2013

Towards Independent Particle Reconstruction from Cryogenic Transmission Electron Microscopy

W. Lewis Collier

Jean Yves Hervé

University of Rhode Island, jyh@uri.edu

Lenore Martin

University of Rhode Island, martin@uri.edu

Follow this and additional works at: https://digitalcommons.uri.edu/cmb_facpubs

The University of Rhode Island Faculty have made this article openly available.
Please let us know how Open Access to this research benefits you.

This is a pre-publication author manuscript of the final, published article.

Terms of Use

This article is made available under the terms and conditions applicable towards Open Access Policy Articles, as set forth in our [Terms of Use](#).

Citation/Publisher Attribution

W. Lewis Collier, Jean Yves Hervé, and Lenore Martin. 2013. Towards Independent Particle Reconstruction from Cryogenic Transmission Electron Microscopy. In *Proceedings of the International Conference on Bioinformatics, Computational Biology and Biomedical Informatics (BCB'13)*. ACM, New York, NY, USA, , Pages 525 , 10 pages. DOI=10.1145/2506583.2506622
<http://doi.acm.org/10.1145/2506583.2506622>

This Article is brought to you for free and open access by the Cell and Molecular Biology at DigitalCommons@URI. It has been accepted for inclusion in Cell and Molecular Biology Faculty Publications by an authorized administrator of DigitalCommons@URI. For more information, please contact digitalcommons-group@uri.edu.

Towards Independent Particle Reconstruction from Cryogenic Transmission Electron Microscopy

W. Lewis Collier^{*}
i-Particle Systems
86 Indian Trail
Saunderstown, RI, 02874 USA
LCollier@i-Particle.com

Jean Yves Hervé
University of Rhode Island
Department of Computer
Science and Statistics
Kingston, RI 02881 USA
jyh@cs.uri.edu

Lenore Martin
University of Rhode Island
Department of Cell and
Molecular Biology
Kingston, RI 02881 USA
martin@uri.edu

ABSTRACT

Coronary heart disease is the single largest killer of Americans so improved means of detecting risk factors *before* arterial obstructions appear are expected to lead to a improvement in quality of life with a reduced cost. This paper introduces a new approach to 3-D reconstruction of *individual* particles based on statistical modeling from a sparse set of 2-D projection images. The method is in contrast to the current state of practice where reconstruction is performed via signal processing or Bayesian methods that use averaged images acquired from an ensemble of particles. As such, this new approach has its impetus in use for novel diagnostic tests such as LDL and HDL particle shape characterization. The approach is also expected to have uses in areas such as quality assurance for drug delivery nano-technologies and for general proteomic studies.

The individual particle reconstruction algorithm is based on hidden Markov models. Higher order Markov chain statistics are generated from the a priori model of the target of interest. This model can be generated from traditional methods such as single particle reconstruction and/or the underlying physical properties of the particle. The basic approach can reconstruct simple models from a single image but can be extended to include a sparse set of images taken at small rotation angles. Reconstruction results from a simple model and a sparse set (where $n=1$) of simulated projection images are presented.

Categories and Subject Descriptors

I.4.7 [Image Processing and Computer Vision]: Feature Management—*size and shape*; I.4.8 [Image Processing and Computer Vision]: Scene Analysis—*object recognition*; I.4.10 [Image Processing and Computer Vision]: Image representation—*volumetric*; J.3 [Life and Medical Sciences]: Biology and genetics

^{*}Primary contact author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM-BIB '13 Washington, D.C. USA

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

General Terms

Algorithms, Experimentation, Measurement

Keywords

Hidden Markov Model (HMM), individual particle reconstruction (IPR), single particle reconstruction (SPR), proteomics, low density lipoprotein (LDL), high density lipoprotein (HDL)

1. INTRODUCTION

This paper presents initial results from research directed towards creating a 3-D reconstruction of transparent objects from a sparse set of 2-D images. Here, a limited number of orthonormal projection images of an individual transparent object is available to be used in the reconstruction processing. A classic example of the reconstruction of transparent objects is imaging of biological and nano-particles via cryogenic electron microscopy (*cryo-em*). See [30] for a historical and theoretical review of *cryo-em*. In this processing, the targets are flash-frozen from their natural state so that structure may be preserved. A *cryo-em* example from [28] of a low density lipoprotein (LDL) particle (approximately 25 nm in diameter and 10 nm in height) imaged at relatively large rotation angles is shown in Figure 1.

Figure 1: Example LDL *Cryo-em* Images of an individual particle viewed from a) $+45^\circ$, b) 0° , and c) -45° .

This sparse projection image approach is in contrast to the currently available 3-D reconstruction techniques for transparent objects. These non-sparse methods, which include acquisition and processing of many thousands [12, 3] of individual images taken at varying spherical angles around a large population of targets, allow for a representation of the target to be reconstructed. The two methods, however, are somewhat symbiotic since the new method can provide initial models for the iterative processing for the multiple im-

age approaches while the prior methods can provide detailed models to be used in the individual particle reconstruction technique.

The remainder of this paper is organized as follows. The rest of this section provides background information. Section 2 describes the processing in the new approach. Initial results and discussion are provided in sections 3 and 4.

1.1 Single Particle Reconstruction

The current three-dimensional reconstruction techniques in electron microscopy have led to what is called “single particle reconstruction” (SPR). The visualization of biological macromolecules via SPR utilizes a uniformity assumption in order to achieve high-fidelity rendering of 3-D representations from multiple 2-D projection images. In essence, these techniques align the images and average out the sample noise and variations in the population samples by processing the various pictures of the “same” object into a single result space. As noted by Frank [10], this SPR nomenclature can be confusing since the reconstruction process traditionally uses many thousands of projection images. The name implies that a single particle class (e.g. a virus) is reconstructed and not that pictures of a just single item, are used to perform the reconstruction. It is a set of distinct images, however, that are used to reconstruct a class average from a representative ensemble of distinct particles.

While the problem of solving a SPR is a complex procedure, there are solutions. “The problem of reconstructing a 3D object from a set of 2D projections was in principle solved analytically by Radon at the beginning of the twentieth century (Radon, 1917). The principles of 3D reconstruction were re-invented in the late 1960s and early 1970s, with the introduction of computerized tomography in medicine and 3D reconstruction techniques in electron microscopy.” [29]. SPR utilizes many aspects of Radon’s principles in order to reconstruct nano-scale biological particles. See [16] for a treatment of the Radon transform and its application to electron microscopy processing. As is seen, the complexity is not so much in the math, per se, but in the alignment and filtering of all the images that are required to provide coverage of the entire object 3-D spatial extent. The solutions derived from SPR research allows for some simplifying assumptions in the IPR approach.

A first consideration when processing *cryo-em* images is alignment of the images. This alignment can be used in many ways but two main objectives are orientation classification for averaging and angular coverage. Averaging is used to reduce the tens of thousands of images to a subset that has reduced noise characteristics. Angular coverage is used in SPR to prevent “holes” in the reconstruction analysis by ensuring that samples provide coverage across a suitable portion of the full range of spherical angle space. Alignment of such images can be handled via the Fourier-slice theorem. A complete description of the underlying mathematics can be found in [9] and other descriptions can be found in resources such as [22]. Thus, the general alignment problem has solutions that can be used for IPR.

A second issue with *cryo-em* images is correction for the imaging modality. Images generated by *transmission electron microscopy*, regardless of whether or not the sample is frozen, are not true 2-D projections of the 3-D sample space. Various artifacts and noise are present that degrade the information in the projection images. Noise can be induced

from many sources but with appropriate care, these can be ignored in the presence of other factors. The two largest contributors to distortion are the contrast transfer function (CTF) and the envelope function of the microscope [7, 11]. While the shape of the CTF (see, e.g., Figure 1 in [19]) depends on several parameters such as defocus, spherical aberration coefficient, source size, and defocus spread, only defocus is a parameter that changes between images [6]. Despite the fact that information is lost at the CTF’s zero crossings, there are ways to overcome this effect. A simple method of correcting is to truncate the data at the first zero crossing. However, many more robust techniques have been developed [19, 15] to estimate the CTF so that the true projection image can be recovered. Thus, for reconstruction from a limited number of images, acquisition of the proper defocus images should allow for recovery of the underlying projections to be used in the reconstruction processing from sparse images.

The several algorithms used in variations of SPR also offer model generation support for IPR. Single particle tilt series (SPTS) often do use just an individual particle as a target. By utilizing a much lower electron beam energy, SPTS allows for capture of 60 to 280 images of the same particle [8]. However, the lower energy restricts the achievable resolution so other means may be more applicable to providing an a priori model for the IPR approach. The random conical tilt (RCT) process takes advantage of the target’s preferred orientation in the sample holder [21]. RCT only takes two images of the sample space; one at some desired rotation angle and the second one at a nominally untilted angle. Since the particles are in the same basic orientation in the untilted image, this image provides a “known” orientation angle for alignment processing. With only two exposures, the electron beam can be more energetic to achieve higher resolutions. In fact, only the rotated image, which has but a single exposure, is used for reconstruction so sample damage is vastly minimized as an input into the reconstruction process. The angular reconstitution (AR) method, which is summarized in [24], uses the random nature of particles in suspension in order to glean the requisite FT slices across a range of spherical angles without requiring a physical tilting of the *cryo-em* sample. An alternate version of this is the Albany Zero Tilt (AZT), which uses but a single image of the sample. The alignment and further processing is performed via a series of post processing steps. Thus, there are several existing techniques for gleaning a model of a target particle that can be used in statistical-based reconstruction techniques.

1.2 Bayesian Reconstruction

Statistical methods may also be used for reconstruction of particles from *cryo-em* images. A recent paper [13] utilizes a Bayesian method for inferring macromolecular structure from SPR-style microscopy images. Rather than modeling the target itself, this method creates an estimation of the 3-D space by observing how the projection image pixels are created by the electron beam passing through the material. The vertices of the voxels are used as the definition points of the coulomb density of the 3-D space. A trilateral interpolation of the vertices along each ray to the pixels is used as a means of statistically connecting the pixel values to the voxel densities. A probabilistic model for the pixel values from the line integration along the rays is then used to feed an iterative process for maximizing the probability that the

vertices give rise to the given pixel values.

The images used for this processing are averaged images, much like those used for standard single particle reconstruction processing. This method also assumes that the alignment of the particles is sufficiently accurate to not unduly influence the estimate of the 3-D space. As with SPR, the averaging of the images implies a uniformity of all the particles in the ensemble class. But, the averaging can also affect the statistics of the pixel values. A pure Bayesian approach needs to understand the standard deviation of the pixel values in the images so that the probabilities of the voxels can be estimated. The paper makes an assumption that the standard deviation of the noise is the same for all pixels. As with the rotational alignment, an assumption of images corrected for microscope issues such as CTF makes sense since these are well understood. But, an assumption of constant sigma across all pixels does not seem plausible since the number of electrons received for each pixel changes based on the voxels involved in the pixel. The electron flux will clearly have an effect on the variance of the pixel values. As can be easily measured, and was shown in [5], the variance of the pixel values is not constant. Sigma of *cryo-em* pixel images was shown to range from 1 to 14 gray levels for a 3×3 region when the mean ranged from 50 to 200. As with the size of the region over which the variance is computed, the number of images in the average also plays an important role in the measured variance. The study [13] notes that differing numbers of images used for averaging at various rotation angles was a possible cause of errors in the estimates.

This approach also has other constraints that limit its applicability. The paper uses a constraint about there only being small changes in intensity between neighboring grid points [13]. From biological inference, we know that there can be significant changes in coulomb density from voxel to voxel as atoms and molecules with various densities exist side by side. Independence is a primary concern in simplifying statistical computations. While the results shown in [13] were fairly good approximate reconstructions, from the above constraints, it is not clear if this method can be used in general practice. As with SPR techniques, it does *not* provide details of individual particles and the above issues may preclude the use of its reconstructed model as a basis prior for the new approach.

1.3 LDL Particles

While the IPR method has applications in many areas, the initial biologic targets of interest are lipoprotein particles. Thus, an understanding of the LDL particle and related research is provided to motivate the new sparse image approach.

As coronary heart disease (CHD) is the single largest killer of Americans [2], improved means of detecting risk factors *before* arterial obstructions appear is expected to lead to a improvement in quality of life with a reduced cost. Geometric parameters other than average diameters or density measurements may one day prove more important to achieving a better understanding of and prediction of heart disease risk. By allowing for efficient generation of a histogram of geometric parameters from LDL and HDL particles found in blood samples, information derived from *cryo-em* images may be correlated to observed cardiovascular state in order to assist in the determination of essential relation-

ships between lipoprotein geometry and overall cardiovascular health. Tools that can provide these parameters can also be used for other proteomic and nano-particle studies.

The characteristic parameters of LDL and HDL packages are important to a better understanding and prediction of heart disease risk. Past research has suggested that geometric parameters of lipoprotein macromolecules may better correlate with cardiovascular risk than just the particle counts and the associated density of particles in the serum. In general, LDL particles that fall on the large end of the LDL particle size spectrum are tied to better health while smaller LDL particles correlate to worse health. In a study of diabetics [26], changes in the LDL sizes and plasma lipid levels accounted for part of the antiatherogenic effect of fenofibrate in type-2 diabetes. In this study, increases of only 0.98nm in average LDL diameters (which is in the range of %5 to %10 of expected particle size) showed a detectable change in atherogenesis. LDL packages that are at the small end of the overall LDL size spectrum are considered to be an important cause of arterial plaque initiation that can eventually lead to heart attacks and stroke. And somewhat surprisingly, even though conventional wisdom states that HDL is “good” cholesterol, another study [18] showed that high concentrations of smaller sized HDL particles resulted in a 15-fold increase in the risk of heart disease. Newer test methods that are based upon Nuclear Magnetic Resonance (NMR) spectroscopy can provide a count of the number of LDL particles in presorted bin sizes but not actual sizes of individual LDL particles [14]. These studies imply that more is at play than just the size of the particles: shape and/or other geometric properties may be important.

Initial *cryo-em* micrograph studies [27] reported that the LDL particles may be more discoid in shape than spherical as had been previously assumed. Figure 1 shows example *cryo-em* images of an LDL particle at rotations of $+45^\circ$, 0° , and -45° . This study measured the diameter and height of the discs but did not measure the thickness of the wall as seen in the projection in frame A of Figure 1. These studies, however, were difficult to carry out since they did not apply computer vision to provide a way to perform measurements on samples from a large number of objects from each image. Such a study is necessary to evaluate the effect of geometric shape parameters on overall health in a statistically significant patient sample but, to date, the necessary processing tools were not available.

Other studies have been performed in an attempt to resolve the LDL shape question, but these did not provide actual metrics on the shapes. In one study [17], computer-based analytical methods were applied to *cryo-em* images by applying a homogeneity requirement on subclasses of LDL particles. In this analysis, approximately 5600 individual LDL particles were manually divided into subclasses. Further classification reduced the data to approximately 4000 particles. These particles were further analyzed and were determined to be generally ellipsoidal in shape. Another study [25] used volume/mass ratios of High Performance Gel Chromatography (HPGC) to indirectly determine LDL shape. This study fitted data based on a geometric model of a discoid object with varying heights rather than directly from images of the macromolecules.

One of the more recent full-featured reconstructions of LDL particles was reported in 2010 [23]. This study reports a reconstructed 3-D volume that shows the shape of

LDL particles to be somewhat like a flattened walnut with a slight bump on the narrow end. In this analysis, ≈ 8500 particle images were selected from a pool of ≈ 48000 particle images to be processed. The resultant reconstructed shape is a slightly flattened ellipsoid and a size of approximately 250 Å by 240 Å by 166 Å. The detailed reconstruction from [23] (See Figure 2) also shows internal layers of alternating cholesterol esters (CE) and fatty acyl chains. The reconstruction shows a very advanced model of LDL

Figure 2: 2010 LDL *Cryo-em* Reconstruction

that can be used to understand the structure of the particle. But, this study also demonstrates a major concern with heterogeneous sample sets: about ≈ 40000 (or 82%) of the particles were *not* included since they were of different size or shape. These studies all used large numbers of images but resulted in but a single class average description of the biological target of interest. They do not provide a pathway to understanding the individual parameters of a large sample of proteomic particles. But, the underlying model maybe useful in processing of individual particles.

2. PROCESSING

The over-arching goals of processing with a limited set of images are to provide spatial parameter information for each individual particle and to do so in a reasonable processing architecture. As noted above, SPR processing provides a detailed ensemble average rather than details about each particle. These methods also can require multiple passes of $O(N \log(N))$ computations for each image times the tens of thousands of images processed. The statistical-based processing shown below can provide details on each particle within very reasonable computational constraints.

2.1 Software Framework

A software framework was developed to test and analyze various reconstruction methods. The software suite is called **ARTEMIS** (**A**-priori **R**econstruction from **T**ransmission **E**lectron **M**icroscope **I**mage **S**ets). The major parts of this framework are a projection simulation engine and the reconstruction processing. Both processing streams make use of a common voxel space. The utility allows both the color and opaqueness of each density to be specified so that the target can be viewed showing the densities. A simple three layer disc, which emulates a simplified version of the expected LDL shape from [28], is shown in the figure but this simulated shape can be any set to any configuration programmatically or from a file. The 3-D space is then filled with the expected density (transmission coefficient) for each

voxel based on the simulated target. Voxels that are not part of the object can be left unfilled or can be set to a value to simulate the expected statistics of the buffer solution in the *cryo-em* setup. Examples of the target model and a simulated projection image is shown later in Figure 5a. The voxel space and resulting projection images are scale independent to allow for reconstruction from images taken at any scale.

The images and statistics are created from rays passing through the voxel space. The pixel value for each location in the 2-D projection is computed from the transmission losses “seen” at each pixel as the ray trace passes through the simulated target. Thus, as with the *cryo-em* equipment, darker and lighter pixel values can be determined by the amount of density in the target through which the ray passes. Likewise, the voxels intersected by the rays can be analyzed to compute statistics.

The 2-D projections of the simulated 3-D targets are generated as orthonormal projections of the voxel space onto the images. The center of the voxel space is mapped to the rotated center of the projection image. Similarly, each rotated pixel center location is projected through the target voxel space. The projection ray for each pixel is used to compute intersection lengths for each voxel that contributes to the given pixel. This process is repeated for each desired pixel in the projection image.

The computation of the simulated pixels accounts for the voxel intersection lengths, camera gain, and camera DC offset. This allows the pixel values to be scaled to match the camera output range. The computation treats each voxel as a region of constant transmission medium. The length of the intersection through a voxel is also used to adjust the transmission loss so that shorter traversal lengths have less effect than a long length. Each simulated pixel value, p_j , is expressed in Equation 1

$$p_j = E_j \prod_{i=0}^I (T_i^{(L_i/L_{max})}) + C_j \quad (1)$$

where L_i is the length of the intersection of the current beam through the i^{th} voxel and L_{max} is the maximum possible voxel intersection for the overall voxel size (e.g. opposing corners). This approach follows the physics of *cryo-em* where an intersection length that approaches 0 through a region of constant density results in a transmission coefficient approaching 1 so that there is very little loss. And when the intersection length nears the maximum intersection length, a value nearer to the base transmission coefficient T_i of the voxel is used. The constants E_j and C_j provide the camera scale and offset adjustments. These constants are expected to be the same for all pixels in a projection image but are provided on a per-pixel basis in case camera distortions are to be included in the projection images. Given the transmission coefficients are all less than or equal to 1, the initial energy coefficient, E_j , also helps prevent underflow from the product of many small numbers. Additional issues such as electron beam damage, imaging noise, and CTF issues can be added to the model as required.

The utility uses projection images and reference data created by the program and generates the sequence of observations required for reconstruction of the 3-D voxel space via the HMM approach. By convention, stacks of the voxel space are used as the state sequences for the HMM. These stacks may be oriented along any of the three axes in either

the plus or minus direction. Since the most probable hidden state sequence is dependent upon the direction, this provides for six different renderings of the voxel space. Observation generation follows the same basic idea as projection image creation but the rays are based on the voxel centers rather than the pixel centers. Once the intersection with the image is determined, the observation is created from the pixel values neighboring the intersection point.

2.2 Hidden Markov Models

Hidden Markov Modeling (HMM) methods of reconstructing 3-D voxel spaces offer algorithms that do not require the exponentially scaled computations of correlation processing nor the extensive collection of images required for averaging methods such as SPR. Relationships between the observables and the underlying object can lead to estimates of the 3-D structure utilizing the relationship. Some methods, such as Bayesian methods seek to determine the most probable “explanation” of the observations given a physical model of how the observed images are created. An example of this is discussed in Section 1.2. Other methods, such as hidden Markov models, utilize an understanding of the 3-D target to help find the most probable set of states that explain the observations. An example of this processing that uses tree ring sizes to infer the most probable weather is discussed in [1].

A HMM is a form of a Bayesian network that has been in use since the late 1960’s for a variety of state-machine modeling applications. See [20] for a seminal paper that provides a tutorial on HMM processing and its use in speech recognition. In a HMM, the system is modeled as a state machine that gives rise to a set of observed measurements. A representation of an HMM system is shown in Figure 3. The Markov process is denoted by the sequence of unknown

Figure 3: Representative Hidden Markov System

states s_t and the corresponding T observations are denoted by O_t . While the representation of a Markov system is shown as states and observations versus time, this could just as easily be a spatial relationship so the T subscript could be K , a traditional spatial or wave number index. The probabilities of the process transitioning from state to state are denoted by A_{s_i,s_j} and the probability of a given state giving rise to an observation is shown as B_{s_i,O_t} . A first order HMM (FOHMM) must also satisfy additional requirements of a) each current state only depends on prior state, b) time independence (stationary) of the model, and c) observation independence.

Note that the model as given does *not* account for all things that could affect the hidden states. In the given example [1] of tree ring widths used to infer prior weather, the tree ring sizes could be affected by the density of the forest,

or rainfall, or cloudiness. While this could be a limitation of the technique, it can also be an asset. By ignoring the other factors, the most probable state sequence is the sequence that is best explained by the factors of concern. Stamp [1] alludes to this in the example he presents based on [4]. In this example, a set of characters from the English language is modeled with just two states. As may be expected, a 50000 character sample of text is divided into vowels and consonants. But, the work in [4] was able to decipher results for models with up to 12 hidden states. Thus, other factors do not necessarily affect the results when a smaller number of factors are considered.

2.3 3-D Reconstruction via HMM

HMM processing was investigated because of the inherent match between the requirements of the HMM processing and the desired goals of the individual particle reconstruction. These models are based on an underlying state to state relationship. The nature of biological particles implies some sort of order. For the example target particle, the well-known “states” of phosphate head group (red), lipid (green), protein (blue), and buffer (white) were chosen as a first set of states. The shape was deduced from the original images in [28]. Assuming the sample is not imaged to the point where damage occurs, the frozen nature of the imaged samples ensures that the states do not change. Due to the structure of HMM processing, state and observation independence is needed. The construction of the problem space can guarantee the necessary constraints while the efficient nature of HMM processing leads to a computationally tractable solution for reconstructing individual particles.

Basic first order Markov sequences, however, are not useful for this 3-D reconstruction. The greedy statistical nature of a FOHMM means that, in general, the voxel state that fills the majority of voxels in a stack will most likely dominate for a given observation even if the projection image is at an angle to the voxel stack. As can be seen in Figure 2, a current model of the LDL particle consists of several layers of alternating fat and cholesterol encased in the phospholipid head group shell. Which ever density prevails in a voxel stack will dominate the simple first order statistics processing. This results in a reconstruction that cannot determine the height of the particle. The reconstruction is simply a cylinder where the top slice is *red* then the remaining slices of the middle region are *green* since the *green* statistics dominate. It is this internal structure, however, that helps make use of higher order HMM (HOHMM) processing feasible.

HOHMM systems generally rely on additional state to state information while a FOHMM only uses the prior state; both use the observation statistics. HOHMM’s can utilize more than just the prior state and have also included the chain’s duration in a given state to provide more robust results. From the basic model provided, it is obvious that just using the prior two states will not provide any additional information. For example, the *white* to *red* transition in figure 5a would simply result in requiring a *white-white-red* sequence in order to make the jump from the *white* state to the *red* state at the top of the reconstructed particle. But this will only result in the top two reconstruction slices being *white* then the same pattern cylindrical will continue as before. Likewise, if a *red-red-green* sequence is required to

get from *red* to *green*, the top band of *red* will now be two rows high instead of just one. Given the preponderance of *green* in the middle, the *green-green-green* sequence will be more probable so the same resulting stack of *green* in the middle will result. Even higher order models can be conjured up but this defeats the purpose of the HMM. Another metric is needed in order to help the hidden Markov chain be found.

The external buffer as well as the shell and internal layers of the simple model can be reconstructed by using a HOHMM that includes state duration in addition to prior state information. Thus, the HMM A matrix is not simply a probability that a state transitions to another state but it also includes an extent or duration component. This HMM denotes the elements of state transition matrix A by Equation 2.

$$A(S_t, S_{t-1}, S_{t-2}) = P(S_{t-1} \rightarrow S_t | S_{t-1} \text{ is length } L_1 \text{ and } S_{t-2} \text{ is length } L_2) \quad (2)$$

For this model, the two prior state sequence lengths provide sufficient detail to allow the HMM to determine the structure of the particle. By including external unknown states (which is effectively the vacuum outside of the frozen buffer) of length 1, the *white* buffer sequence is statistically determined by the top of the particle model (red). As can be seen in Figure 2, the top of the particle is not truly flat so a statistical model of the buffer lengths to the *top* of the particle needs to be deduced from the model. Likewise, all internal sections (states) of the particle can be described statistically by the preceding hidden state layer lengths. For our simple model, the two preceding state lengths were used but more complicated structures may require a higher order model.

The second order HMM that included state duration that was found to be effective in reconstructing the simple discoid target uses an extended A matrix. In this schema, the elements of the A matrix were defined as

$$A'_{f',s_n} = P\{s_f, L_f, s_p, L_p \rightarrow s_n\}, \quad (3)$$

where f' is the 4-tuple of the *from* node that is defined as s_f , the *from* state, L_f , the length the chain was in the *from* state, s_p , the *previous* state, L_p , the length the chain was in the *previous* state, and s_n , which is the *next* state (*next* = n is used instead of $to = t$ in order to prevent confusion with the observation step t). The new $a(f' = \{f, L_f, p, L_p\}, n)$ values can be easily computed from the model voxel space. As with the original A matrix, the elements of A' are normalized so that the probability of transitions *from* each prior-state tuple state is row stochastic.

Using the state identifiers of *next_state* :: $n = j$ and *set_of_states* :: $s = i$ and converting the A subscripts to generic array indices, the basic α -pass equation [1, 20] can be rewritten as

$$\alpha_{t+1}(n) = \left[\sum_{s=1}^N \alpha_t(s) a(s, n) \right] b_n(\mathcal{O}_{t+1}), \quad (4)$$

for $1 \leq t \leq T - 1$ and $1 \leq n \leq N$. The summation part of the equation processes all of the edges of the n^{th} trellis sub-structure and the loop over the set of n states computes all of the fans in the trellis to the *next* step. Thus, the alpha term for each possible state at the *next* step in the chain

($\alpha_{t+1}(n)$) is given by the probability that the possible *next* state gives rise to the observation at the next step in the chain ($b_n(\mathcal{O}_{t+1})$) times the sum of all the α terms *from* the prior step in the chain multiplied by the probability that the chain can transition *from* state $s = f$ to the n^{th} state. Note that while not changed here, the s index is really the *from* index given the trellis structure, which can be converted to make the processing a HOHMM.

The α pass processing was modified to include the *from* and *previous* information. Figure 4 shows how this can be flattened to convert the HOHMM into a FOHMM. This expansion occurs since there now is a node for each tuple of {*from* state, *from* length, *previous* state, and *previous* length} instead of just the N basis states. The left side of

Figure 4: Expanded HOHMM Trellis Structure

the single trellis fan shows example *from*, *previous* nodes as 4-tuples denoted by {*from*, *previous*} state and length information. Note that some 4-tuples cannot coexist at each step in the overall trellis. For example, the first state-tuple shown cannot occur concurrently with the second state-tuple since S_1 cannot be of length 1 and 2 at the same time step. However, the A probability may exist since in one area of the voxel space the $(S_1, 1, S_2, 1) \rightarrow (S_3, 1, S_4, 1)$ transition may occur and in another area the $(S_1, 2, S_2, 1) \rightarrow (S_3, 1, S_4, 1)$ transition may happen. The f, p state node listed at the bottom of the *from* side of the sub-trellis represents the extra “unknown” states beyond the voxel space. Here, S_{n+1} and S_{n+2} represent the states outside of the known voxel space and these are expected to only have lengths of 1 so there are no further 4-tuples needed for these states.

3. RESULTS

A test of this processing was performed that reconstructed the example target exactly. Figure 5 shows the original model (rotated at 45° and zoomed in to show the internal structure) in the left panel, the simulated projection image (not rotated) in the center panel, and the reconstructed particle (also rotated and zoomed) in the right panel. The HMM statistics for matrix A modeled the fact that the *white* buffer region was of length T , T_{top} , or T_{bottom} states. These statistics also captured the effect that the sequence of unknown of length 1 followed by a *white* buffer region of length T_{top} gave rise to a transition to the *red* state. Likewise, a sequence of *white* to *red* with appropriate lengths predicted

a transition to either *blue* or *green*, where the selected state depended upon the associated observation value. Similarly the *blue* or *green* regions transitioned back to *red* then *white* based on the modeled sequence lengths.

This test shows that the basic processing is functioning but it does not provide any significant insights into the overall reconstruction issue for arbitrary angles. This single-pass reconstruction relies on the fact that the target was square to the reconstruction voxel space. The statistics match exactly “from the top” and there is no ambiguity of observations and states.

Additional reconstruction results were obtained by performing the HMM processing in four orthogonal directions as depicted in Figure 6. In effect, the reconstruction space was

Figure 6: Extended HOHMM Processing Directions

rotated around the z axis. The reconstruction voxel space was constructed at $\pm 45^\circ$ to the orientation of the modeled target in order to produce independent observations for each step in the chains (yellow boxes). The directions were chosen so as to keep the reconstruction space aligned to the rows and columns of the projection images. This artificial constraint was chosen to help with analysis; a rotation around the y axis could have also been performed. The same HMM model parameters, which were deduced from the $+y$ to $-y$ direction for all 4 images, were used for the 4 reconstruction passes. Thus, the HMM state chains followed the directions of the black arrows.

The results of these 4 reconstructions are shown in Figure 7. As is shown, the basic particle shape is also reconstructed but with a “tail” in the trailing direction of the HMM chain. There is no “tail” on the leading edge since the statistics of the “white space” say that there can be no particle states until at least the minimum distance from the edge has been traversed by the state chain. On the trailing edge, however, the state to state probabilities *can* result in particle states, thus the tails when the observations support it. As with the external tails, the internals are also subject to distortion along the HMM state chain since the view is at an angle to the direction of the state-to-state probabilities. The detailed voxel to voxel comparison is provided in

Table 1. As is shown, the reconstruction space matches the original space to 98.76%.

Orginal / Recon	Blue	Green	Red	White
Blue	7770	420	0	30
Green	150	10365	0	0
Red	0	0	12470	20
White	1665	0	1110	240625

Table 1: 45° Reconstruction Voxel Accuracy

Since the 4 reconstructions share the same aligned voxel space, the separate results can be easily combined without the need for state interpolation or other schemes. Given the orthogonal method by which the reconstructions were generated, a simple majority scheme can be employed to deduce the final reconstruction. Since, for the most part, the tails extend along the trailing edge, each voxel is defined by the three images that are not being affected by the trailing effect. Figure 8 shows an example projection image with noise. Figure 9 shows the reconstruction merged with a

Figure 8: Projection image at 45° with noise added .

simple majority scheme. This reconstruction also matches the original model better than 98% but the tail artifacts have been reduced.

While the similar results with non-averaged noise may seem counter-intuitive, a look at the state-to-observation statistics gives insight into why HMM processing works. Figure 10 shows the observation to state probabilities. The left side shows the base statistics for no rotation of the voxel space. The peak at the right is for the pure buffer case; all observations are high in the gray level space. The lower two peaks are for the combined projection pixels. Thus, only a handful of observation values are allowed. The spread denotes the noise level added; if there was no noise, the allowable regions would be delta functions at the resulting projection gray levels. The right side of the figure shows the effects of rotating the voxel space. The buffer region stays clearly separated but the two peaks for the particle voxels start to merge. The HMM processing, however, uses *both*

Figure 5: a) Original Discoid model, rotated to 45° and sliced to show internal structure, b) Resulting projection image taken along the Y axis, and c) HMM reconstruction results, also rotated to 45° and sliced.

Figure 7: Reconstruction results for the four rotated and non-rotated cases of a) X $+45^\circ$, b) X -45° , c) Z $+45^\circ$, and d) Z -45° as viewed from the $+y$ axis.

Figure 10: State (color) to observation probabilities for a) aligned voxel space and b) rotated voxel space.

Figure 9: Merged reconstruction results shown at 45° and sliced open to show the internals.

the state-to-state and state-to-observation statistics so the effect of noise is minimized.

4. CONCLUSIONS

There are many classes of problems that need individual particle reconstruction. The presented approach offers promise that a simple method that requires significantly less computational processing than the current state-of-practice methods exists. The processing shown requires on the order of a minute of processing on a modest laptop computer versus the hours of computations for existing methods. Even processing of 50,000 separate images would require only a half day or so of computing and the results would be for each individual particle, not the ensemble average. The method described in this paper can provide reconstruction results from modeled targets and simulated projection images that match the original target better than 98% on a voxel to voxel basis. The individual particle reconstruction shown here also has an ability to work through the noise found in non-averaged images and with care in setting the contrast range of the images, it is expected that even better per-

formance can be achieved. For problems such as medical diagnostics and quality assurance for drug delivery vesicles, the proposed IPR approach may lead to better quality and new understandings in various proteomic studies.

Further study is needed to achieve these goals. A partial list of topics for further study includes angle dependencies, size scaling, and merging schema. These, and other issues, are currently being investigated.

5. ACKNOWLEDGMENTS

Thank you to the reviewers for their helpful comments.

Special thanks to Mr. Peter Morley for his assistance in development of some of the software used in this study.

6. REFERENCES

- [1] A revealing introduction to hidden markov models, Apr. 2012.
- [2] Heart diseases, Mar. 2013.
- [3] CARRAGHER, B., KISSEBERTH, N., KRIEGMAN, D., MILLIGAN, R. A., POTTER, C. S., PULOKAS, J., AND REILEIN, A. Leginon: An automated system for acquisition of images from vitreous ice specimens. In *Journal of Structural Biology* (Oct. 2000), pp. 33–45.
- [4] CAVE, R. L., AND NEUWIRTH, L. P. Hidden markov models for english. *Hidden Markov Models for Speech* (Oct. 1980).
- [5] COLLIER, W. L. Svm analysis of cryogenic transmission electron microscopy images. In *URI CSC-581, Topics in AI: Support Vector Machines, Spring 2007, Final Project* (May 2007), pp. 1–22.
- [6] Contrast transfer function (ctf) correction, Aug. 2012.
- [7] ERICKSON, H. P., AND KLUG, A. The fourier transform of an electron micrograph: Effects of defocusing and aberrations, and implications for the use of underfocus contrast enhancement. In *Berichte der Bunsengesellschaft für physikalische Chemie* (Nov. 1970), pp. 1129–1137.
- [8] FERNÁNDEZ, J., AGULLEIRO, J., BILBAO-CASTRO, J., MARTÍNEZ, A., GARCÍA, I., CHICHÓN, F., MARTÍN-BENITO, J., AND CARRASCOSA, J. Image processing in electron tomography. In *Microscopy*:

Science, Technology, Applications and Education: A. Mendez-Vilas and J. Diaz (Eds.) (2010), pp. 19–28.

- [9] FESSLER, J. *Chapter 3. Analytical Tomographic Image Reconstruction Methods*. web.eecs.umich.edu/fessler/course/516/1/c-tomo.pdf, TBD, 2009.
- [10] FRANK, J. Single-particle reconstruction of biological macromolecules in electron microscopy Ū 30 years. In *Q Rev Biophys* (Aug. 2009), pp. 139–158.
- [11] HANSZEN, K. J. The optical transfer theory of the electron microscope: Fundamental principles and applications. In *Advances in Optical and Electron Microscopy* (1971), Academic Press, New York, pp. 1–84.
- [12] HENDERSON, R. The potential and limitations of neutrons, electrons, and x-rays for atomic resolution microscopy of unstained biological macromolecules. In *Quarterly Review of Biophysics* (May 1995), pp. 171–193.
- [13] JAITLEY, N., BRUBAKER, M. A., RUBINSTEIN, J. L., AND LILIEN, R. H. A bayesian method for 3d macromolecular structure inference using class average images from single particle electron microscopy. *Bioinformatics* 26, 19 (Oct. 2010), 2406–2415.
- [14] JEYARAJAH, E. J., CROMWELL, W. C., AND OTVOS, J. D. Lipoprotein particle analysis by nuclear magnetic resonance spectroscopy. In *Clinics in Laboratory Medicine* (Dec. 2006), pp. 847–870.
- [15] LUDTKE, S. J., BALDWIN, P. R., AND CHIU, W. Eman: semiautomated software for high-resolution single-particle reconstructions. In *J. Struct Biol* (Dec. 1999), pp. 82–97.
- [16] MIDGLEY, P. A., AND WEYLAND, M. 3d electron microscopy in the physical sciences: the development of z-contrast and efem tomography. *ultramicroscopy* 96 (2002), 413–431.
- [17] ORLOVA, E. V., SHERMAN, M. B., CHIU, W., MOWRI, H., SMITH, L. C., AND JR., A. M. G. Three-dimensional structure of low density lipoproteins by electron cryomicroscopy. In *Proc. Natl. Acad. Sci. USA* (July 1999), pp. 8420–8425.
- [18] OTVOS, J., MARABINI, R., MASEGOSA, I., MARTÍN, M. S., MARCO, S., FERNÁNDEZ, J., DE LA FRAGA, L. G., VAQUERIZO, C., AND CARAZO, J. Radio signals give new spectrum for cholesterol lipoprotein readings. In *Arteriosclerosis, Thrombosis and Vascular Biology: Journal of the American Heart Association* (July 1998), pp. 2223–2231.
- [19] PENCZEK, P., ZHU, J., SCHRÖDER, R., AND FRANK, J. Three dimensional reconstruction with contrast transfer compensation from defocus series. In *Scanning Microscopy* (1997), Scanning Microscopy International, pp. 147–154.
- [20] RABINER, L. R. A tutorial on hidden markov models and selected applications in speech recognition. In *Proceedings of the IEEE* (1989), pp. 257–286.
- [21] RADERMACHER, M., WAGENKNECHT, T., VERSCHOOR, A., AND FRANK, J. Three-dimensional reconstruction from a single-exposure, random conical tilt series applied to the 50s ribosomal subunit of escherichia coli. In *Journal of Microscopy* (May 1987), pp. 113–136.
- [22] The central section theorem, Feb. 2010.
- [23] REN, G., RUDENKO, G., LUDTKE, S. J., DEISENHOFER, J., CHIU, W., AND POWNALL, H. J. Model of human low-density lipoprotein and bound receptor based on cryoem. In *PNAS* (Jan. 2010), pp. 1059–1064.
- [24] SCHATZ, M., ORLOVA, E. V., DUBE, P., STARK, H., ZEMLIN, F., AND VAN HEEL, M. Angular reconstituion in three-dimensional electron microscopy: Practical and technical aspects. In *Scanning Microscopy Vol. 11* (1997), pp. 195–210.
- [25] TEERLINK, T., SCHEFFER, P. G., BAKKER, S. J. L., AND HEINE, R. J. Combined data from ldl composition and size measurement are compatible with a discoid particle shape. In *Journal of Lipid Research* (May 2004), pp. 954–966.
- [26] VAKKILAINEN, J., STEINER, G., ANSQUER, J.-C., AUBIN, F., RATTIER, S., FOUCHER, C., HAMSTEN, A., AND TASKINEN, M.-R. Relationships between low-density lipoprotein particle size, plasma lipoproteins, and progression of coronary artery disease : The diabetes atherosclerosis intervention study (DAIS). In *Circulation* (www.circulationaha.org) (Mar. 2003), pp. 1733–1737.
- [27] VAN ANTWERPEN, R., AND GILKEY, J. C. Cryo-electron microscopy reveals human low density lipoprotein substructure. In *Journal of Lipid Research* (Dec. 1994), pp. 2223–2231.
- [28] VAN ANTWERPEN, R., LABELLE, M., NAVRATILOVA, E., KRAUSS, R. M., LUNA-CHAVEZ, C., FORTE, T. M., AND GILKEY, J. C. Structural heterogeneity of apob-containing lipoproteins visualized using cryo-electron microscopy. In *Journal of Lipid Research* (Oct. 1999), pp. 1827–1836.
- [29] VAN HEEL, M., GOWEN, B., MATADEEN, R., ORLOVA, E. V., FINN, R., PAPE, T., COHEN, D., STARK, H., SCHMIDT, R., SCHATZ, M., AND PATWARDHAN, A. Single-particle electron cryo-microscopy : towards atomic resolution. In *Quarterly Reviews of Biophysics* (Apr. 2000), pp. 307–369.
- [30] VAN HEEL, M., ORLOVA, E. V., HARAUZ, G., STARK, H., DUBE, P., ZEMLIN, F., AND SCHATZ, M. Angular reconstituion in three-dimensional electron microscopy: Historical and theoretical aspects. In *Scanning Microscopy Vol. 11* (1997), pp. 195–210.