CLASSIFICATION AND AVERAGING OF ELECTRON TOMOGRAPHY VOLUMES

A. Bartesaghi[1]*, P. Sprechmann*[2]*, G. Randall*[2]*, G. Sapiro*[3] *and S. Subramaniam*[1]

[1]Center for Cancer Research, National Institutes of Health, Bethesda, MD 20892
[2]Instituto de Ingeniería Eléctrica, Universidad de la República, Montevideo, Uruguay, 11800
[3]Electrical and Computer Engineering Department, University of Minensota, Minneapolis, MN 55455

ABSTRACT

Electron tomography provides opportunities to determine three-dimensional cellular architecture at resolutions high enough to identify individual macromolecules such as proteins. Image analysis of such data poses a challenging problem due to the extremely low signal-to-noise ratios that makes individual volumes simply too noisy to allow reliable structural interpretation. This requires using averaging techniques to boost the signal-to-noise ratios, a common practice in electron microscopy single particle analysis where they have proven to be very powerful in elucidating high resolution molecular structure. Although there are significant similarities in the way data is processed, several new problems arise in the tomography case that have to be properly dealt with. Such problems involve dealing with the *missing wedge* characteristic of limited angle tomography, the need for robust and efficient 3D alignment routines, and design of methods that account for diverse conformations through the use of classification. We hereby present a computational framework for alignment, classification and averaging of volumes obtained from limited angle electron tomography, providing a powerful tool for elucidation of high resolution structure and description of conformational variability in a biological context.

***Index Terms*—Tomography, Image registration, Image classification, Clustering methods.**

1. INTRODUCTION

Transmission electron microscopes can be used to determine three-dimensional structure using principles that are very similar to those used in technologies such as computerized axial tomography. A series of projection images spanning a limited angular range is obtained by tilting the specimen relative to the electron beam, and combine them using back projection algorithms to generate a three-dimensional volume of the imaged object. Electron tomography technology provides opportunities to determine three-dimensional cellular architecture at resolutions of approximately 50 Å or better, it is especially applicable for the structural analysis of organelles and other macromolecular assemblies that are too heterogeneous to be investigated by NMR or X-ray crystallographic techniques. The development of reliable computational approaches for the analysis of electron tomography data is a challenging problem due to the extremely low signal-to-noise ratios and high structural complexity that are inherent to biological electron microscopic images. One of the factors limiting signal strength comes from the potential of electrons to damage organic matter, which necessitates the use of electron doses that are high enough to obtain measurable contrast, but low enough to minimize structural damage. This is particulary true for cryo-microscopy where reconstructed volumes are simply too noisy to allow direct interpretation of high resolution structural information. The idea is to register and then average multiple noisy copies of the same structure together so that signal strength can be improved. The benefits of using averaging techniques for boosting the signal-to-noise ratio are well established and routinely used in electron microscopy single particle analysis [1, 2]. There are though, fundamental differences in the type of data being processed which is now three-dimensional, very heterogeneous and affected by the missing wedge. In this context, averaging will only make sense within a classification framework that prevents copies coming from different conformational structures to be averaged together. Current techniques for 3D averaging of cryo tomographic data require availability of an initial reference, to which all volumes are aligned and averaged regardless of the conformational heterogeneity [3, 4]. Many times, selection of the reference is done on a subjective basis resulting in high resolution maps that are strongly biased and poorly representative of the underlying data.

We present a framework for alignment, classification and averaging of tomographic volumes that deals with noisy, heterogenous and wedge affected volumes and does not require initial references. The idea is to identify and register volumes that represent similar conformations so they can be averaged together for improving the signal-to-noise. As a key component of the framework, we introduce in Section 2 a fast 3D alignment technique that takes the missing wedge into account, enabling for the first time computation of complete distance matrices for a more accurate approach to classification

This work has been partially supported by ONR, NSF, DARPA, and NGA. PS and GR performed part of this work while at the University of Minnesota visiting ECE and IMA.

than current multivariate statistical analysis techniques [4].

2. ALIGNMENT OF VOLUMES AFFECTED BY THE MISSING WEDGE

Volumes reconstructed from limited angle tomography are intrinsically affected by a missing wedge of information in Fourier space. According to the central slice theorem, see *e.g.* [5], each measured projection provides frequency information in a slice through the origin (orthogonal to the incident beam) typically covering a ± 70 degree range, see Fig. 1 (left). Given two such volumes, the alignment problem consists in finding the rigid transformation between the two that minimizes a given dissimilarity criteria. In three-dimensions this is a six degree-of-freedom (DOF) problem, one rotation and one translation per coordinate axis. Standard 3D alignment techniques (including widely used cross-correlation) are known to give biased estimates as they tend to align the areas of missing frequency information to one another [6].

To deal with this problem we restrict the computation of the dissimilarity measure to the overlapping regions in Fourier space between the two volumes. Since the overlap region changes as volumes are rotated, this has to be done differently for each tentative 3D rotation requiring the use of exhaustive search to determine the best matching rotation. This however, will have no practical interest due to the huge computational burden involved. To overcome this difficulty we solve the alignment in two stages, rotations are first recovered within the Spherical Harmonics (SH) framework and then spatial shifts are obtained by Fourier based cross-correlation. The reasons for using SH are two, one is the existence of a fast convolution theorem that allows direct calculation of the rotations (avoiding the brute force search) [7]. Second, the missing wedge will have a trivial decomposition in this representation (see below). In order to decouple the six DOF problem into two problems of three DOF each, for the first stage we use the amplitude of the Fourier Transform (FT) which is invariant to translations. In the SH framework we represent each volume as a scalar image defined on the unit sphere S^2, obtained by integrating the amplitude of the 3D FT along rays though the origin. Using this representation, the missing wedge simply appears as an occluding mask on the sphere as seen in Fig. 1 (right). If we denote the two scalar images on the sphere by $a(\mathbf{x}), b(\mathbf{x}) : S^2 \rightarrow \Re$, their matching score in the presence of occluding masks $m_a(\mathbf{x}), m_b(\mathbf{x}) : S^2 \rightarrow \{0, 1\}$ is defined for each $\mathcal{R}$ in the 3D rotation group as:

$$S(\mathcal{R}) = \int_{x \in S^2} \left[a(x) - b(\mathcal{R}^t x)\right]^2 m_a(x)\, m_b(\mathcal{R}^t x)\, dx \quad (1)$$

and the goal is to find $\mathcal{R}$ that minimizes this quantity. Using the convolution theorem on the sphere, Eq. 1 can be expressed as a series of convolutions:

$$S = a^2 m_a \star m_b - 2(a m_a) \star (b m_b) + m_a \star b^2 m_b, \quad (2)$$

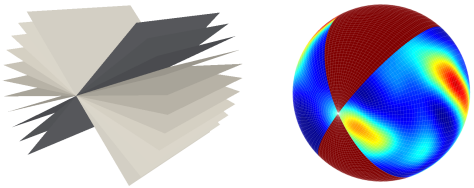

Fig. 1. Missing wedge in limited angle tomography. *Left*: Projections are taken in a limited angular range providing slices of information in Fourier space (central slice theorem). *Right*: Spherical Harmonics representation of a wedge affected volume obtained by projecting (onto the unit sphere) the magnitude of its 3D Fourier Transform along rays through the origin.

which can be computed rapidly as products of FTs. The best $\mathcal{R}$ is readily obtained by localizing the minima within the correlation volume S. The accuracy of this solution is limited by sampling of the SH grid, the number of coefficients used in the SH expansion and also because we are projecting the volumetric data into a lower dimensional space (the sphere). We therefore refine this estimate using off-the-shelf optimization routines that locally maximize the 3D cross-correlation (computed exclusively on overlapping frequency components), also yielding the shifts in the three coordinate axis.

2.1. Performance of the alignment routine

In general, success of the alignment routine depends on the size of the missing wedge with respect to image features and also the amount of noise present. Under certain circumstances, the available information in Fourier space may not be enough to allow correct matching of volumes even in the noise-free case. We evaluate the performance of our alignment technique under different amounts of noise and missing wedge sizes on a three-legged non-symmetric 3D phantom model. For each *(noise, wedge)* pair we produced randomly rotated copies under the same noise and wedge conditions and computed the average alignment error (comparing the Frobenius norm of the linear transformation matrices to the correct values). In Fig. 2 we show alignment error plots using the usual cross-correlation alignment and the technique described in Section 2. In the first case, alignment performance degrades quickly for tilt ranges smaller than ± 70 and modest noise levels, while in the second case the low error area extends significantly up to ± 55 degrees and shows better tolerance to noise. To illustrate the potential implications of this, we built a three-fold symmetric phantom with a simulated extension in one of the legs mimicking the situation we expect to find when studying the interaction between macromolecular complexes of dissimilar sizes by cryo-electron tomography. Seventy randomly oriented copies of the phantom were generated. Projections of each volume were taken in a ± 60 range, and Simultaneous Iterative Reconstruction Technique (SIRT) was used for reconstruction. We aligned all volumes in the set using regular cross-correlation alignment as

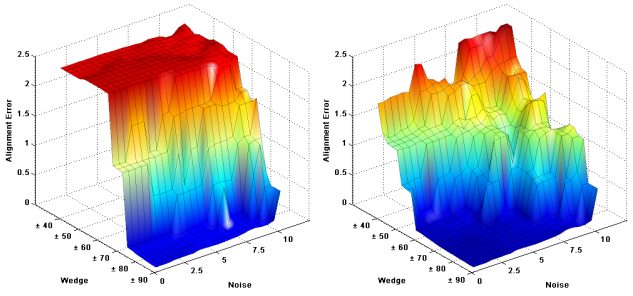

Fig. 2. Average alignment error for different sizes of missing wedge and noise levels. *Left*: Using the usual cross-correlation alignment, total MSE = 3.72. *Right*: Using the technique introduced in Section 2, total MSE = 1.37.

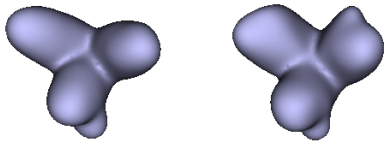

Fig. 3. Effect of alignment errors due to the missing wedge in volume averaging. *Left*: Average of 70 volumes aligned using the technique presented in Section 2 produces correct localization of the small feature in one of the legs. *Right*: Same average obtained after cross-correlation alignment results in spreading of the feature among the three legs due to alignment errors.

well as our technique and computed the corresponding volume averages. Results are shown in Fig. 3, observe how using cross-correlation alignment (even in the noise-free case) the small feature spreads out in all three legs due to frequent alignment errors while our proper accounting of the missing wedge maintains the correct geometry.

3. VARIABILITY ANALYSIS THROUGH CLASSIFICATION

In order to deal with the great heterogeneity characteristic of biological material and at the same time be able to achieve signal-to-noise ratios that permit high resolution interpretation, a classification strategy to averaging must be adopted making sure averages are computed strictly from homogeneous sets of volumes. In practice, some degree of true density variability will be sacrificed in favor of improving the signal level. The value of classification techniques is that of providing a glimpse of the landscape of conformations so that an educated decision can be made on what level of averaging is most appropriate. The classification problem can be stated as the following: given a set of volumes that are noisy and rotated copies (affected by the missing wedge in different orientations) of many different conformations, the goal is to recover class averages that represent each of the conformations. The number of representative classes is initially unknown and we assume that for each class there are enough copies in uniformly distributed orientations so that isotropic resolution can be attained in the average maps.

Our approach to classification is divided in three stages. First, we compute a distance matrix by aligning all pairs of volumes in the set. The definition of distance is that introduced in Section 2 that measures similarity between volumes based only on overlapping frequency components. Although still computationally demanding, this is now feasible because we have overcome the need for brute force search. More importantly, this will enable a more accurate approach to classification that is unbiased and does not require selection of a reference. Next, a hierarchical clustering technique using complete linkage (farthest neighbor) is used to build an initial dendrogram [2]. Cutting the tree at a small enough threshold, will yield classes containing volumes that are very similar to one another. The idea here is to get an initial and very conservative set of references that will later be refined. Note that many volumes may be left out this initial classification (depending on the number of neighbors and how close they are to each other), and homogeneous classes may be split down into sub-classes of stronger similarity. In the third stage, these classes are used as seeds to an iterative procedure (reminiscent of k-means) that progressively refines class averages by re-assigning volumes to references. In each iteration, volumes are assigned to the closest reference using an elimination step to get rid of outliers. To allow early associations to be changed as the loop advances, volumes may be assigned to a different reference in each iteration. References are continuously compared to each other to determine if they actually represent the same conformation so they can be merged together.[1] We use *bundle alignment* to guarantee that volumes within a class are optimally aligned using an iterative reference-free algorithm [2]. The classification loop is terminated when no significant changes are observed indicating convergence was achieved.

The refinement step is somewhat similar to that used in single particle analysis [1], as before however, there are some fundamental differences due to the fact that these are 3D wedge affected volumes. Progressive averaging will not only increase the signal-to-noise, but will also effectively decrease the size of the missing wedge further improving the accuracy of subsequent alignments. This however, requires careful computation of averages in Fourier space to prevent missing wedge regions to be averaged with areas of valuable information (this will happen if averages are computed directly in image space). Unlike the 2D case, here we can incorporate a number of geometric constraints within the classification loop that result in improved robustness. For example, we can use the linear transformations between all pairs of volumes (computed and stored during the first step of the classification), to enforce alignment consistency checks within sets of volumes. That is, if volumes (A, B) are related by the linear transformation $T_{A,B}$ and (B, C) by $T_{B,C}$, then A and C should be related by $T_{A,B} \bullet T_{B,C}$, otherwise, this may be an indication

[1] References are merged together if the distance between them falls below a given threshold.

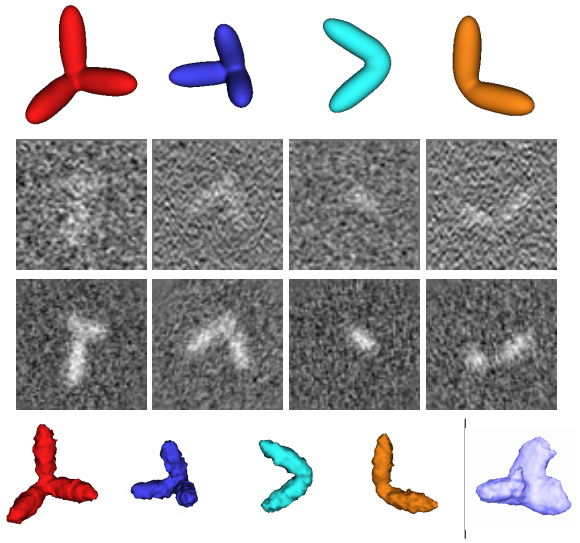

Fig. 4. Classification experiment on a set of 80 volumes reconstructed from noisy ± 60 degree projections. *Top*: Four original phantom conformations. *Center*: Slices through some of the initial SIRT reconstructed volumes and final averages after classification. *Bottom*: 3D class averages after classification (left) and average of all volumes without classification (right).

that alignment errors have occurred. Also, when assigning a volume to a reference we consider not only its distance to the reference, but also the distances to all volumes contributing to that reference, as well as the amount of wedge overlap, so associations can be made much more accurately.

3.1. Classification experiments

We analyze the performance of the clustering method with two different experiments. In the first, four different phantoms were generated, Fig. 4 (top). Phantom configurations were deliberately selected to produce high degree of similarity between densities despite being different from one another, *e.g.*, one of the trimer configurations is obtained by adding an extra leg to the dimer configuration. Twenty randomly rotated copies of each model were produced giving a total of 80 volumes. As routinely done in cryo-tomography, forty projections of each were taken in a ± 60 range and noise was added to each projection. Volumes were then reconstructed using the SIRT reconstruction technique. The distance matrix, as well as all alignments within the loop, were computed using bandpass filtered versions of the input volumes. Nine classes were extracted from the initial hierarchical clustering and used to initialize the loop. After three iterations, four class averages were obtained, shown in Fig. 4. Also shown in this figure is an average obtained without proper classification as currently done in cryo-tomography studies, showing how inaccurate this is for describing heterogeneous datasets. For the second experiment, we generated phantoms in three stoichiometry configurations with different quantities in each group: monomers (55), dimers (75) and trimers (45). We also added to the set 25 volumes generated as random noise making a total of 200 volumes. The distance matrix was computed and twenty initial classes obtained by thresholding the hierarchical clustering dendrogram. After running the loop, the three class averages were recovered with all outlier volumes left out (results not shown). The code for the experiments was parallelized to take advantage of the high performance computational capabilities of the Biowulf Linux cluster at the National Institutes of Health, Bethesda, MD.

4. CONCLUDING REMARKS

We presented a framework for alignment, classification and averaging of volumes obtained from limited angle electron tomography. We showed the potential misinterpretations arising from not carefully considering the wedge and the variability in this type of analysis. To illustrate the importance of our framework, we used realistic 3D phantoms since ground-truth was needed. Experiments with real data from our lab are currently underway. Availability of such image processing tools has enormous implications as it will confer new capabilities to electron tomography technology, allowing a whole new spectrum of discoveries in molecular and cellular biology.

5. REFERENCES

[1] van Heel M, Gowen B, Matadeen R, Orlova EV, and Finn R, "Single-particle electron cryo-microscopy: Towards atomic resolution," *Q Rev Biophys*, vol. 33, 2000.

[2] Joachim Frank, *"Three-Dimensional Electron Microscopy Of Macromolecular Assemblies : Visualization Of Biological Molecules In Their Native State"*, Oxford Univ Press, February 2006.

[3] F. Forster, O. Medalia, N. Zauberman, W. Baumeister, and D. Fass, "Retrovirus envelope protein complex structure in situ studied by cryo-electron tomography," *PNAS*, vol. 102, no. 13, pp. 4729–4734, 2005.

[4] H. Winkler, "3D reconstruction and processing of volumetric data in cryo-electron tomography," *Journal of Structural Biology*, August 2006, In Press.

[5] F. Natterer and F. Wuebbeling, *Mathematical Methods in Image Reconstructions*, SIAM, Philadelphia, 2001.

[6] A. S. Frangakis, J. Bohm, F. Forster, S. Nickell, D. Nicastro, D. Typke, R. Hegerl, and W. Baumeister, "Identification of macromolecular complexes in cryoelectron tomograms of phantom cells," *PNAS*, vol. 99, no. 22, pp. 14153–14158, 2002.

[7] P. J. Kostelec and D. N. Rockmore, "FFTs on the rotation group," Tech. Rep., Santa Fe Institute, 2003.