



Published in final edited form as:

Nat Methods. 2014 January ; 11(1): 63–65. doi:10.1038/nmeth.2727.

The Local Resolution of Cryo-EM Density Maps

Alp Kucukelbir¹, Fred J. Sigworth^{1,2}, and Hemant D. Tagare^{1,3}

¹ Department of Biomedical Engineering, Yale University, New Haven, United States

² Department of Cellular and Molecular Physiology, Yale University, New Haven, United States

³ Department of Diagnostic Radiology, Yale University, New Haven, United States

Abstract

Using local sinusoidal features in a standard statistical testing framework, we propose a definition of local resolution for 3D density maps. The resulting algorithm has no free parameters and may be extended to other imaging modalities. Evaluating the local resolution of single particle reconstructions and sub-tomogram averages shows variable resolution across a 4 to 40Å range.

Various resolution measures for electron cryo-microscopy (cryo-EM) have been proposed in the past three decades¹. Unlike the classical “Rayleigh” resolution which characterizes instruments, these measures characterize features present in the data. A commonly used cryo-EM resolution measure is the Fourier shell correlation (FSC) procedure. It quantifies the strength, relative to noise, of sinusoidal features across the entire density map. FSC produces a single resolution for the entire density map. FSC cannot assess locally varying resolution, which may be caused by sample heterogeneity and image processing errors². The goal of this paper is to overcome this limitation of FSC by presenting a definition of local resolution which can assess variable resolution.

As a resolution measure, FSC has other limitations, too. FSC resolution depends on the computational stage where the data are split³. Further, calculating the resolution from FSC requires a threshold, whose value and interpretation has been debated¹. Alternative approaches^{4;5} address some of these shortcomings, but do not define local resolution.

Recent structural studies^{6;7} use windowed FSC for local resolution⁸. Windowed FSC masks the split-dataset density maps with a window and calculates FSC resolutions as the window moves through the map. This requires a window size parameter, whose value is often arbitrary. While this approach implicitly conducts multiple tests on the density map, it does not control the false discovery rate (FDR) in the thresholding of the FSC. FDR control is

Users may view, print, copy, download and text and data- mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use: http://www.nature.com/authors/editorial_policies/license.html#terms

Corresponding Author: Hemant D. Tagare (hemant.tagare@yale.edu).

Author Contributions

A.K. and H.D.T. conceived the theory. A.K. developed the algorithm and performed experiments. A.K., F.J.S., and H.D.T. designed the experiments and wrote the manuscript.

Competing Financial Interests

The authors declare no competing financial interests.

critical because local resolution tests are repeated at many points in the volume. Additionally, there is data dependency between neighboring points which windowed FSC does not account for.

We propose a mathematical theory and an efficient algorithm for measuring local resolution that address all of the above limitations. The theory is based on the following idea: A λ Å feature exists at a point in the volume if a 3D local-sinusoid of wavelength λ is statistically detectable above noise at that point. A likelihood-ratio hypothesis test of the local-sinusoid vs. noise can detect this feature at a given p-value (typically $p = 0.05$). We define the local resolution at a point as the smallest λ at which the local-sinusoid is detectable, accounting for multiple testing with an FDR procedure.

Our algorithm, named ResMap, implements this theory. ResMap begins by initializing a local-sinusoid model at $\lambda = 2\mu$, where μ is the voxel spacing in Å. Likelihood ratio tests are conducted at all voxels in the volume, with explicit FDR control that accounts for data dependency. Voxels that pass the test are assigned resolution λ , while those that fail are tested at a larger λ . The algorithm produces a local resolution map with a number assigned to every voxel in the density map (Fig. 1a). There are no algorithm parameters to tune and we may unambiguously speak of finding local resolution at the given p-value.

In ResMap, local-sinusoids of wavelength λ are approximated by a set of functions called H2. This set is derived from Gaussian windowed second-order Hermite polynomials^{9;10}, with window size proportional to the wavelength λ (Fig. 1b and Online Methods). ResMap results with H2 are specifically denoted as ResMap-H2. H2 functions are steerable, so their linear combination can locally model any arbitrarily oriented local-sinusoid in 3D (Supplementary Note 1).

At a fixed wavelength λ , the standard likelihood-ratio test¹¹ can detect whether a local-sinusoid is present in the steerable function approximation. The test requires an estimate of the noise variance, which we obtain from the region surrounding the particle. The likelihood-ratio test does not depend on how this variance is estimated. Other noise estimates, such as those obtained by analyzing split-dataset reconstructions, can also be used. The smallest λ at which the likelihood-ratio test passes at a given p-value defines the resolution. We control for false discoveries using a method that takes into account the dependencies between tests¹² (Online Methods).

We first evaluated ResMap using a simulated density map of a radially symmetrical ‘chirp signal’ whose wavelength decreased with radius. We added white and non-white noise with two different variance levels (Supplementary Fig. 1). ResMap-H2 estimates show an intuitive relation to the underlying signal features (Fig. 1c). Further, increasing the noise worsens the resolution at every point. ResMap-H2 results for this simulation exhibit a ripple in the transitions between the peaks and valleys of the signal. This is because transitions have more energy in the higher frequencies and are thus detectable with local sinusoids of smaller scale.

We then tested ResMap with four different density maps ranging from near-atomic (~ 4 Å) single particle reconstructions to typical sub-tomogram averages (~ 40 Å). All results were

obtained with a p-value of 0.05. We compare ResMap-H2 results to regular and gold-standard³ FSC plots, and windowed FSC maps.

First, we analyzed a single particle 80S ribosome reconstruction (EMD-2275)¹³. The original publication estimates a resolution of 4.5Å (gold-standard FSC at 0.143) and notes the blurring from the heterogeneity in the 40S subunit (Fig. 2a). ResMap-H2 resolution estimates fall between 4.5 and 5.5Å in the 60S subunit and between 4.5 and 9Å in the 40S subunit. Some parts of the 40S are just as resolved as the 60S, which ResMap-H2 results show in the portion of 40S adjacent to 60S. The median ResMap-H2 resolutions in the 40S and 60S subunits are 6.5 and 5Å, respectively, which agree with a map-vs-atomic-model FSC plot (Fig. 3J in the original publication¹³). ResMap-H2 results additionally point to a decrease in resolution near the edges of the particle. This may be due to image alignment errors or the interaction of the ribosome with the solvent.

Second, we analyzed a single particle Tulane virus reconstruction (EMD-5529)¹⁴. The original publication estimates a 6.3Å (gold-standard FSC at 0.143) resolution for the entire particle and highlights the significant flexibility of the protruding domains of the virus. ResMap corroborates these findings, estimating the resolution of the shell between 6 and 7Å and the resolution of the protruding domains between 7 and 9Å (Fig. 2b).

Third, we analyzed a sub-tomogram average of GroEL (EMD-2221)¹⁵. The original publication reports a 8.4Å (FSC at 0.5) resolution. ResMap-H2 estimates suggest that many α -helices are resolved up to 7.5Å (Fig. 2c). This is evident in the close-up rendering that displays the central part of a helix at 7.5Å but the end and adjoining loop at about 9.5Å. These results are corroborated in Supplementary Video 1 where the central part of the helix is shown to maintain its tubular structure under a range of surface threshold values.

Finally, we analyzed a sub-tomogram average of ATP synthase dimers (EMD-2161)¹⁶. The original publication estimates a 37Å (FSC at 0.5) resolution. ResMap-H2 resolution estimates are between 30 and 42Å in the central dimers, which are better resolved than the neighboring dimers and membrane (Fig. 2d). The edges of the central dimers appear to be at a higher resolution than their cores. This is likely due to the strong dark bands surrounding the particle, as is typical in particles reconstructed without contrast transfer function correction.

ResMap is consistent with windowed FSC, but differs in some important aspects. Windowed FSC can be slow, taking anywhere between 25 minutes and 4 hours to compute, depending on the window size. ResMap usually requires a few minutes. More importantly, windowed FSC results appear to be sensitive to the fixed size of the user-defined window (Supplementary Fig. 2). Too large a window size may include the solvent in the FSC computation and lead to underestimation of resolution (Supplementary Fig. 3). ResMap does not suffer from this effect because the spatial-frequency domain localization uncertainty is minimized by the adaptively varying widths of the H2 functions. The behavior of ResMap for typical p-values is shown in Supplementary Figure 4.

For all cases above, ResMap-H2 local resolutions within the reconstructed particle are almost always between the 0.5 and 0.143 threshold of the FSC in the original publications.

This is consistent with the idea that the 0.5 threshold may be too conservative while the 0.143 threshold too optimistic^{3;13}. Moreover, ResMap-H2 results agree with published flexibility analyses and also visually match the level of detail in the density maps.

We have released ResMap as a cross-platform executable package with a simple graphical user interface (Supplementary Fig. 5). The software and test data are publicly accessible (Supplementary Software; <http://resmap.sourceforge.net>). ResMap can also be applied to other fields by choosing features other than local sinusoids. For instance, a 2D Gaussian feature may be appropriate for optical nanoscopy, while rotated 2D arcs may be of use in radio telescopes.

The anticipated² recent increase in heterogeneity studies^{6;7;13} highlights the pressing challenge of evaluating the local resolution of cryo-EM density maps. A local resolution method that is both statistically rigorous and practical is therefore a critical step in enabling researchers to assess the quality of cryo-EM density maps.

Methods

Methods and any associated references are available in the online version of the paper.

Online Methods

This section presents mathematical details of the theory and algorithm. The mathematical formulation is in 3D, but actual computations are performed on column vectors, where the elements in 3D are inserted in order into the vector. To highlight this difference, 3D variables are displayed in ordinary type (A) and their vectorized counterparts in bold type (\mathbf{A}).

Modeling the signal locally

We first describe how the density map can be approximated locally by any basis. Then, we introduce the 3D sinusoid-like feature basis used in our algorithm.

The 3D density map S is in a $V \times V \times V$ voxel array. The voxels are indexed by discrete-valued coordinates x, y, z . We refer to a voxel in the array as $v = (v_x, v_y, v_z)$ where v_x, v_y, v_z are its coordinates. The vectorized density map of S is a $V^3 \times 1$ column vector \mathbf{S} .

Suppose $W_{v, \alpha}$ is a spherically symmetric Gaussian function centered at voxel v with scaling parameter α ,

$$W_{v, \alpha}(x, y, z) = \exp\left(-\frac{\alpha}{2} \left[(x - v_x)^2 + (y - v_y)^2 + (z - v_z)^2\right]\right), \quad (1)$$

and that $\phi_{v, \alpha}^k, k = 1, \dots, K$ are basis functions centered at v . We then have

$$\sqrt{\mathbf{W}_{v, \alpha}^D} \mathbf{S} = \sqrt{\mathbf{W}_{v, \alpha}^D} \Phi_{v, \alpha} \beta + \eta, \quad (2)$$

which locally approximates the density map S by basis functions $\phi_{v,\alpha}^k$. Here $\mathbf{W}_{v,\alpha}^D$ is a diagonal matrix with $\mathbf{W}_{v,\alpha}$ along the principal diagonal, $\Phi_{v,\alpha}$ is a matrix whose columns are the basis functions $\phi_{v,\alpha}^k$, β is a column vector of the coefficients of the basis functions, and η is zero-mean Gaussian noise with variance σ^2 . Note that the weighting function $W_{v,\alpha}$ determines the spatial extent of the local model.

To fit the local model to S we minimize the weighted residual sum of squares (WRSS $_{v,\alpha}$)

$$\text{WRSS}_{v,\alpha} = (\mathbf{S} - \Phi_{v,\alpha}\beta)^T \mathbf{W}_{v,\alpha}^D (\mathbf{S} - \Phi_{v,\alpha}\beta) = \left\| \sqrt{\mathbf{W}_{v,\alpha}^D} (\mathbf{S} - \Phi_{v,\alpha}\beta) \right\|^2, \quad (3)$$

with respect to β . The minimizing coefficient vector is denoted as $\hat{\beta}$.

3D sinusoid-like features

A natural basis for density maps is one containing rotated 3D sinusoids with wavelength λ Å. Unfortunately, describing all orientations in 3D requires an infinite number of basis functions. A computationally tractable alternative is to use steerable filters^{9;10}, which we refer to as “steerable functions”. Steerable functions are a finite set of functions with the property that every 3D rotation of any of the functions is produced by linear combinations of the functions⁹ (Supplementary Note 1).

The steerable functions we use are the second-order Hermite polynomial and its approximate quadrature, multiplied by a Gaussian function. We call this set H2. The elements of H2 match cosine and sine functions up to their second order Taylor expansion terms. They can also be scaled such that their spectral peak occurs at any desired wavelength. The H2 steerable functions are constructed from a pair of functions

$$\begin{aligned} G_{v,\alpha}(x, y, z) &= P_{G_{v,\alpha}}(x, y, z) W_{v,\alpha}(x, y, z) \\ &= \left[4(\alpha(x - v_x))^2 - 2 \right] W_{v,\alpha}(x, y, z), \\ H_{v,\alpha}(x, y, z) &= P_{H_{v,\alpha}}(x, y, z) W_{v,\alpha}(x, y, z) \\ &= \left[(\alpha(x - v_x))^3 - 2.254(\alpha(x - v_x)) \right] W_{v,\alpha}(x, y, z), \end{aligned} \quad (4)$$

where $G_{v,\alpha}$ is the cosine-like function, $H_{v,\alpha}$ is the sine-like function. The scalar α controls the peak frequency and the width of the Gaussian function. Setting $\alpha = 2\pi/\lambda$ 2/ 5 gives a spectral peak for $G_{v,\alpha}$ and $H_{v,\alpha}$ at wavelength λ (Supplementary Note 2).

The functions in Equation (4) are each composed of the Hermite polynomial ($P_{G_{v,\alpha}}$) or its quadrature ($P_{H_{v,\alpha}}$) multiplied by $W_{v,\alpha}$, the spherically symmetric Gaussian weighting function. Their spectral peak occurs on the x -axis of the frequency domain because the functions are oriented along the the spatial x -axis. Rotating $G_{v,\alpha}$ and $H_{v,\alpha}$ so that their spectral peaks occur along the vertices and faces of an icosahedron respectively gives 6 + 10 3D steerable functions. The linear combination of these 16 functions produces all possible rotations of $G_{v,\alpha}$ and $H_{v,\alpha}$ in 3D¹⁰, thereby covering a shell in the 3D Fourier domain (Supplementary Note 1).

We denote the rotated Hermite polynomials as $P_{G_{v,\alpha}}^i$, $i = 1, \dots, 6$ and $P_{H_{v,\alpha}}^j$, $j = 1, \dots, 10$. These polynomials and the constant function $\mathbf{1}$ are our local-sinusoid signal model. Because the weighting function $W_{v,\alpha}$ appears outside of the basis matrix $\Phi_{v,\alpha}$ in Equation (2), the $\Phi_{v,\alpha}$ matrix need only contain the vectorized polynomials,

$$\Phi_{v,\alpha} = \begin{bmatrix} | & | & & | & | & & | \\ \mathbf{1} & \mathbf{P}_{G_{v,\alpha}}^1 & \cdots & \mathbf{P}_{G_{v,\alpha}}^6 & \mathbf{P}_{H_{v,\alpha}}^1 & \cdots & \mathbf{P}_{H_{v,\alpha}}^{10} \\ | & | & & | & | & & | \end{bmatrix}, \quad (5)$$

where the bold face denotes vectorization.

Likelihood ratio testing

Testing whether the data in the neighborhood of voxel v support the local-sinusoid model is equivalent to testing the two hypotheses,

$$\begin{aligned} \mathcal{H}_0 &: \text{The constant term } \beta_0 \text{ is unconstrained, but } \beta_1 = \cdots = \beta_{16} = 0, \\ \mathcal{H}_1 &: \text{All of } \beta_0, \beta_1, \dots, \beta_{16} \text{ are unconstrained,} \end{aligned} \quad (6)$$

where the null hypothesis \mathcal{H}_0 states that the data do not support the local-sinusoid model (the coefficients of all local-sinusoid terms are zero). The alternate hypothesis \mathcal{H}_1 allows the coefficients to take on any finite value.

The likelihood ratio test¹¹ is a standard procedure for comparing such hypotheses. The test works by calculating the β 's that maximize the likelihood (probability) under each hypothesis and then taking the logarithm of the ratio of the maximized likelihoods. The β 's that maximize the likelihood are found by minimizing the WRSS from Equation (3) under \mathcal{H}_0 and \mathcal{H}_1 , respectively. Many common statistical tests such as the Pearson χ^2 test and the F-test are derived from the likelihood ratio test (Supplementary Note 3).

A simple calculation shows that the negative logarithm of the log-likelihood ratio, called the likelihood ratio statistic (LRS), is given by

$$\begin{aligned} \text{LRS}(\mathbf{S}; v, \alpha) &= \frac{1}{\sigma^2} \left(\left\| \sqrt{\mathbf{W}_{v,\alpha}^D} (\mathbf{S} - \mathbf{1}\hat{\beta}_0) \right\|^2 - \left\| \sqrt{\mathbf{W}_{v,\alpha}^D} (\mathbf{S} - \Phi_{v,\alpha}\hat{\beta}) \right\|^2 \right), \\ &= \frac{\mathbf{S}^T (\Gamma_0 - \Gamma) \mathbf{S}}{\sigma^2}, \text{ where} \\ \Gamma_0 &= \mathbf{W}_{v,\alpha}^D - \mathbf{W}_{v,\alpha}^D \mathbf{1} (\mathbf{1}^T \mathbf{W}_{v,\alpha}^D \mathbf{1})^{-1} \mathbf{1}^T \mathbf{W}_{v,\alpha}^D, \\ \Gamma &= \mathbf{W}_{v,\alpha}^D - \mathbf{W}_{v,\alpha}^D \Phi_{v,\alpha} (\Phi_{v,\alpha}^T \mathbf{W}_{v,\alpha}^D \Phi_{v,\alpha})^{-1} \Phi_{v,\alpha}^T \mathbf{W}_{v,\alpha}^D. \end{aligned} \quad (7)$$

The LRS is a difference of weighted residuals between the null model fit and the local-sinusoid model fit. It takes large values when the local-sinusoid model is a better fit than the null model.

The likelihood ratio test is applied by comparing the LRS to a number c , defined by

$$\Pr \left[\frac{\mathbf{S}^T (\Gamma_0 - \Gamma) \mathbf{S}}{\sigma^2} < c \right] = 1 - p, \quad (8)$$

for some p-value p , usually 0.05. If $\text{LRS} < c$, then the data do not support the model and we accept the null hypothesis. Otherwise, we accept the hypothesis that the local-sinusoidal model fits the data.

Calculating the threshold c requires the statistical distribution of the LRS. Unfortunately, because of the weighting function $W_{v, a}$, the LRS does not have a closed form statistical distribution. However, $\mathbf{S}^T (\Gamma_0 - \Gamma) \mathbf{S}$ asymptotically tends to a weighted sum of χ^2 random variables $\sum_r \gamma_r \chi^2$ where γ_r are the eigenvalues¹⁸ of $\Gamma_0 - \Gamma$. Fast and accurate numerical methods are available to compute such distributions¹⁹.

The LRS computation requires the value of the noise variance σ^2 . We estimate this variance accurately by taking non-overlapping cubes of voxels from the region of the density map surrounding the particle. We use the following variance estimator recommended for local signal modeling¹⁸

$$\hat{\sigma}^2 = \frac{1}{B} \sum_{b=1}^B \frac{\mathbf{C}_b^T (\Gamma_0 - \Gamma) \mathbf{C}_b}{\text{trace}(\Gamma_0 - \Gamma)}, \quad (9)$$

where \mathbf{C}_b is a cube of voxels from the background and B is the number of non-overlapping cubes that are available in the background. This estimator is robust to non-white noise as it only requires the noise spectrum to be relatively flat within the shell in 3D Fourier space that the local-sinusoid model, implicitly indicated by Γ , is approximating.

The noise variance may also be estimated from the difference map between split-dataset reconstructions. In this case, the estimator from Equation (9) is adjusted to accept cubes of voxels from the region inside the particle. The noise statistics inside and outside the particle are nearly identical (Supplementary Note 4). Both noise variance estimators are implemented in the accompanying software package (Supplementary Software; <http://resmap.sourceforge.net>).

Multiple testing correction

The likelihood ratio test chooses between two hypotheses at each voxel. Since this test has to be repeated for many voxels in S , some sort of false discovery rate (FDR) control is necessary. The tests in neighboring windows are not independent from each other, therefore we use the Benjamini-Yekutieli FDR procedure¹² that accounts for dependencies between tests.

Summary

ResMap works by applying a hypothesis test at every voxel. The null hypothesis is that the data in the neighborhood of a voxel do not support a local sinusoid. The alternative hypothesis is that the data describe a local sinusoid. These features are modeled by 3D steerable functions. The likelihood ratio statistic is used to decide between the the

hypotheses at a given p-value. Noise variance is estimated from the area surrounding the particle and multiple testing correction is applied to carry out the test at many voxels.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We are grateful to S. Scheres, G. Lander, A. Bartesaghi, and K. Davies for stimulating discussions and sharing their density maps for this study. This work was supported by NSERC award PGS-D3 (A.K.), and NIH grants R01LM010142 (H.D.T.), R01GM095658 (A.K., H.D.T.), and R01NS021501 (F.J.S.).

References

1. Liao H, Frank J. *Structure*. 2010; 18:768–775. [PubMed: 20637413]
2. Leschziner AE, Nogales E. *Annu. Rev. Biophys.* 2007; 36:43–62.
3. Scheres S, Chen S. *Nature Methods*. 2012; 9:853–854. [PubMed: 22842542]
4. Unser M, et al. *J. Struct. Biol.* 2005; 149:243–255. [PubMed: 15721578]
5. Sousa D, Grigorieff N. *J. Struct. Biol.* 2007; 157:201–210. [PubMed: 17029845]
6. Lander GC, et al. *Nature*. 2012; 482:186–191. [PubMed: 22237024]
7. Hashem Y, et al. *Nature*. 2013; 494:385–389. [PubMed: 23395961]
8. Heymann JB, Belnap DM. *J. Struct. Biol.* 2007; 157:3. [PubMed: 17011211]
9. Freeman WT, Adelson EH. *IEEE Trans. Pattern Anal. Mach. Intell.* 1991; 13:891–906.
10. Derpanis KG, Gryn JM. *IEEE Inter. Conf. Image Process.* 2005; 3:553.
11. Lehmann, EL.; Romano, JP. *Testing Statistical Hypotheses*. Springer; New York, USA: 2005.
12. Benjamini Y, Yekutieli D. *Ann. Stat.* 2001; 29:1165–1188.
13. Bai X, Fernandez IS, McMullan G, Scheres SH. *eLife Sci.* 2013; 2
14. Yu G, et al. *PloS ONE*. 2013; 8:e59817. [PubMed: 23533651]
15. Bartesaghi A, Lecumberry F, Sapiro G, Subramaniam S. *Structure*. 2012; 20:2003–2013. [PubMed: 23217682]
16. Davies KM, Anselmi C, Wittig I, Faraldo-Gómez JD, Kühlbrandt W. *PNAS*. 2012; 109:13602–13607. [PubMed: 22864911]
17. Pettersen EF, et al. *J. Comput. Chem.* 2004; 25:1605–1612. [PubMed: 15264254]
18. Loader, C. *Local Regression and Likelihood*. Springer; New York, USA: 1999.
19. Farebrother R. *J. R. Stat. Soc. Ser. C*. 1984; 33:332–339.

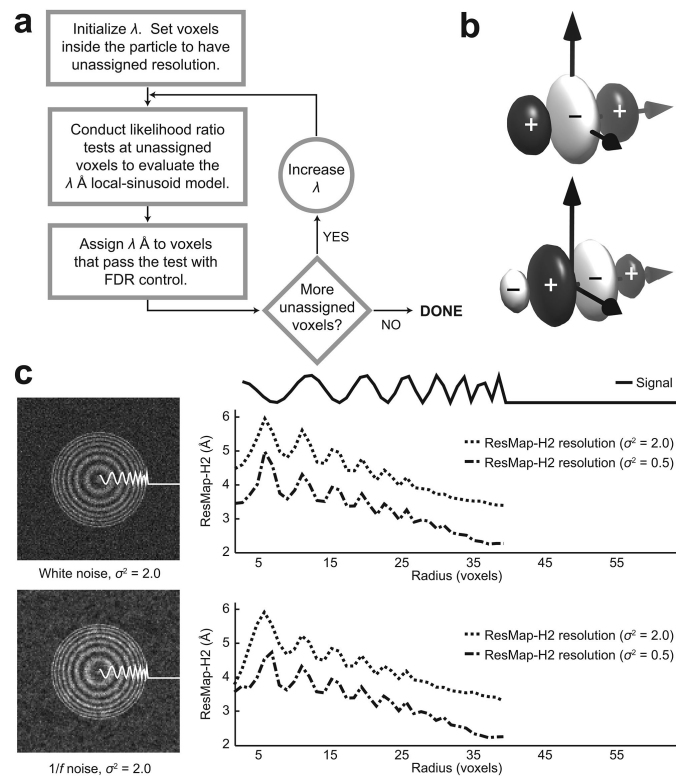


Figure 1. Local resolution. **(a)** The ResMap algorithm. Wavelength λ is initialized to twice the voxel spacing. Likelihood ratio tests decide whether the local-sinusoid model is detectable at each voxel. Voxels that pass the test are controlled for false discoveries. Voxels that fail the test are tested again after increasing λ (Online Methods). **(b)** Cosine- and sine-like H2 functions oriented along an axis. White and black indicate negative and positive parts, respectively (Supplementary Note 1). **(c)** Left, slice through noisy simulated density maps with voxel spacing 1 Å. Right, radial plots. ResMap-H2 resolution estimates show a steady improvement as the simulated signal becomes more finely varying. Bottom, corresponding results for 1/f noise display robustness against non-white noise (Supplementary Fig. 1).

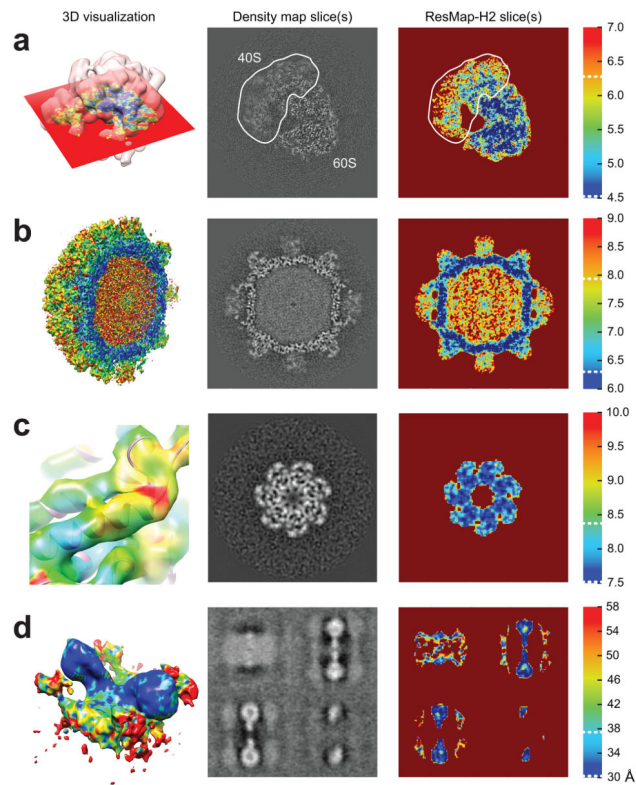


Figure 2.

ResMap-H2 results using experimental density maps. Color bars apply to both volumes and slices. 3D visualizations are rendered using UCSF Chimera¹⁷. White dotted lines in color bars indicate FSC 0.143 and 0.5 thresholds from the original publications. **(a)** 80S ribosome (EMD-2275). ResMap-H2 results indicate a decreased resolution within the 40S subunit and near the edges of the particle. **(b)** Tulane virus (EMD-5529). ResMap estimates lower resolutions in the protruding domains while the shell appears well-resolved. **(c)** Subtomogram GroEL (uncropped version of EMD-2221). ResMap-H2 results show an α -helix with varying levels of resolution (Supplementary Vid. 1). **(d)** Subtomogram ATP synthase dimer (uncropped version of EMD-2161). ResMap delineates the central dimer as better resolved than the neighboring dimers and membrane.