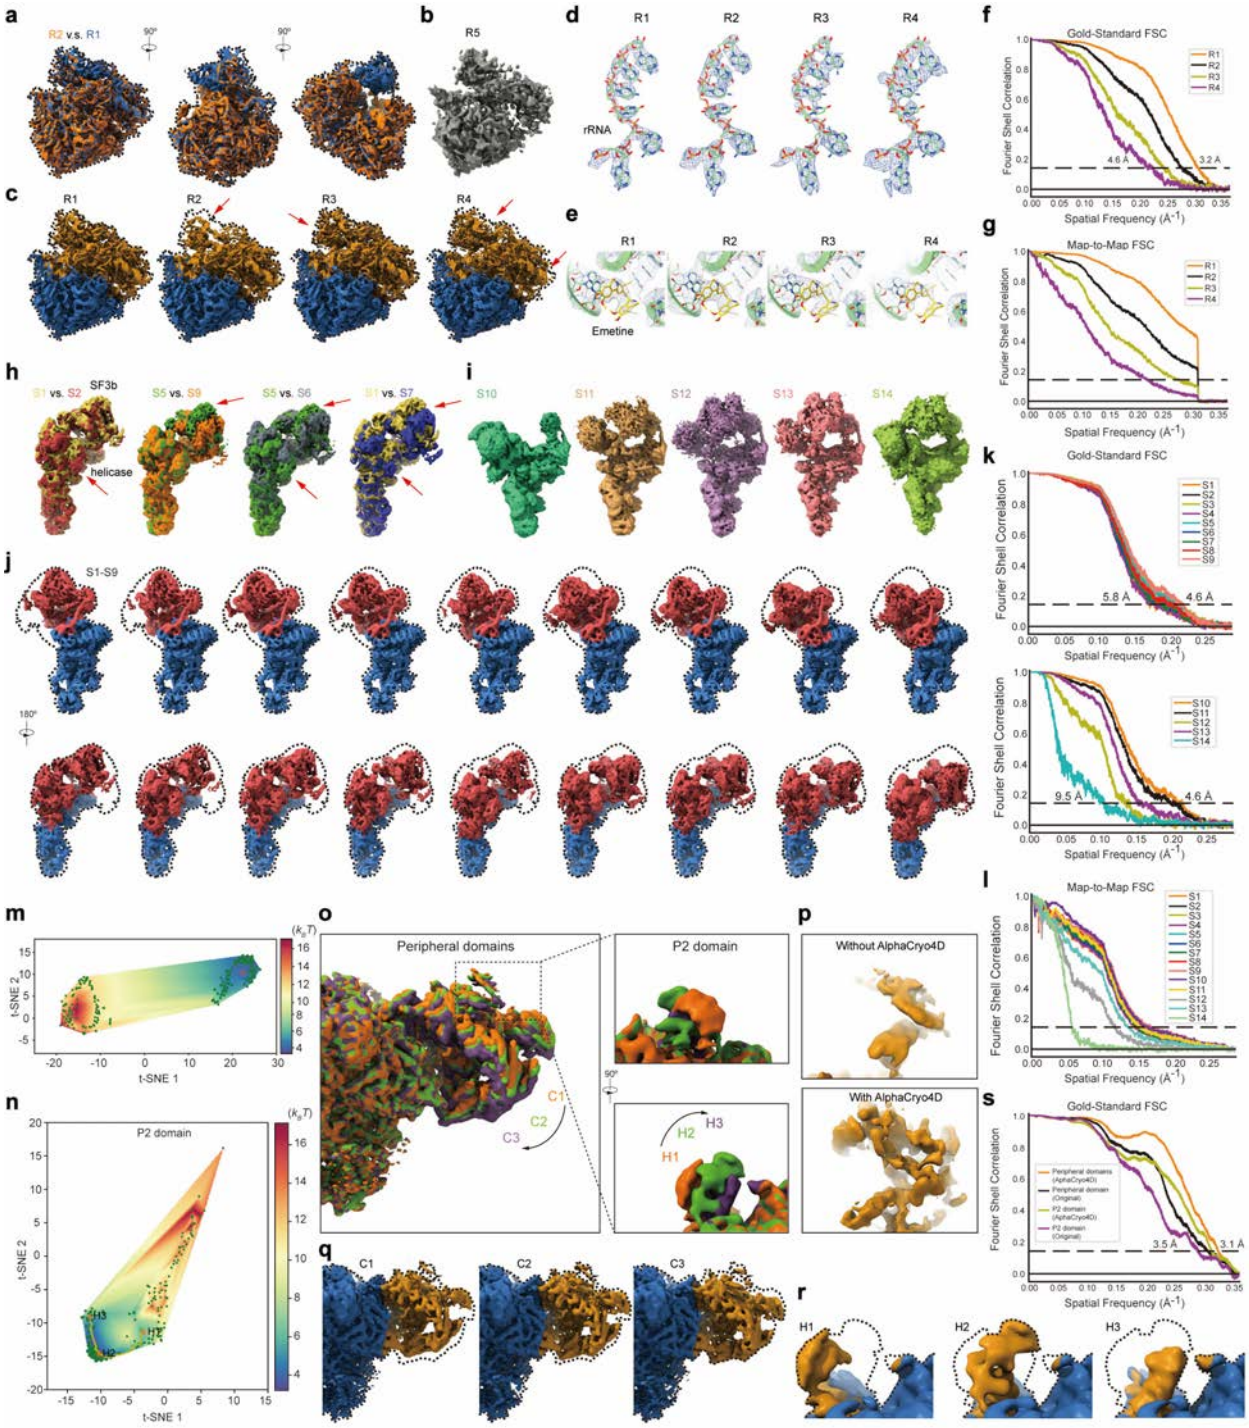

Extended Data Fig. 7. Applications of AlphaCryo4D in analysis of three experimental datasets. **a-g,** Visualizing conformational dynamics of the *Pf*80S ribosome bound to anti-protozoan drug at high resolution by AlphaCryo4D. **a,** Cryo-EM density map of the 40S head missing state R2 superimposed with the major state R1 from three different viewing angles. **b,** Cryo-EM density map of cluster R5 on the energy landscape (Fig. 2a). The particle number of

this cluster limited the resolution of its final reconstructed density map. **c,** Comparison of the density maps of clusters R1-R4. The red arrows point to the conformational difference of states R2-R4 relative to the state R1. **d,** Comparison of local cryo-EM densities of an rRNA strand superimposed with its corresponding atomic model in four conformational states. **e,** Comparison of local cryo-EM densities of the small-molecule drug emetine bound to the ribosome, superimposed with its corresponding atomic model, in four conformational states. **f**, Gold-standard FSC of the R1-R4 states reconstructed by AlphaCryo4D. **g**, Map-to-map FSC of the R1-R4 states with the published cryo-EM structure of the *Pf*80S ribosome (EMD-2660). The cliff of the FSC curve was caused by the low-pass filtering of the previously released cryo-EM maps. **h-l,** Visualizing continuous motions of pre-catalytic spliceosome by AlphaCryo4D. **h,** Different local motion modes of the SF3b and helicase subcomplexes. The red arrows pointed the regions where the conformational changes are observed by comparing the cryo-EM densities of different states. These four comparisons exhibited the local movement of helicase (S1 vs. S2), the local movement of SF3b (S5 vs. S9), the relative motions of the SF3b and helicase subcomplexes in the opposite direction (S5 vs. S6) and the concerted motions of the SF3b and helicase subcomplexes in the same direction (S1 vs. S7), which showed that the continuous motion the SF3b or helicase subcomplexes is independent. **i,** Cryo-EM density maps of clusters S10-S14. Discrete heterogeneity is observed in these states. **j,** Cryo-EM density maps of clusters S1-S9 from other two viewing angles. The dynamic region of the density map is colored red and the rigid region is colored blue. Continuous conformational motion was observed among these states. **k,** Gold-standard FSC of the S1-S14 states classified by AlphaCryo4D. **l,** Map-to-map FSC of the S1-S14 states over the published B4 cryo-EM density map of the pre-catalytic spliceosome (EMD-3685). **m-s**, Improvement of local resolution of the peripheral domains of the RyR1 by AlphaCryo4D. **m,** Energy landscape estimated with the alignment of the peripheral domains. This energy landscape is aimed to improve the quality of the final density map. **n,** Focused energy landscape of the P2 domain of the RyR1. Clusters H1, H2 and H3 were uncovered on this energy landscape. The small P2/RY3&4 domain abutting the helical C-Sol domain via flexible links has low resolution in the refined map shown in Fig. 2k. To improve these features, we computed the zoomed-in energy landscape by focusing on the P2 domain, which also exhibited three conformational clusters, from which we found the P2 domain oscillating relative to the helical C-Sol domain. Further refinement of the AlphaCryo4D-

classified particles improved the clarity of the secondary structure in the density map shown in (**r**), but it did not reach the atomic level likely due to the very small size of this domain limiting the alignment accuracy. **o,** Superposition of the cryo-EM density maps of clusters C1, C2 and C3 from the energy landscape of the peripheral and P2 domains. The dashed line indicates the outline of all density maps of P2 domain superimposed together. **p,** Comparison of the density map of the P2 domain after AlphaCryo4D with the original state 1. **q,** Side-by-side comparison of the cryo-EM density maps of clusters C1, C2 and C3 from the energy landscape of the peripheral domain. The dashed line indicates the outline of all density maps superimposed together. The dynamic region of the density map is colored orange, and the stable region is colored blue. **r,** Local cryo-EM density maps of clusters H1, H2 and H3 on the energy landscape of the P2 domain. **s,** Gold-standard FSC of the peripheral and P2 domains with and without AlphaCryo4D.

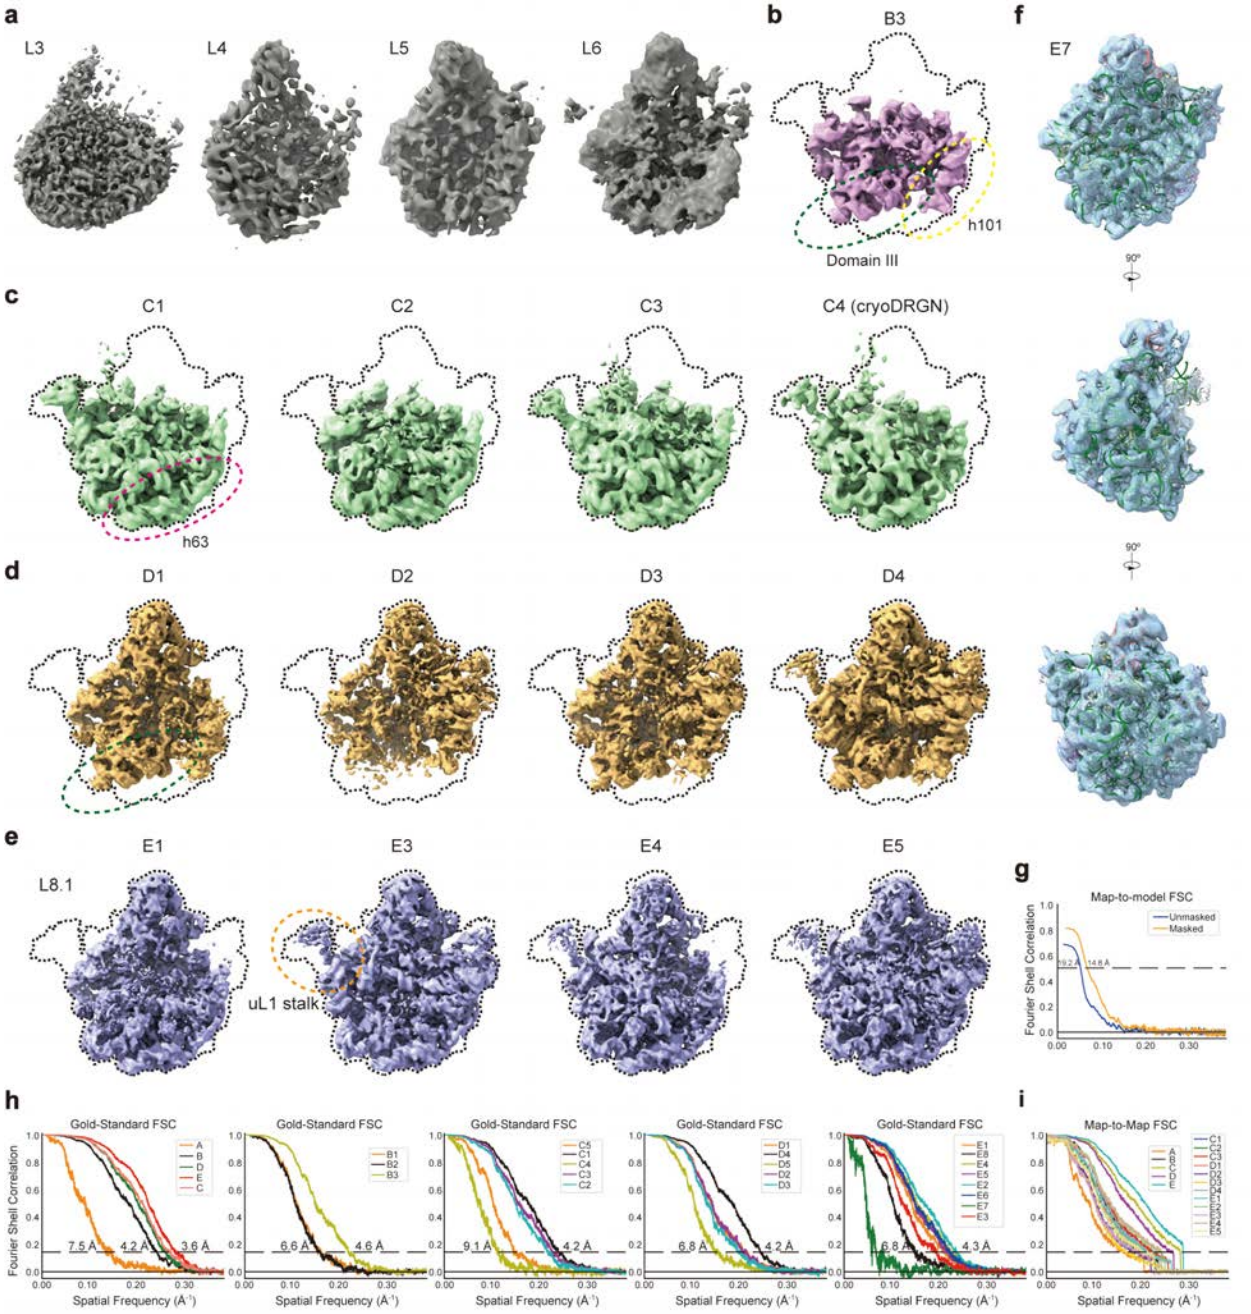

Extended Data Fig. 8. Reconstructions of bacterial ribosomal assembly intermediates by AlphaCryo4D. a, Cryo-EM density maps of clusters L3-L6. The classes contained particles of likely poor quality. **b,** Reconstruction of the cluster B3 from the state B energy landscape (Fig. 3d). The cluster B3 revealed the main conformation identical to state B. **c,** Reconstructions of the state C1, C2, C3 and C4 from the state C energy landscape (Fig. 3f). The state C4 contained the density of rRNA helix 68, which was previously found by cryoDRGN. **d,** Reconstructions of the state D1, D2, D3 and D4 from the state D energy landscape (Fig. 3h). **e,** Reconstructions of the

state E1, E3, E4 and E5 from the state E energy landscape (Fig. 3j). **f,** The cryo-EM map of the state E7 in transparent surface representation superimposed with corresponding atomic model in cartoon representation from three different viewing angles. **g,** Model-to-map FSC plots of the state E7 reconstruction with and without applying a global mask. **h,** Gold-standard FSC of the 5 major states (A, B, C, D and E) and 21 minor states (B1-B3, C1-C5, D1-D5 and E1-E8) reconstructed by AlphaCryo4D. **i,** Map-to-map FSC of the existing states over the published cryo-EM structures of the states A (EMD-8434), B (EMD-8440), C (EMD-8441), D (EMD-8445), E (EMD-8450), C1 (EMD-8442), C2 (EMD-8443), C3 (EMD-8444), D1 (EMD-8446), D2 (EMD-8447), D3 (EMD-8448), D4 (EMD-8449), E1 (EMD-8451), E2 (EMD-8452), E3 (EMD-8453), E4 (EMD-8455) and E5 (EMD-8456) of the bacterial ribosomal assembly intermediates.

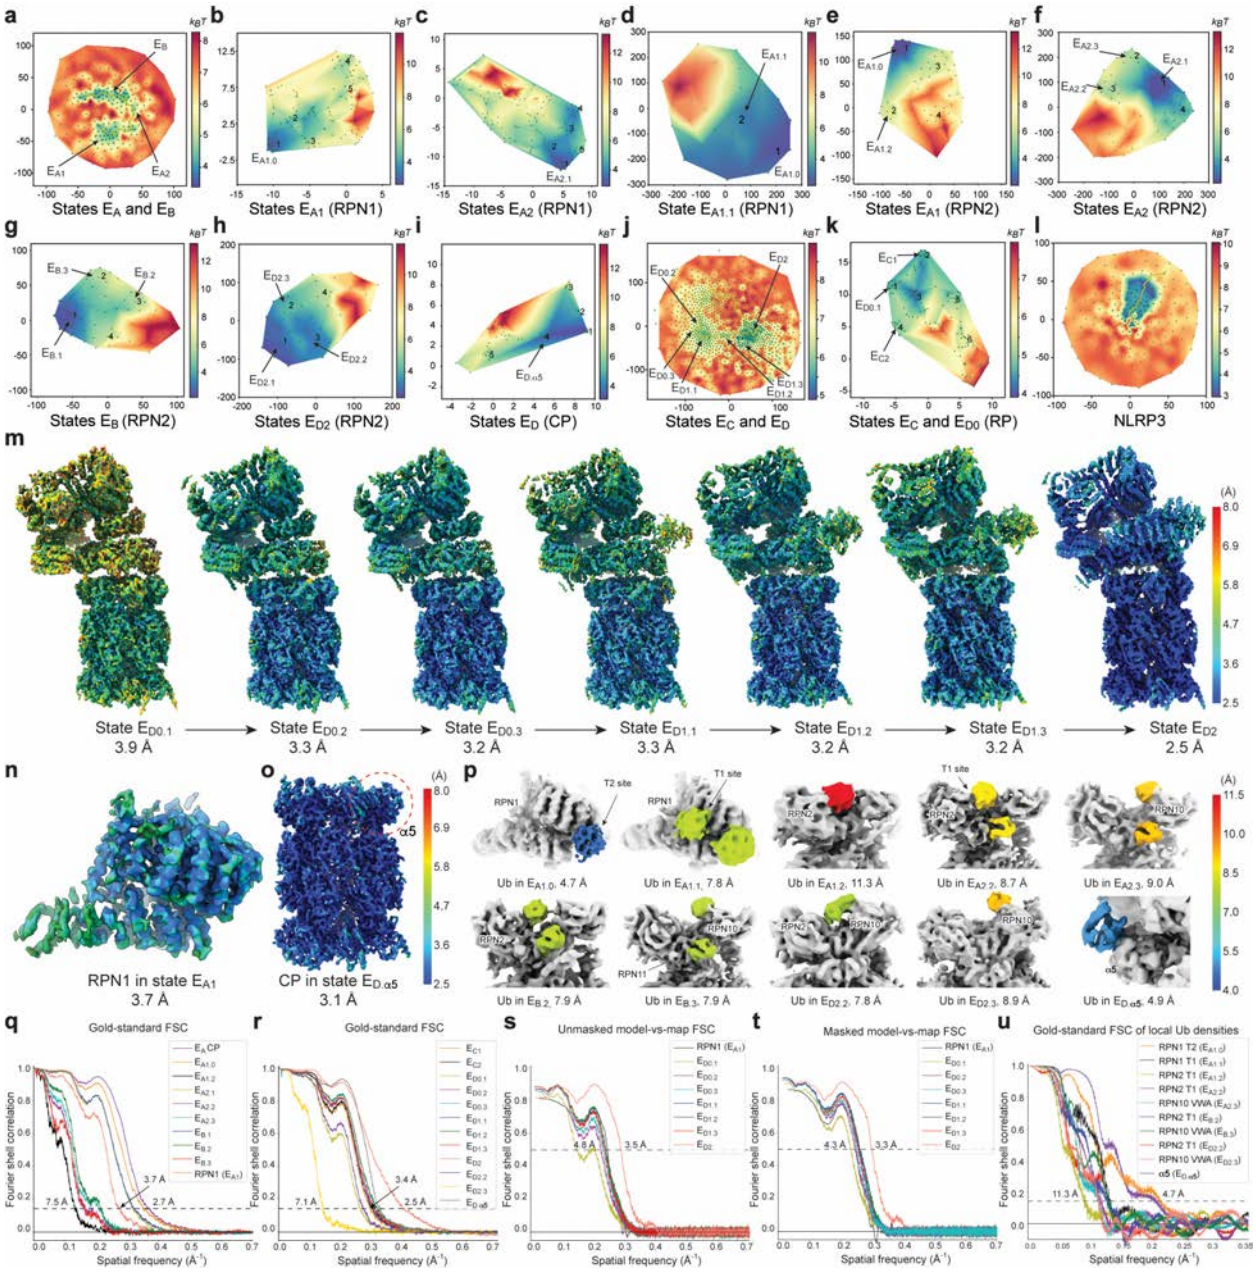

Extended Data Fig. 9. Energy landscapes and cryo-EM reconstructions of novel proteasomal states determined by AlphaCryo4D. a-j, The zoomed-in energy landscape used for focused 3D classifications of AlphaCryo4D on the experimental cryo-EM datasets for finding new states of the substrate-bound proteasome. The zoomed-in local energy landscapes of the substrate-bound human proteasome by AlphaCryo4D are shown, respectively, for state E_A and E_B in (**a**), for states E_{A1} and E_{A2} with 3D mask focusing on RPN1 in (**b**) and (**c**) to visualize the RPN1 T2 site, in (**d**) to visualize the RPN1 T1 site, for states E_{A1}, E_{A2}, E_B and E_{D2} with 3D masking around RPN2 in (**e**), (**f**), (**g**) and (**h**) to visualize the RPN2 T1 site, for combined states

of E_D with masking around CP (**i**) to visualize the ubiquitin-binding sites on the α5-subunit, for states E_C and E_D in (**j**) and for combined states of E_C and E_{D0} with 3D mask focusing on the RP (**k**) to visualize the ATPase motor intermediates during translocation initiation. **l**, The energy landscape computed from the experimental cryo-EM dataset of the NLRP3-NEK7 complex. **m**, Local resolution measurement of the seven sequential intermediate states during translocation initiation measured by ResMap[103]. **n**, The local resolution measurement of the improved density map of RPN1 in state E_{A1} by ResMap. **o**, The local resolution measurement of the CP in state $E_{D.\alpha5}$ by ResMap. **p**, Putative ubiquitin (Ub) densities in ten distinct states colored by the local gold-standard FSC resolutions of the Ub densities measured in panel (**u**), with the rest of densities shown in light grey. The color bar for local resolutions is shown on the right in panels in **m**-**p**, where panels (**n**) and (**o**) share the same resolution color bar. **q** and **r**, The gold-standard FSC plots of the 20 conformers of the RP-CP subcomplex. **s** and **t**, The FSC curves calculated between the experimental cryo-EM maps and their corresponding atomic models for newly discovered states or improved states without (**s**) and with (**t**) masking in the calculations. **u**, The gold-standard FSC curves calculated between the two half maps with local Ub masks in ten distinct states.

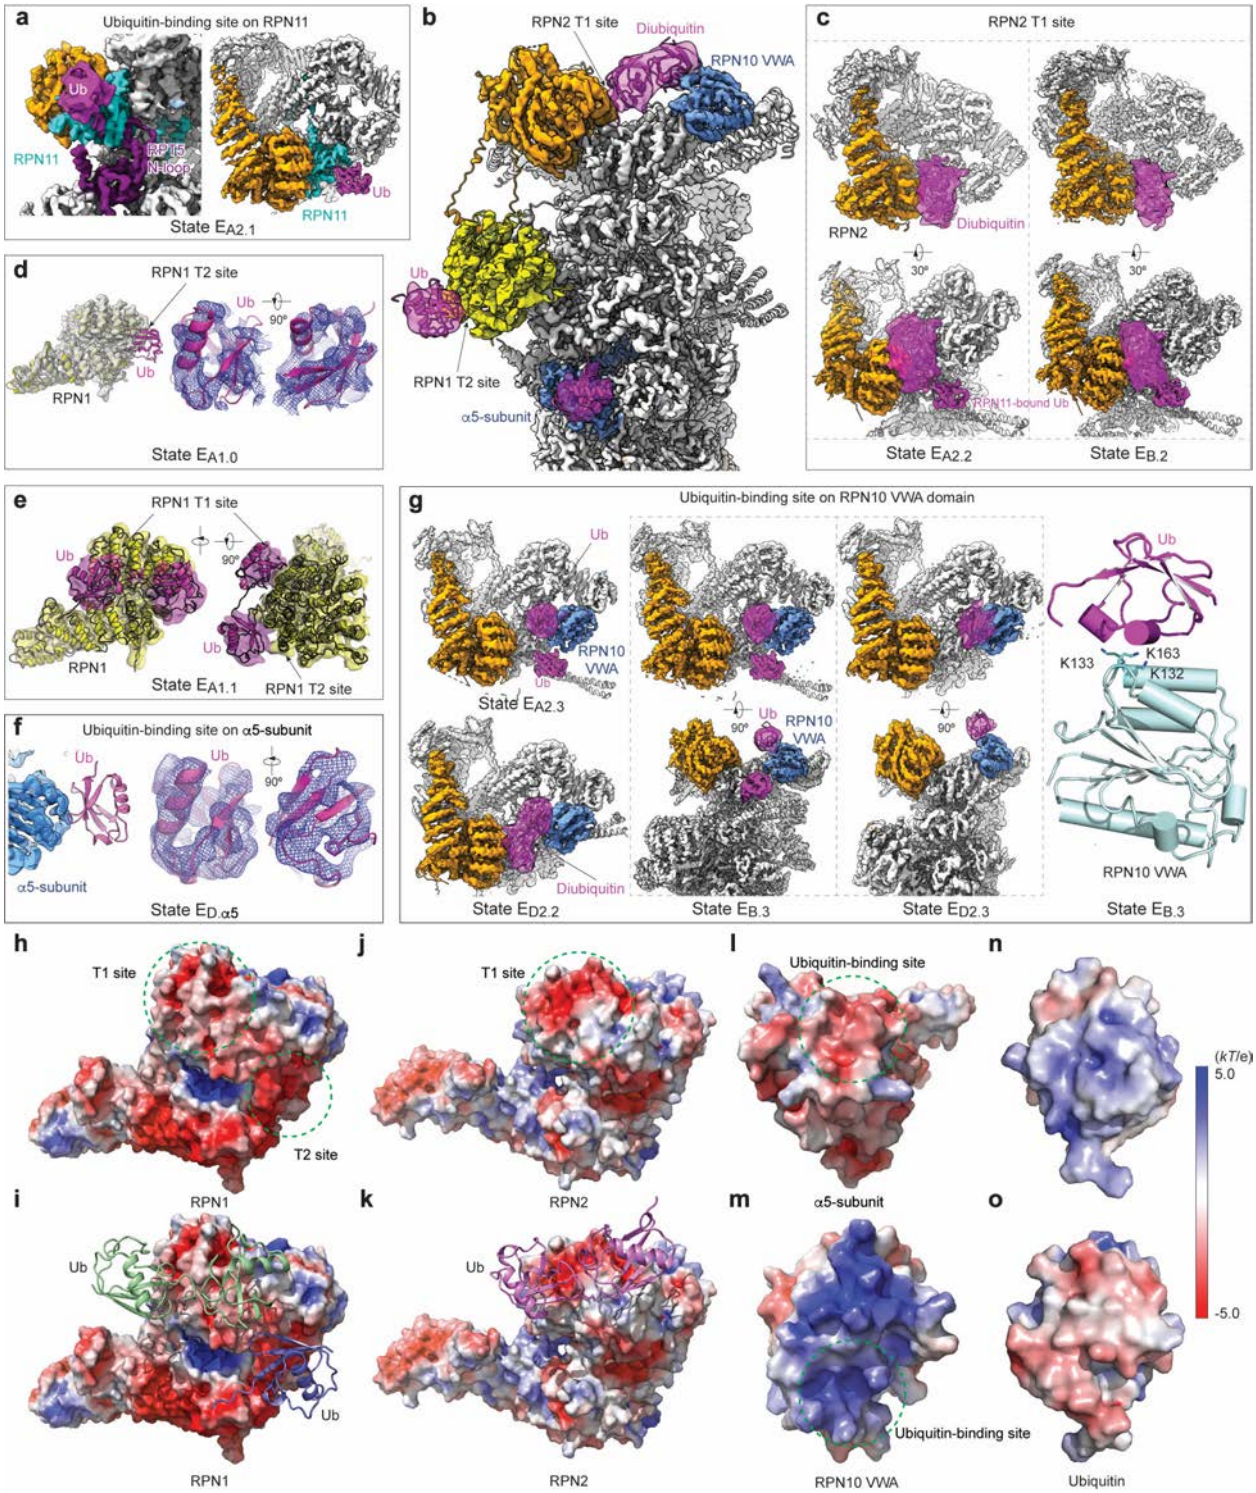

Extended Data Fig. 10. Structural analysis of ubiquitin-proteasome interactions. a, Cryo-EM density of state $E_{A2.1}$ shows the improved resolution localized around RPN11-bound ubiquitin and the RPT5 N-loop as compared to previously determined state E_{A2}. Ub, ubiquitin. **b**, High-resolution cryo-EM density of state E_{D2} superimposed with Ub densities in state $E_{D2.2}$

highlighted as magenta, showing that a diubiquitin chain in contact with both the RPN2 T1 site and the RPN10 VWA domain. The atomic model of state $E_{D2.2}$ in cartoon representation is superimposed with the density map. **c**, Closeup view of the RPN2-bound diubiquitin density in states $E_{A2.2}$ (left column) and $E_{B.2}$ (right column) in two different perspectives, with the lower row showing the coexistence of the RPN11-bound Ub and the RPN2-bound diubiquitin. **d**, The refined 3.7-Å RPN1 density map from state $E_{A1.0}$ in transparent surface representation superimposed with its atomic model with Ub bound near the RPN1 T2 site. The right insets show the local cryo-EM density of the RPN1-bound Ub in state $E_{A1.0}$ superimposed with its atomic model in cartoon representation. **e**, Cryo-EM density of state $E_{A1.1}$ in transparent surface representation superimposed with its pseudo-atomic model of cartoon representation, with one Ub moiety bound to the RPN1 T1 site and another Ub near the T2 site. The two Ub moieties appear to be compatible with a Lys63-linked diubiquitin model in an extended conformation, which is supported by the low-level density connecting the two Ub densities. **f**, Closeup view of the α5-subunit bound with Ub in the refined 3.1-Å density map of state $E_{D.\alpha5}$. The cryo-EM density of the α5-subunit is shown in transparent surface representation superimposed with its atomic model. The right insets show the local cryo-EM density of the α5-bound Ub of state $E_{D.\alpha5}$ in blue mesh representation low-pass filtered at 4.9 Å and superimposed with its atomic model in cartoon representation. **g**, Closeup views of the RPN10-bound Ub densities in states $E_{A2.3}$, $E_{B.3}$, and $E_{D2.3}$. Right panel shows the pseudo-atomic model of Ub interaction with the RPN10 VWA domain. The residues of RPN10 in contact with ubiquitin are labeled and shown as stick representation. Note that the Ub-binding mode of the RPN10 VWA domain here is different from that previously observed in the crystal structure of ubiquitylated yeast Rpn10 that also suggested Rpn10 VWA domain as a ubiquitin-binding domain[65]. In the yeast Rpn10, Ub binds to the Rpn10 surface area interfacing Rpn9 and sterically occluded Rpn10 incorporation into the proteasome[65]. **h-m**, The electrostatic surfaces of the RPN1, RPN2, RPN10 and α5 subunits show the charge complementarity of the RPN1 T1 and T2 (panels **h**, **i**), RPN2 T1 (panels **j**, **k**) and α5-binding (panel **l**) sites of Ub interactions are all acidic. By contrast, the RPN10 VWA site is basic (panel **m**), suggesting that it may bind the acidic side of Ub. **m** and **n**, The electrostatic surface of Ub shows one side that is basic (panel **n**) and the other side acidic (panel **o**).

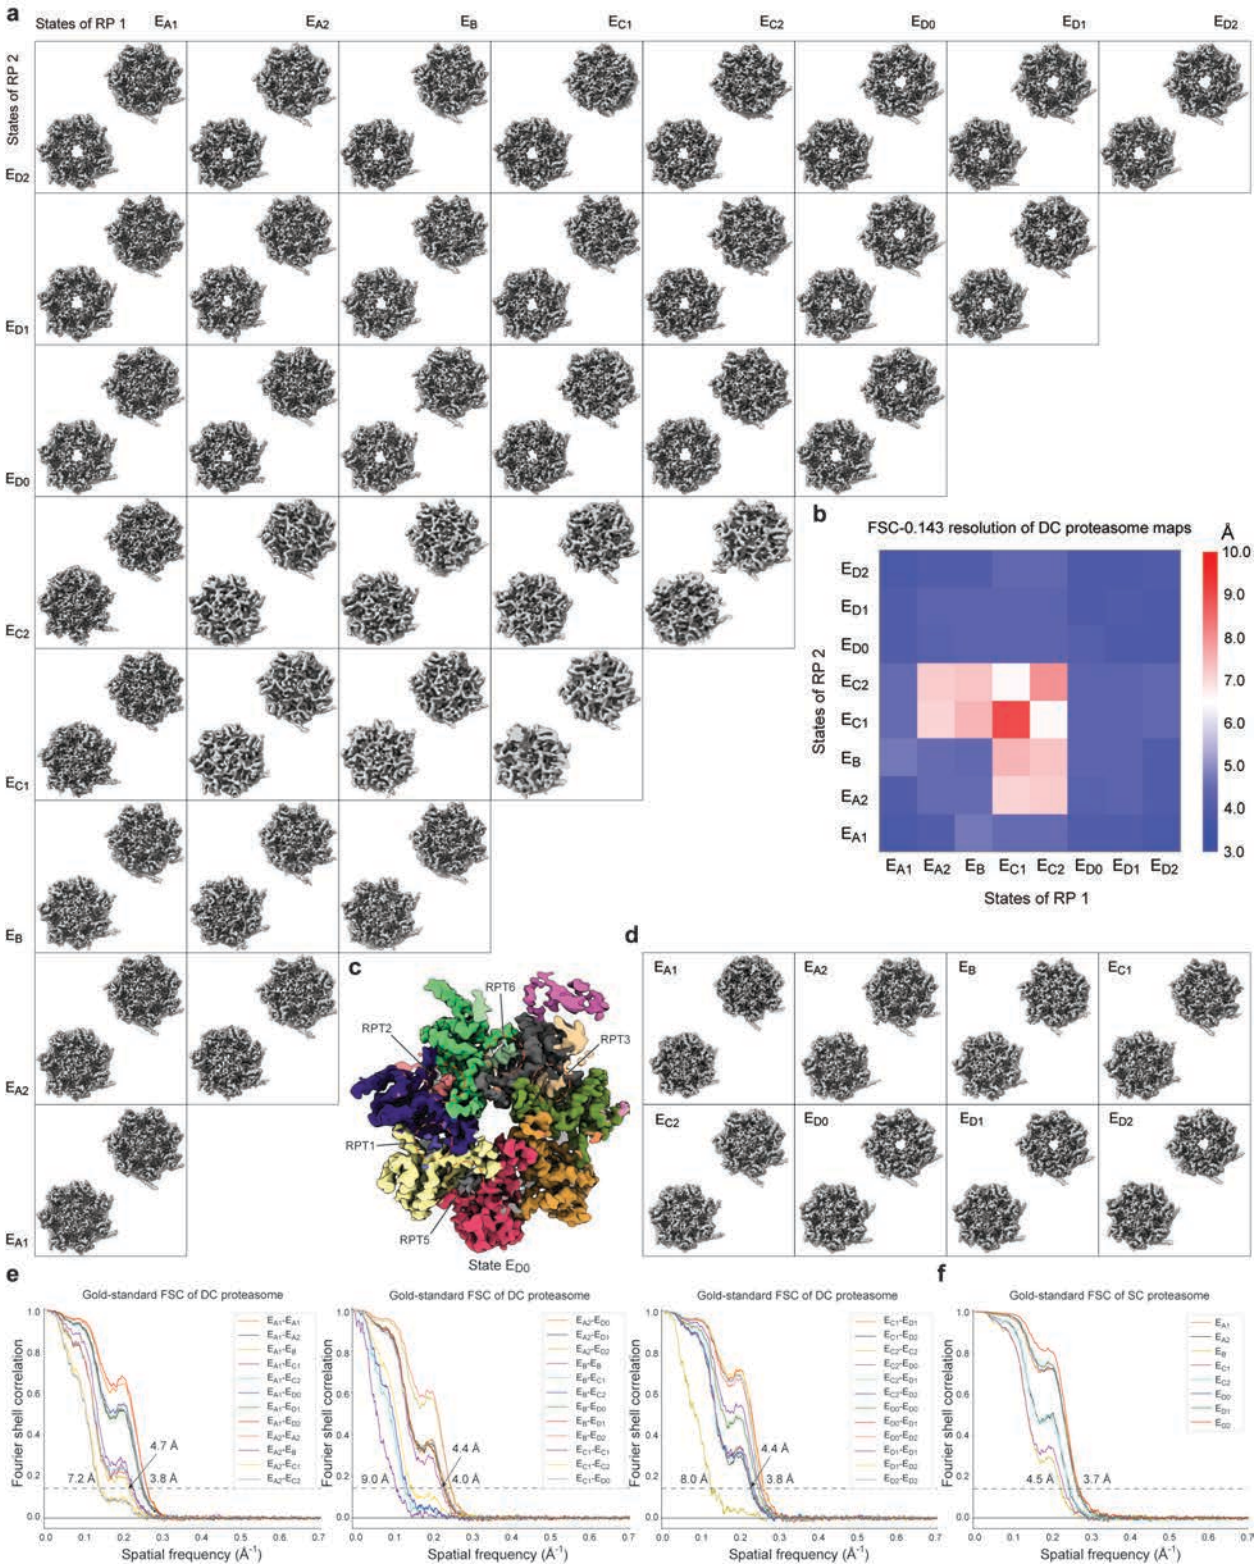

Extended Data Fig. 11. Cryo-EM analysis of the doubly capped (DC) and singly capped (SC) proteasomes. a, The CP gate states of the 36 density maps of distinct conformational states

of the DC proteasome. All density maps shown are low-pass filtered to their measured resolutions by the gold-standard FSC-0.143 cutoff without amplitude correction of B-factors. **b**, The FSC-0.143 resolution matrix of the 36 DC proteasomal states. **c**, The RP-CP interface of state E_{D0}, showing five C-terminal tails inserted into the inter-subunit surface pockets of α-ring. **d**, The CP gate states of the 8 density maps of distinct conformational states of the SC proteasome, with the upper right image in each panel corresponding to the RP-proximal side of the CP gate, and the lower left image to the RP-distal side of the CP gate. **e**, The gold-standard FSC plots of the 36 conformers of the DC proteasome. **f**, The gold-standard FSC plots of the 8 conformers of the SC proteasome.

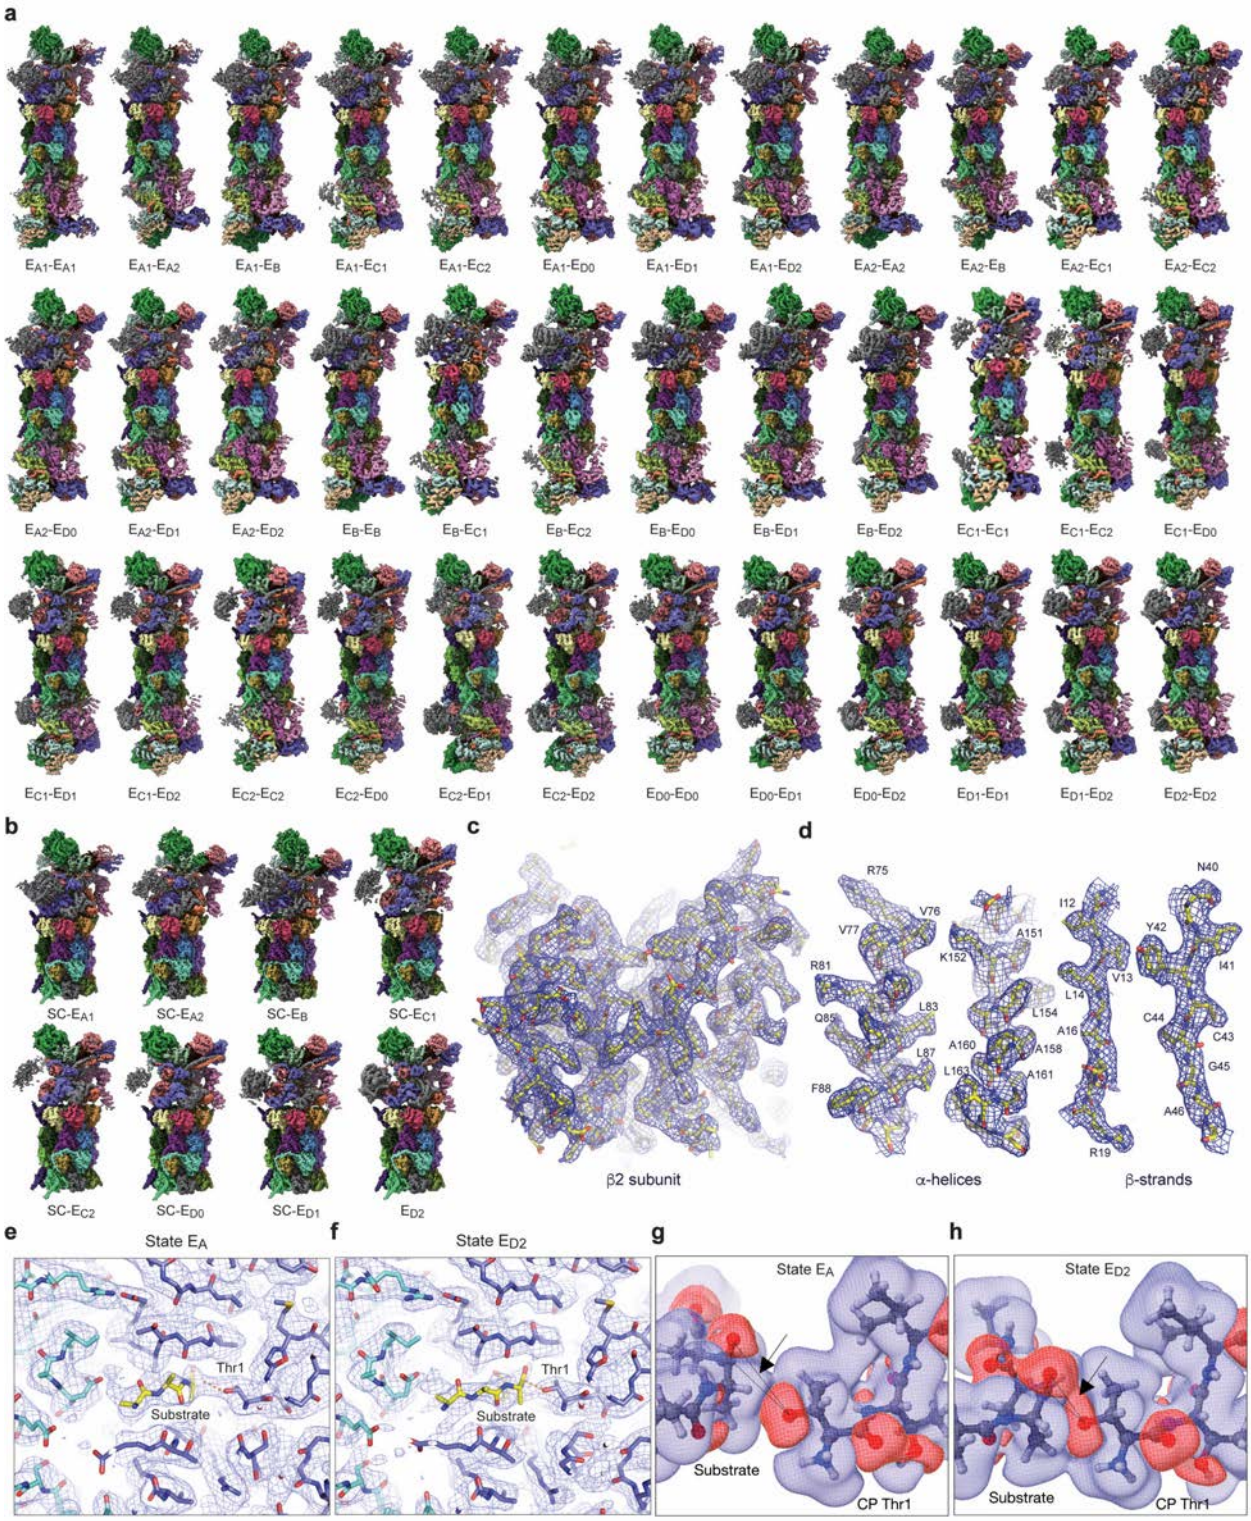

Extended Data Fig. 12. Cryo-EM reconstructions of the SC and DC proteasomes and atomic-level analysis of substrate interactions with the CP active site in distinct states. a, The density maps of 36 distinct conformers of the DC proteasome classified by AlphaCryo4D. **b**,

The density maps of 8 distinct conformers of the SC proteasome. All maps are filtered to their respective gold-standard FSC-0.143 resolutions without amplitude correction of B-factors and are differentially colored by their subunits in ChimeraX[102]. **c**, Closeup view of the cryo-EM density of the β2-subunit in state E_{D2} shown in blue mesh representation superimposed with the atomic model. **d**, Cryo-EM densities of the typical secondary structure elements in the β2-subunit of state E_{D2} shown as blue mesh representation superimposed with their corresponding atomic models. The residues are labelled. The densities are shown at 4σ level and exhibit structural features consistent with 2.5-Å resolution. **e**, The atomic model of substrate-bound catalytic active site of the β2-subunit superimposed with the cryo-EM density map of state E_A at 2.7 Å. **f**, The atomic model of substrate-bound catalytic active site of the β2-subunit superimposed with the cryo-EM density map of state E_{D2} at 2.5 Å. **g** and **h,** 3D Charge density difference map showing the interactions between the substrate polypeptide and the residue Thr1 in the β2-subunit of the CP. The red color labels an iso-surface of charge increase while the blue color labels charge decrease, both at a level of 0.04 e/Å^2. The charge increase region at the Thr1-Oγ atom is elongated sideways, and the substrate position in the E_{D2} state provides both a proximity and orientation alignment for the pair to interact. The dashed line triangle shows the plane in which the detailed charge difference contours are plotted as shown in Fig. 6c, d.

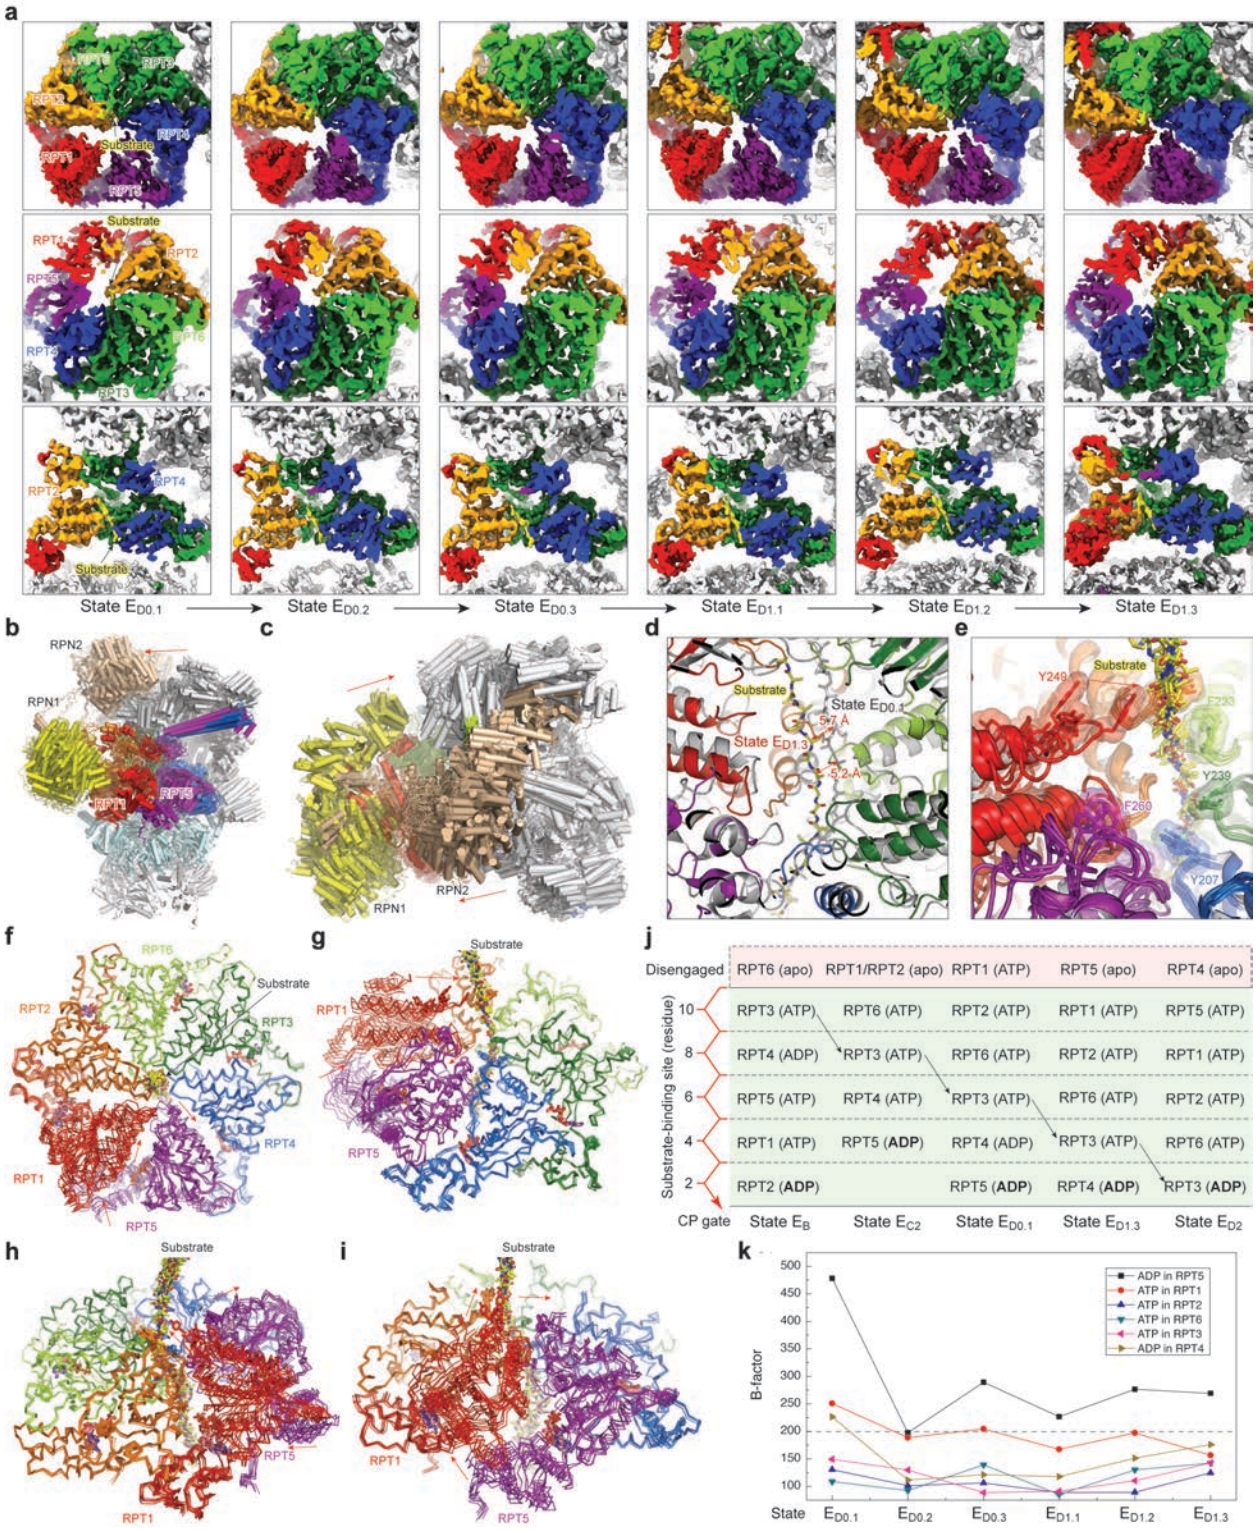

Extended Data Fig. 13. Cryo-EM structures of six sequential intermediate states between E_{C2} and $E_{D1.3}$. a, Cryo-EM density maps of the AAA-ATPase motor in six sequential states from $E_{D0.1}$ to $E_{D1.3}$ from three different viewing angles. The substrate densities are highlighted in

yellow. **b**, Superposition of the seven structures in cartoon representations from a lateral perspective, with all structures aligned together based on their CP structures, showing the RP movements relative to the CP. **c**, Superposition of the six structures in cartoon representations from a top-view perspective, with all structures aligned together based on their CP structures. **d**, Measurement of substrate movement between states $E_{D0.1}$ and $E_{D1.3}$ relative to the CP, suggesting that the substrate is translated ~5-6 Å during the process of ADP release in RPT5 and of RPT1 re-association with the substrate. **e**, The closeup view of the pore-1 loop interaction with the substrate with all six states superimposed after aligned together based on the structures of RPT3, RPT4 and RPT6. **f-i**, Superposition of the six sequential states with their RPT3, RPT4 and RPT6 aligned together. Red arrows show the direction of subunit movements in RPT1 and RPT5. **j**, A diagram illustrating the axial stepping of the substrate-interacting pore-1 loops that is coupled with ATP hydrolysis in the RPT subunits, revised from the previously published results based on the present study[30]. **k**, The average B-factor of the nucleotides in six ATPase subunits fitted to the cryo-EM densities of states $E_{D0.1}$ to $E_{D1.3}$ and computed by Phenix in real-space refinement procedure[99]. The high B-factor of ADP fitted in RPT5 indicates its low occupancy or unstable association and suggests that the ADP in RPT5 undergoes dissociation in this process.

Extended Data Table 1. Hyperparameters of the deep residual networks in the 3D autoencoder. Optimizer: Adam. Epochs: 50. Initial learning rate: 0.01. A factor of 0.1 and patience of 3 means that the learning rate will times 0.1 (factor) if the loss function does not improve in 3 (patience) epochs. Batch size: 32 (for the three simulated datasets), 8 (for the substrate-bound proteasome dataset).

Layer name	Output size (simulation)	Output size (RyR1)	Output size (Proteasome)	Output size (pf80S Ribosome)	Output size (Spliceosome)	Output size (Bacterial Ribosomal Intermediates)	Network structure
Conv1	$200 \times 200 \times 200$	$360 \times 360 \times 360$	$300 \times 300 \times 300$	$360 \times 360 \times 360$	$320 \times 320 \times 320$	$320 \times 320 \times 320$	$5 \times 5 \times 5, 2$
Conv2_x	$100 \times 100 \times 100$	$180 \times 180 \times 180$	$150 \times 150 \times 150$	$180 \times 180 \times 180$	$160 \times 160 \times 160$	$160 \times 160 \times 160$	$3 \times 3 \times 3, 2$, stride 2 $3 \times 3 \times 3, 2$
Conv3_x							$3 \times 3 \times 3, 2$ $3 \times 3 \times 3, 2$
Conv4_x	$50 \times 50 \times 50$	$90 \times 90 \times 90$	$75 \times 75 \times 75$	$90 \times 90 \times 90$	$80 \times 80 \times 80$	$80 \times 80 \times 80$	$3 \times 3 \times 3, 1$, stride 2 $3 \times 3 \times 3, 1$
Conv5_x							$3 \times 3 \times 3, 1$ $3 \times 3 \times 3, 1$
Conv6_x (encoding)							$3 \times 3 \times 3, 1$ $3 \times 3 \times 3, 1$
TransConv7_x	$100 \times 100 \times 100$	$180 \times 180 \times 180$	$150 \times 150 \times 150$	$180 \times 180 \times 180$	$160 \times 160 \times 160$	$160 \times 160 \times 160$	$3 \times 3 \times 3, 1$, stride 2 $3 \times 3 \times 3, 1$
TransConv8_x							$3 \times 3 \times 3, 1$ $3 \times 3 \times 3, 1$
TransConv9_x	$200 \times 200 \times 200$	$360 \times 360 \times 360$	$300 \times 300 \times 300$	$360 \times 360 \times 360$	$320 \times 320 \times 320$	$320 \times 320 \times 320$	$3 \times 3 \times 3, 2$, stride 2 $3 \times 3 \times 3, 2$
TransConv10_x							$3 \times 3 \times 3, 2$ $3 \times 3 \times 3, 2$
TransConv11							$5 \times 5 \times 5, 1$

Extended Data Table 2. Approximate computational time of AlphaCryo4D on the simulated dataset with the SNR of 0.01 and three experimental datasets.

	Consensus alignment (h)	3D volumes bootstrapping (h)	Autoencoder training (h)	Manifold learning (h)	High-resolution refinement (h)
Spliceosome (EMPIAR-10180)	\	48	4	0.1	30
*Pf*80S Ribosome (EMPIAR-10028)	80	30	2	0.1	70
Bacterial Ribosome (EMPIAR-10076)	160	38	4	0.2	105
Simulated dataset with the SNR of 0.01	75	10	3	15	110

Extended Data Table 3. Cryo-EM data collection, refinement and validation statistics. The complete refinement and validation statistics for all 64 proteasome conformers are provided in the Supplementary Information Table. This table shows structure determination statistics of seven high-resolution structures most relevant to the new findings of this work.

	$E_{D0.1}$	$E_{D0.2}$	$E_{D0.3}$	$E_{D1.1}$	$E_{D1.2}$	$E_{D1.3}$	$E_{D2.1}$
Data collection and processing							
Magnification	105,000	105,000	105,000	105,000	105,000	105,000	105,000
Voltage (kV)	300	300	300	300	300	300	300
Electron exposure (e–/Å^2)	44	44	44	44	44	44	44
Defocus range (μm)	-0.6 to -3.5	-0.6 to -3.5	-0.6 to -3.5	-0.6 to -3.5	-0.6 to -3.5	-0.6 to -3.5	-0.6 to -3.5
Pixel size (Å)	0.685	0.685	0.685	0.685	0.685	0.685	0.685
Symmetry imposed	C1	C1	C1	C1	C1	C1	C1
Initial particle images (no.)	3,254,352	3,254,352	3,254,352	3,254,352	3,254,352	3,254,352	3,254,352
Final particle images (no.)	25,755	105,081	169,522	92,777	175,114	146,506	856,683
Map resolution (Å)	3.9	3.3	3.2	3.3	3.2	3.2	2.5
FSC threshold	0.143	0.143	0.143	0.143	0.143	0.143	0.143
Map resolution range (Å)	3.0-8.0	2.5-8.0	2.5-8.0	2.5-8.0	2.5-8.0	2.5-8.0	2.5-4.8
Refinement							
Initial model used			6MSH			6MSJ	6MSK
Model resolution (Å)	4.3	4.1	4.0	3.9	3.8	3.9	3.2
FSC threshold	0.5	0.5	0.5	0.5	0.5	0.5	0.5
Model resolution range (Å)	3.0-8.0	3.0-8.0	3.0-8.0	3.0-8.0	3.0-8.0	3.0-8.0	2.5-5.8
Map sharpening *B* factor (Å^2)	-20	-20	-20	-20	-30	-30	-20
Model composition							
Non-hydrogen atoms	106475	106475	106475	106475	106448	106448	106499
Protein residues	13540	13540	13540	13540	13540	13540	13540
Ligands	11	11	11	11	10	10	11
B factors (Å^2)							
Protein	155.88	155.03	155.53	126.38	112.34	133.02	88.54
Ligand	268.45	173.9	201.79	134.06	115.01	144.41	88.86
R.m.s. deviations							
Bond lengths (Å)	0.006	0.008	0.007	0.006	0.007	0.007	0.007
Bond angles (°)	1.022	1.095	1.076	1.041	1.047	1.072	1.074
Validation							
MolProbity score	1.77	1.82	1.77	1.82	1.45	1.78	1.75
Clashscore	5.05	5.58	5.08	5.57	5.13	5.14	4.49
Poor rotamers (%)	0.29	0.47	0.39	0.44	0.35	0.36	0.37
Ramachandran plot							
Favored (%)	91.56	90.99	91.47	90.98	91.71	91.5	91.24
Allowed (%)	8.15	8.73	8.24	8.75	7.91	8.23	8.47
Disallowed (%)	0.29	0.28	0.28	0.27	0.39	0.28	0.29