

Detecting circular and rectangular particles based on geometric feature detection in electron micrographs

Zeyun Yu* and Chandrajit Bajaj

*The Center of Computational Visualization, Department of Computer Sciences and Institute of Computational Engineering and Sciences,
The University of Texas at Austin, Austin, TX 78712, USA*

Received 21 July 2003, and in revised form 10 October 2003

Abstract

Accurate and automatic particle detection from cryo-electron microscopy (cryo-EM images) is very important for high-resolution reconstruction of large macromolecular structures. In this paper, we present a method for particle picking based on shape feature detection. Two fundamental concepts of computational geometry, namely, the distance transform and the Voronoi diagram, are used for detection of critical features as well as for accurate location of particles from the images or micrographs. Unlike the conventional template-matching methods, our approach detects the particles based on their boundary features instead of intensities. The geometric features derived from the boundaries provide an efficient way for locating particles quickly and accurately, which avoids a brute-force searching for the best position/orientation. Our approach is fully automatic and has been successfully applied to detect particles with approximately circular or rectangular shapes (e.g., KLH particles). Particle detection can be enhanced by multiple sets of parameters used in edge detection and/or by anisotropic filtering. We also discuss the extension of this approach to other types of particles with certain geometric features.

© 2003 Elsevier Inc. All rights reserved.

1. Introduction

The single particle technique (Baker et al., 1999; Frank, 1996; Heel et al., 2000) has been widely used for 3D reconstruction of large molecular complexes from electron micrographs. In contrast to X-ray diffraction technique, the single particle method does not require formation of crystals. However, the signal-to-noise ratio (SNR) in most cryo-EM images is very low due to various reasons, such that high-resolution single particle analysis often has to rely on averaging of a large number of identical particles (Frank, 1996; Heel et al., 2000). Therefore, locating most, if not all, of the particles in the digitized cryo-EM images is a crucial step in high-resolution single particle reconstruction. This task, commonly known as *particle picking* or *particle detection* in single particle analysis, can certainly be carried out manually (by mouse clicks). However, as the reconstructed resolution approaches the atomic level, hundreds of thousands of particles may be necessary

(Henderson, 1995), which makes it impractical to manually pick the particles. In addition, particle detection by visual observation may be inaccurate and fairly subjective.

Several methods have been proposed for automatic or semi-automatic particle detection (see Nicholson and Glaeser, 2001 for a recent review). The first automatic method was proposed by Heel (1982). In this approach the local variance over a small area around each point is computed, and each local maximum of the variance map is then considered as a particle. Another commonly used approach is based on template-matching (Frank and Wagenknecht, 1984; Saad et al., 1998; Thuman-Com-mike and Chiu, 1995, 1996), where the template is chosen as a rotationally averaged image of manually picked particles. The template is cross-correlated with the entire image and the “peaks” of the resulting cross-correlation map are identified as particles. This method, however, may fail for non-spherical particles or for multiple-view particles. A more robust approach is to use multiple references, each of which stands for one of the non-symmetric orientations or one of the views of the same object/particle (Ludtke et al., 1999; Stoschek

* Corresponding author. Fax: 1-512-471-8694.

E-mail address: zeyun@cs.utexas.edu (Z. Yu).

and Hegerl, 1997). This approach is computationally intensive when multiple views or orientations are considered (Zhu et al., 2002). A more recent approach is based on the fast local correlation technique (Roseman, 2003). Due to the numerical scaling, this method proves to be more sensitive than the conventional cross-correlation methods, in discriminating peaks from the cross-correlation maps.

The above-mentioned techniques, as well as some other methods such as the crosspoint method (Martin et al., 1997), the texture analysis method (Lata et al., 1995), the ring-filter based method (Kivioja et al., 2000), and the neural network approach (Ogura and Sato, 2001), are all explicitly or implicitly based on intensity comparison. In other words, all the intensities within a local window around a pixel have to be considered and compared in order to determine whether or not a particle is located at that pixel.

Another group of particle detection algorithms are based on edge detection. In Andrews et al. (1986), the authors employ edge detection to estimate the perimeter of the particle being considered, from which the mass of the particle is calculated as a measure for particle selection. However, this method only works for dark field electron micrographs. A more general approach, proposed by Harauz and Fong-Lochovsky (1989), uses edge detection followed by component labeling and symbolic (high-level) processing. This method relies on the connectivity of the detected edges. As the signal-to-noise ratio becomes too low, inaccurate edges may make it difficult to deal with some situations (e.g., edges of two particles are incorrectly connected or edges of one particle are incorrectly split). Recently, Zhu et al. proposed a method, based on edge detection followed by Hough transforms, for automatic detection of particles (Zhu et al., 2001, 2002). This method is more robust to noise because using the Hough transform one can extract critical features from the edges, and all the features are then integrated to make the decision.

In this paper we present a new method for automatic particle detection. Our method, similar to the one proposed by Zhu et al. (2001, 2002), is based on edge detection followed by feature extraction from the edge map. However, our method uses two fundamental constructions from computational geometry, namely, the Voronoi diagram and the distance transform. Unlike conventional template-matching techniques, where the templates are “real” intensity images, we only need to know the geometric information of the true particles (e.g., the radius of circular particles or width/length of rectangular particles). A “virtual” shape is generated from the given geometric information and matched against the edge map of the original electron micrograph. The scoring function of this feature-based approach is evaluated with the help of the distance transform. The speed of the detection process can be

largely improved by a good initialization of the “virtual” shape in the edge map with the help of the Voronoi diagram and the distance transform. This avoids a brute-force search for the best position/orientation. We also describe how to refine the centers/orientations of the detected particles by moving the particles along the distance map. In the case of circular and rectangular particles (e.g., *Keyhole Limpet Hemocyanin* (KLH) particles Zhu et al., 2002, as we consider in the following), the competitions between particles (*circle–circle*, *rectangle–rectangle*, and *circle–rectangle*) are introduced to guarantee a very small number of false positives in most of the images we investigated. We also discuss several other issues such as the use of multiple sets of parameters for Canny edge detection, the signal-to-noise improvement by anisotropic filtering and the extension of our geometric-feature-detection based method to other shapes.

The rest of this paper is as follows. Section 2 gives a brief review of the Voronoi diagram and the distance transform. We describe the details of our method in Section 3, where circular and rectangular particles are considered. Results and discussions on several other issues will be given in Sections 4 and 5, respectively. Finally we describe our conclusions and future work.

2. Related work: distance transform and Voronoi diagram

Given a set of discrete points (called *feature points*) P_i , $i = 1, 2, \dots, n$ in 2D space \mathcal{R}^2 , we define the *nearest neighbor*, denoted as $NN(A)$, of an arbitrary point A as one of the feature points P_k such that $d(A, P_k) \leq d(A, P_i)$ for any $i = 1, 2, \dots, n$, where the function $d(*, *)$ is Euclidean distance function between two points. Therefore, given a set of feature points in \mathcal{R}^2 , we can define a function value for any point A as the Euclidean distance between A and its nearest neighbor $NN(A)$. The obtained function map is known as the *distance transform* (DT) of the feature points. An example of the distance transform can be seen in Fig. 1A, where the feature points are shown by white dots for better illustration although the distance map is always zero at any feature point.

If we recall the definition for the *nearest neighbor*, we would notice that some points in \mathcal{R}^2 might have two or more nearest neighbors. Those points are of great interest in computational geometry. In fact, all of those points form a well-known diagram, called the *Voronoi diagram* (VD). A Voronoi diagram generally consists of a set of polygons, each of which corresponds to one feature point, and all polygons form a *partition* of the space \mathcal{R}^2 . Fig. 1B shows an example of a VD.

Both the distance transform and the Voronoi diagram have been widely used in computer vision, image analysis as well as many other fields. In this paper, we shall show how they are applied to particle detection

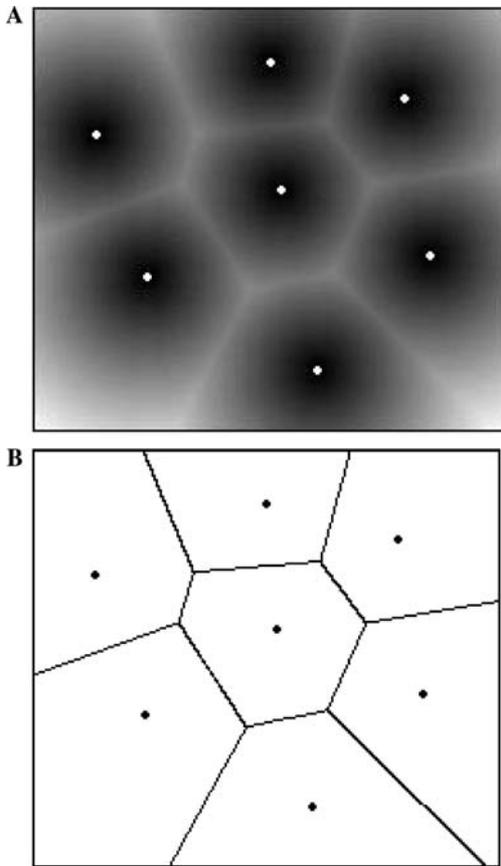


Fig. 1. An example of distance transform and Voronoi diagram. (A) Distance transform and (B) Voronoi diagram.

from electron micrographs. Although many algorithms exist for fast computations of DT and VD, we found that the list-processing approach (Guan and Ma, 1998) was most suitable for our application. This approach calculates the nearest neighbor of every pixel in the digitized space of \mathcal{R}^2 and stores the results in a number of lists of segments, from which the DT and VD can be easily calculated.

3. Approach

In this section we shall describe the details of our particle detection algorithm. We consider only two types of particles here: circular particles and rectangular particles (e.g., KLH data). Particles with other shapes can be detected in a similar way if certain geometric features of the particles are available, as discussed in Section 5. Our method is based on geometric feature detection from particle boundary edges and template matching against the edge map. We do not need the “real” intensity templates. This avoids manual picking of a number of particles and the generation of templates (e.g., by projections from a 3D reference map). What we need is just the geometric information of the templates.

Specifically, we assume that the template circle has a fixed radius Rd and the template rectangle has a fixed width Wd and a fixed length Ln . As an example, a small portion of a KLH micrograph, shown in Fig. 2A, will be considered in this section.

3.1. Pre-computations

Edges are first computed using a Canny edge detector (Canny, 1986). To improve the robustness to the noise, an edge-cleaning operation is applied to the detected edge map, where all small connected components of edges are removed. In other words, a connected component of edges (based on an 8-neighbor connectivity) are removed if the total number of edge pixels in that component is less than a given threshold. As seen in Fig. 2B, the most important features are roughly kept in the cleaned edge map. Note that the “features” here are the geometric features that depend only on the shapes (or contours) of the particles. The edges inside the shapes are considered as “noise,” although they may correspond to some internal structures from the biological point of view.

It has been pointed out that distance transform is very useful for boundary-based template matching (Gavrila, 1998; Huttenlocher et al., 1993; Rucklidge, 1995). The essential idea is to calculate the distance transform of the edge map of the target image. The average distance value calculated along the template contour at a certain location in the target image gives the goodness of the matching between the target image and the template at that location. Obviously a smaller average distance value means better matching. We also exploit this idea in our method. In addition, as we will see in the following, the distance transform is used in our approach for two additional reasons. First, it is used for the initial location of the circles in the circle detection. Second, it is used for the center and/or orientation refinement in both circle detection and rectangle detection. The distance transform of Fig. 2B is shown in Fig. 2C.

The Voronoi diagram is computed in order to estimate the initial locations and orientations of the rectangular particles. More details on how to do this will be given in the following. Fig. 2D shows the Voronoi diagram of the edge map. Both the distance transform and Voronoi diagram are computed using the list-processing approach (Guan and Ma, 1998).

3.2. Circle detection

3.2.1. Initial circle detection

From the distance transform of the edge map (shown in Fig. 2C), we can see that most circular particles have a local maximum near their centers in the distance map. Therefore, we begin circle detection by detecting the

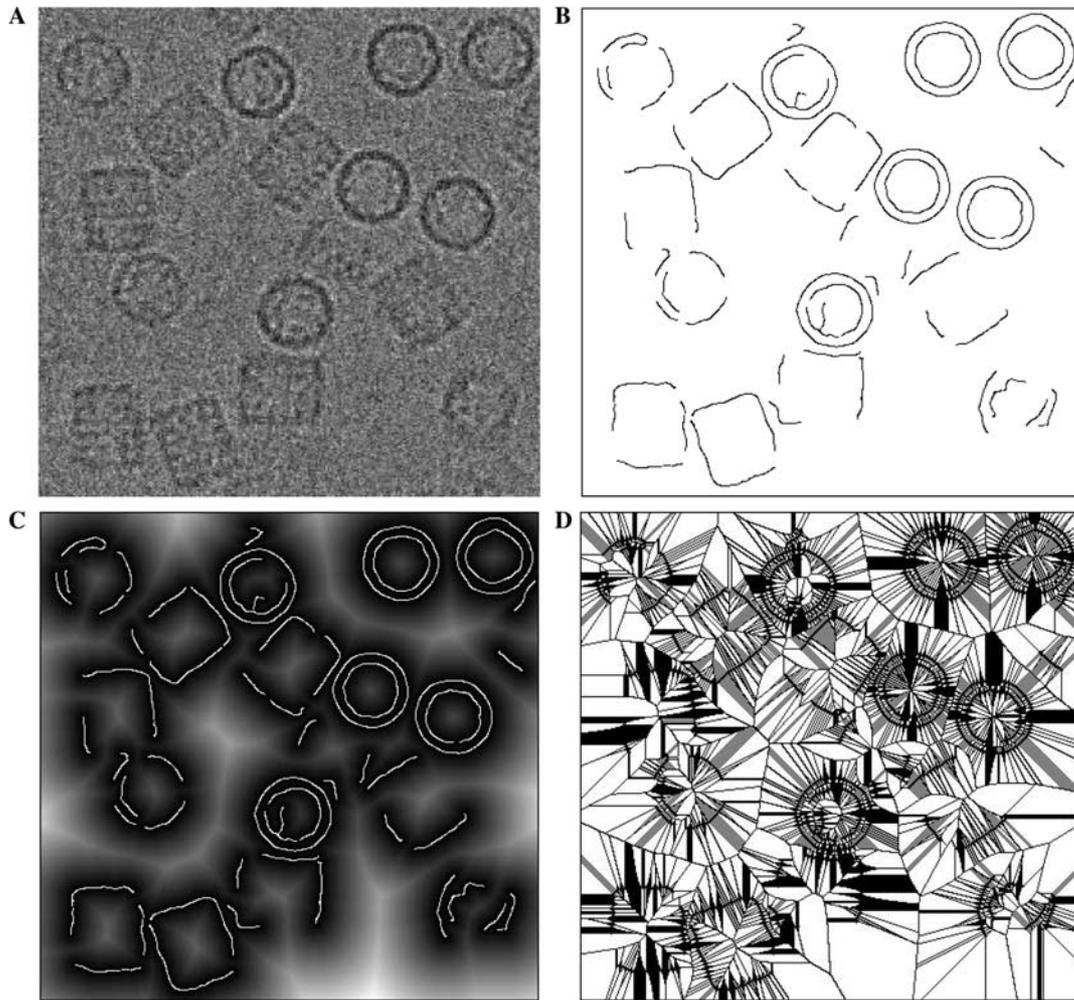


Fig. 2. Illustration of the edges, distance transform, and Voronoi diagram. The edge map is obtained by Canny edge detector (Canny, 1986) followed by edge cleaning (we remove short edges with length below 20 pixels). The distance transform and Voronoi diagram are calculated by list-processing approach (Guan and Ma, 1998). (A) Original image, (B) edge map, (C) distance transform, and (D) Voronoi diagram.

local maxima of the distance map. Our goal here is to find all possible circles, although local maxima in the distance map do not always mean the center of a circle (for example, the rectangles also have local maxima near their centers). The local maximum is defined in our algorithm within a 3×3 local window, which means that a pixel is said to be a local maximum if its distance value is no less than the distance values of its surrounding (eight) neighbors. Besides this, the distance values at (valid) local maxima must be less than the radius Rd of the template circle and greater than a small fixed value (say, 5). Fig. 3A shows the initial circles that we detect using these criteria. Remember that the reason we see so many “local maxima” in Fig. 3A is that the edges obtained by Canny detector do not show perfect circles or rectangles due to the “noise” seen in the original images.

3.2.2. Circle–circle competition

From the initial circle detection, we see too many circles that do not correspond to the true circular par-

ticles. Hence we need to delete most of them. One of the criteria for deleting false circles is the assumption in the *single particle reconstruction* that all particles should be isolated. This assumption immediately gives the following rule for deleting false circles: if two circles intersect, then we must delete one of them. It is quite easy to determine whether two circles intersect or not. But the remaining question is how to choose the “winner” from two intersecting circles. As we said before, the average distance value along a template contour indicates the goodness of the matching. The average distance value of a circle with center C is defined by

$$Adv(C) = \frac{1}{n} \sum_{|d(P,C)-Rd| \leq 1} DT(P), \quad (1)$$

where n is the number of pixels satisfying $|d(P,C)-Rd| \leq 1$ and $d(*,*)$ is Euclidean distance between two pixels. $DT(P)$ is the distance value at P in the distance map. Therefore, the “winner” of two intersecting circles should be the one with smaller average distance value.

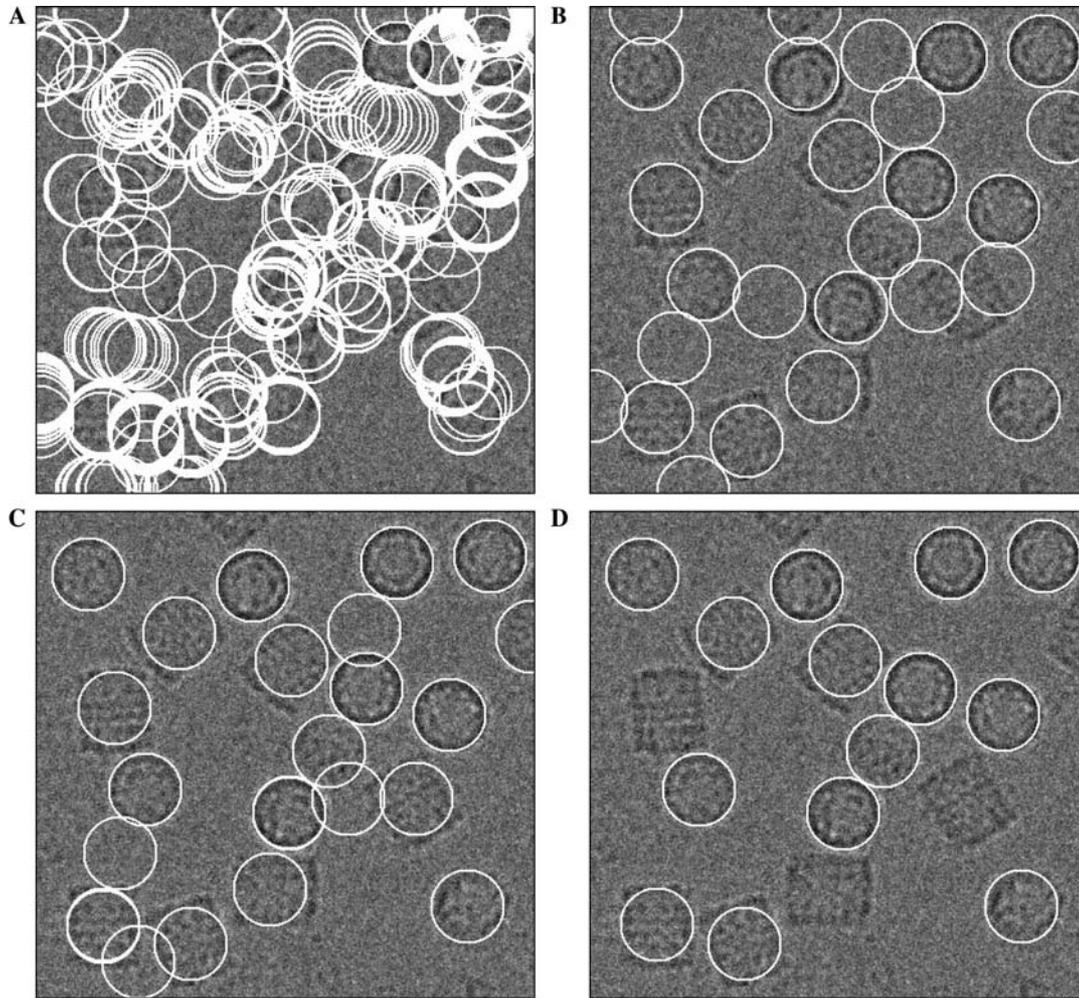


Fig. 3. Illustration of the four steps seen in circle detection. (A) Initial circles, (B) after *circle-circle* competition, (C) after center refinement, and (D) after further refinement.

By this way, we can delete most of the false circles seen in the initial circle detection. The remaining circles are shown in Fig. 3B, from which we can see that the number of circles is largely reduced. But we still see two problems. One is that several circles are more or less away from their true centers due to the noise in the edge map. The other is that some rectangles are wrongly recognized as circles. Both problems will be addressed in the following. Although we said above that all circles must be completely isolated, in Fig. 3B we allow the circles to partially intersect with each other. The reason is that, at this moment, some circles are not well centered. If all survival circles were required to be completely isolated, then some true circles that are not well centered might be wrongly deleted.

3.2.3. Center refinement

As we said before, some circles seen in Fig. 3B do not have correct centers. It is necessary to refine the centers of those circles. How do we do that? Recall the definition of distance transform that we give in Section 2.

Before we calculate the distance transform, we need to identify the nearest neighbor of every pixel in the image domain. Therefore, given a circle with center C , we can define a force on that circle as follows

$$\vec{F}(C) = \frac{1}{n} \sum_{|d(P,C)-Rd| \leq 1} \overrightarrow{P, NN(P)}, \quad (2)$$

where n and $d(*,*)$ are defined the same as in Eq. (1). $NN(P)$ is the nearest neighbor of P . $\overrightarrow{P, NN(P)}$ is the vector from P to $NN(P)$. The force $\vec{F}(C)$ indicates how far the circle should be moving to the correct center. This is an iterative process. Once the circle moves to a new position C' , the new force $\vec{F}(C')$ is calculated and the circle keeps moving in this way until $|\vec{F}(C')|$ is smaller than a given value. Generally two or three iterations should be enough for a circle to reach its correct center. Fig. 3C shows the result after center refinement. An interesting observation is that some circles may move to the same location such that the total number of circles is reduced.

3.2.4. Further refinement

From Fig. 3C we can see that some circles intersect again due to the center refinement. So we need to run the circle–circle competition again. But this time we allow no intersection at all between any two circles. In addition, we delete those circles that have too large average distance value. The final result for the circle detection is shown in Fig. 3D. Remember that we still have the problem that some rectangles are wrongly recognized as circles. This problem will be easily resolved after the rectangle detection described below.

3.3. Rectangle detection

The rectangle detection is performed separately from the circle detection. It also includes four steps: *initial rectangle detection*, *rectangle–rectangle competition*, *center/orientation refinement*, and *further refinement*. We will explain them step-by-step, but special treatment will be given to the first step, the most different part from the circle detection.

3.3.1. Initial rectangle detection

As we saw before, the local maxima of the distance map indicate not only the centers of the circles but the centers of the rectangles as well. Hence, a simple way to detect the initial rectangles is again by detecting the local maxima of the distance map. However, since every rectangle has its center as well as its own orientation, it is too much time-consuming for us to try every possible orientation at each local maximum for the best-matched rectangle. Therefore, we need to seek for other efficient ways for initial rectangle detection. In the following we will consider a method based on Voronoi diagram of the edges (see Fig. 2D). This method begins with *corner detection* followed by conversion from corners to centers/orientations.

A corner of a rectangle is illustrated in Fig. 4A, where the edges (dark dots) are discrete points forming a right angle at *C*. We construct the Voronoi diagram (dashed lines in Fig. 4A) of those edge points and check every point *A* on the Voronoi diagram. If *A* happens to be on the “bisector” of the corner as shown in Fig. 4A, then *A*, together with its two nearest neighbors *B* and *D*, should roughly form a right angle (that is, $\angle BAD \approx 90^\circ$). In addition, the directions of the edges at *B* and *D* should be roughly perpendicular to the vector from *A* to *B* (denoted by $\vec{A,B}$) and the vector from *A* to *D* (denoted by $\vec{A,D}$), respectively. In this case, we can compute the corner *C* by $C = B + D - A$. On the other hand, if *A* is not on the “bisector” of the corner, then *A* and its nearest neighbors do not satisfy the above conditions. Therefore, we have the following algorithm to detect all the corners in an edge map based on Voronoi diagram.

Algorithm for detecting corners:

1. Check every point *A* on the Voronoi diagram (VD).

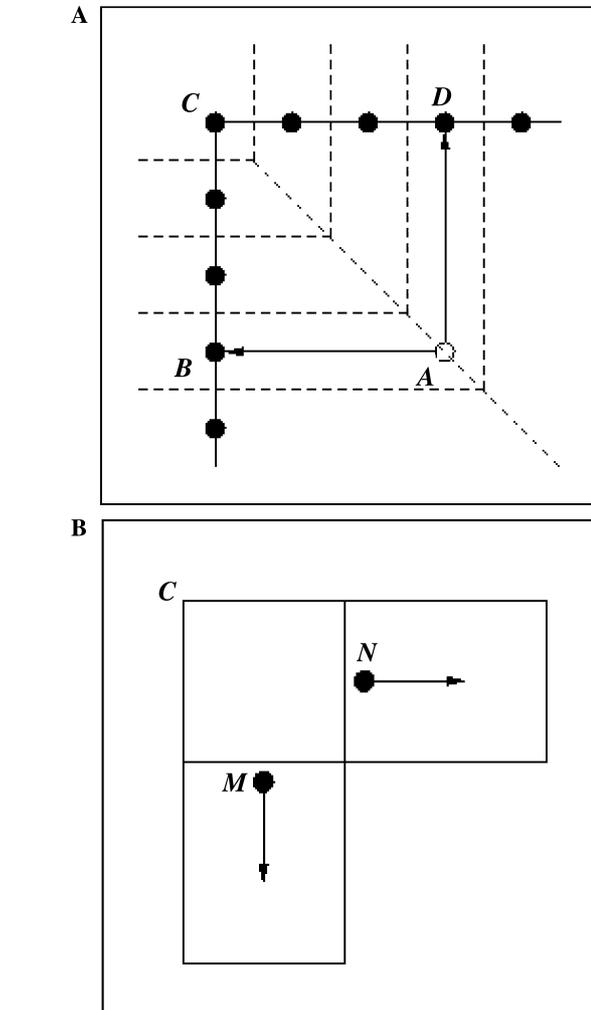


Fig. 4. Illustration of corner detection by Voronoi diagram and the two cases for converting corners to centers. (A) Corner detection and (B) from corner to center.

2. Find its nearest neighbors *B* and *D*. This is readily available after we compute VD (Guan and Ma, 1998).
3. If $\angle BAD \approx 90^\circ$ and both vectors, $\vec{A,B}$ and $\vec{A,D}$, are roughly perpendicular to the edges at *B* and *D*, respectively, then mark *C* ($C = B + D - A$) as a corner.
4. Repeat (1)–(3) until all points on the VD are checked.

It is worth noting that many corners detected by the above algorithm often correspond to the same corner although they may stay a little bit away from the true corner due to the noisy edges. The above method for corner detection is quite robust to noise because, even if part of the edges around a corner are damaged, other undamaged edges can still contribute to the detection of that corner.

After we detect the corners, we then need to convert each of the corners into the center and orientation that uniquely represent each of the rectangles (remember that we assume the width and length of all the rectangles are

fixed). Given a corner, we actually have two cases, as shown in Fig. 4B, which correspond to two possible centers/orientations. Hence we need to choose one of them. Again, the “winner” is chosen by calculating the average distance value on each of these two rectangles and the one with smaller average distance value is the “winner.” It is worth noting that every rectangle has four corners and thus a rectangle can be detected from any one of its four corners even if the others are damaged. Fig. 5A shows the result of the initial rectangles detected by our method. Similar to the initial circle detection, there are many false rectangles due to the noisy edges detected by Canny edge detector. Therefore, we need to let all the initial rectangles compete with each other as described in the following. Compared to the simple approach that searches for the local maxima of the distance map and checks every possible orientation to detect the initial rectangles, the method described above tremendously reduces the number of the candidate rectangles for further processing and hence improves the performance.

3.3.2. Rectangle–rectangle competition

The competition between rectangles is similar to what we saw for circle–circle competition. However, determining whether two rectangles intersect or not is more complicated than the circle–circle intersection. In addition, the average distance value for each rectangle is now computed along the rectangular contour.

3.3.3. Center/orientation refinement

The center refinement of the remaining rectangles is similar to the center refinement of circles seen above. It is possible that the orientations of the rectangles should also be refined. But our experience shows that the orientation refinement (and sometimes even the center refinement) is often unnecessary since the results after the first two steps always give accurate orientation and/or center for each detected rectangle.

3.3.4. Further refinement

Similar to the circle detection, we need to run the rectangle–rectangle competition again for those rectan-

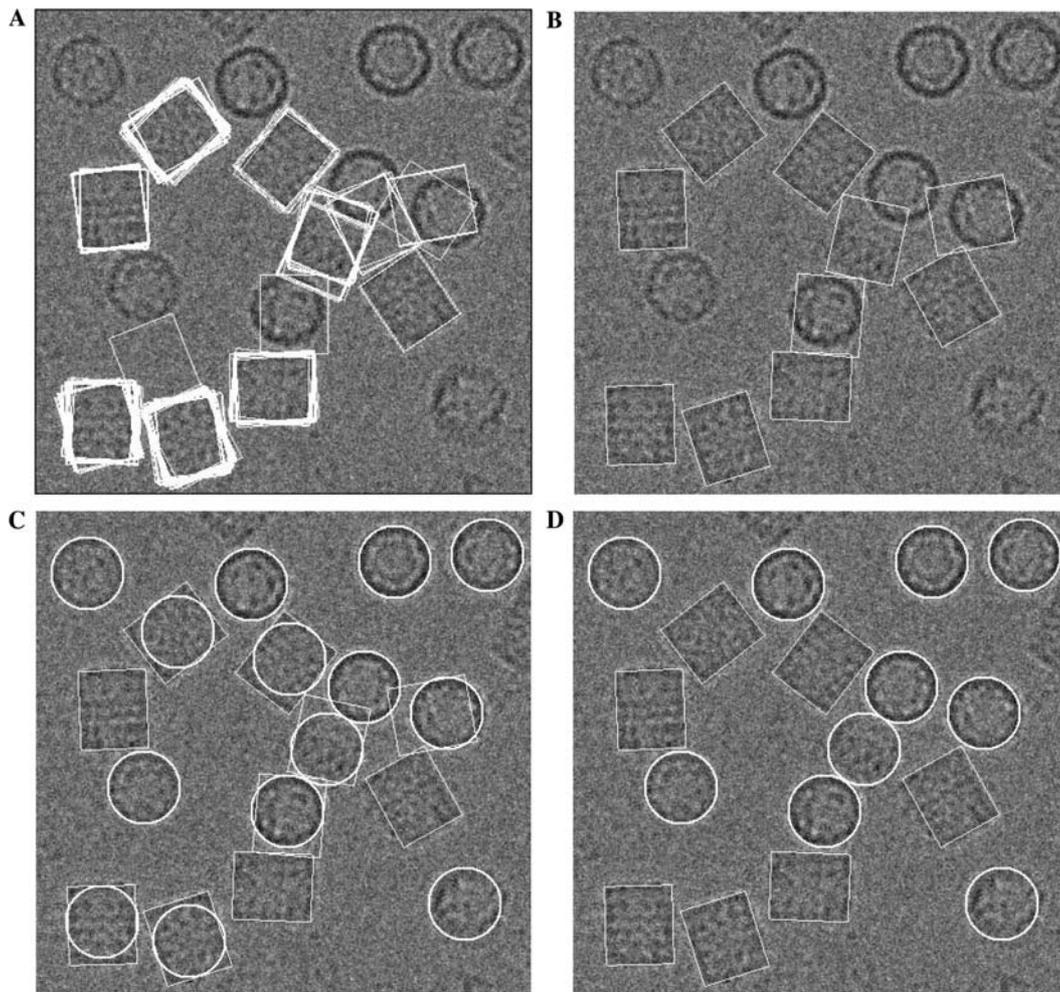


Fig. 5. Illustration of rectangle detection and the circle–rectangle competition. (A) Initial rectangles, (B) refined rectangles, (C) circle–rectangle combined, and (D) final results.

gles that intersect after the above steps. Again, it is useful to delete the rectangles that have too large average distance value. The final result of the rectangle detection is shown in Fig. 5B, from which we can see that there are still a small number of false rectangles. Remember that in circle detection we have an unsolved problem that some rectangles are wrongly recognized as circles. Now it is time for us to solve both problems together by *circle–rectangle competition* as described below.

3.4. Circle–rectangle competition

We first put together the detected circles and rectangles as shown in Fig. 5C. Then we check each of the circles and see if it intersects with any of the surrounding rectangles. If it does, then the circle–rectangle competition is applied. The circle–rectangle competition is similar to the circle–circle competition or rectangle–rectangle competition seen before. The “winner” is also chosen to be the one that has the smaller average distance value. The final result is shown in Fig. 5D.

4. Results

In this section we demonstrate the performance of our particle detection algorithm on several examples of Cryo-EM images, known as *Keyhole Limpet Hemocyanin* (KLH) (Zhu et al., 2002). The original datasets, consisting of 82 defocus pairs, were kindly provided by National Resource for Automated Molecular Microscopy (NRAMM). They were collected at the magnification of 66 000 \times using Philips CM200 transmission electron microscope equipped with a 2048 \times 2048 CCD Tietz camera (Carragher et al., 2000; Potter et al., 1999) (also see the bakeoff paper in this issue). The images we work on in this paper were acquired at far-from-focus conditions ($-2\ \mu\text{m}$).

The original image (2048 \times 2048 pixels) is first down sampled to a smaller image (1024 \times 1024 pixels). The edges are detected using Canny edge detector with the following parameters: $tlow = 0.9$, $thigh = 0.8$, $\sigma = 4$. Here both $tlow$ and $thigh$ are in the range of [0.0–1.0]. These two parameters, however, must be converted into the “true” thresholds before they can be applied to the edge magnitude maps of the images. Two thresholds in Canny method are set as follows (Heath, 1996). We first count the number (denoted by n) of pixels that pass the non-maximal suppression. The higher threshold is set as the magnitude value of the pixel whose magnitude ranks the $\lceil n \times (1 - thigh) + 0.5 \rceil$ th among all the pixels that pass the non-maximal suppression. The lower threshold is simply set as the higher threshold multiplied by $tlow$. We clean the edge maps by removing all small connected components of edges. In other words, a connected

component of edges (based on an 8-neighbor connectivity) are removed if the total number of edge pixels in that component is less than 20. The radius of the template circle is 35 pixels. The width and length of the template rectangle are 65 and 80 pixels, respectively.

The false negatives/positives for each example, evaluated against visual detection, can be found from Figs. 6–8. From these examples we can see that our approach gives accurate centers and/or orientations for most of the detected particles. The false positive rate (FPR) is defined as the number of false positives divided by the total number of particles in the test set. The false negative rate (FNR) is defined as the number of false negatives divided by the total number of particles in the truth set. The FPR and FNR for the side views (the rectangular particles) of the entire far-from-focus KLH datasets are 24.7 and 8.3%, respectively, evaluated against Mouche’s criteria (a manual method for particle picking, also see the bakeoff paper in this issue). Note that the false positive rate is a little bit high due to the fact that our method recognized many “long” rectangular particles (about half of the total false positives), which, however, are considered as false positives in Mouche’s criteria. In case of the top views (the circular particles), the total number of recognized particles by our method is 1351, among which 184 particles are false positives. In addition, 31 true particles are missing in our results. Therefore, the FPR and FNR for the top views are 13.6 and 2.6%, respectively, based on the above definitions of FPR and FNR. Here both the false positives and the false negatives are evaluated against visual detection.

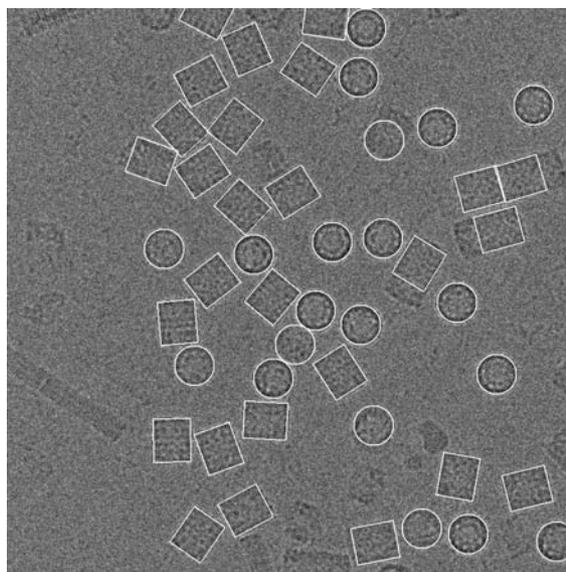


Fig. 6. Example I: *Circles*: 20 particles detected; 1 false positive; 0 false negative. *Rectangles*: 27 particles detected; 0 false positive; 1 false negative.

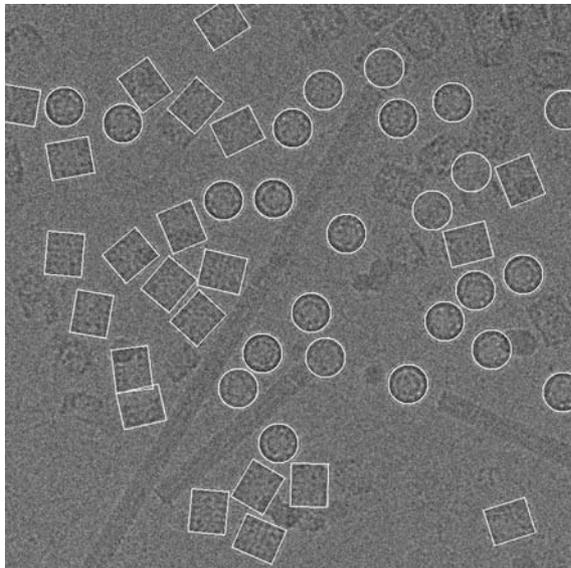


Fig. 7. Example II: *Circles*: 24 particles detected; 2 false positives; 1 false negative. *Rectangles*: 22 particles detected; 1 false positive; 7 false negatives.

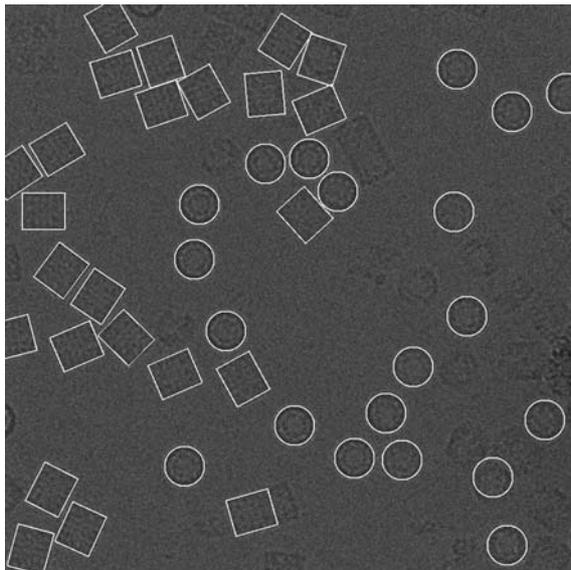


Fig. 8. Example III (with low contrast): *Circles*: 20 particles detected; 0 false positive; 4 false negatives. *Rectangles*: 24 particles detected; 0 false positive; 3 false negatives.

We tested our method on SGI Onyx2 with single processor (400 MHz MIPS R12000) and the total computational time is about 20 s for each image with a size of 1024×1024 pixels. About 18%, 2%, and 6% of the total time are used for edge detection, edge cleaning, and computations of DT & VD, respectively. The particle detection (including circle detection, rectangle detection, and circle–rectangle competition) takes about 74% of the total time, which is roughly four times of that for Canny edge detection. The source code is available for

free download under the GNU public license, from <http://www.ices.utexas.edu/CCV/software>.

5. Discussions

5.1. Using multiple sets of parameters

The edges detected by Canny method are sensitive to the selection of the three parameters (*thigh*, *tlow*, and *sigma*). Although a certain level of noise in the edge map does not affect the performance of our feature-based particle picking algorithm, the edges with completely wrong parameters may exclude some geometric features and hence some particles may be misrecognized. One way to reduce the sensitivity-to-parameter is to utilize multiple sets of parameters and the results from each of them are then combined for better results. We demonstrate this strategy in Fig. 9, where two sets of parameters are used in Canny method. It is obvious that each set of parameters alone cannot correctly detect particles (one rectangular particle is wrongly recognized as circular particle). However, these two sets of parameters are complementary to each other in the sense that wrong particles in one case are correctly detected in the other. The combined particles shown in Fig. 9 give correct result. Two sets of results are combined by inter-competition, which is similar to the competitions we introduced in Section 3. Each of the particles from one result competes with each of the particles from the other one if both particles intersect. The “winner” is again the one with smaller average distance value along its own shape contour. Fig. 10 shows particles detected using three sets of parameters. One can compare it with Fig. 7 and see the improvements. A disadvantage of using multiple sets of parameters is that the computational time may increase by several times depending on how many sets of parameters are considered.

5.2. Anisotropic filtering

Edges play a crucial role in our particle detection approach but they are often corrupted by noise commonly seen in the electron micrographs. Therefore, it is expected that noise reduction would be in great demand as a pre-processing step. Traditional image filters include Gaussian filtering, median filtering, and frequency domain filtering (Gonzalez and Woods, 1992). Most of recent research has been devoted to anisotropic filters that smooth out the noise without blurring the geometrical details such as edges and corners. Several categories of anisotropic filters have been proposed in the area of image processing. Bilateral Filtering (Durand and Dorsey, 2002; Elad, 2002; Tomasi and Manduchi, 1998) is a straightforward extension of Gaussian filtering by simply multiplying an additional term in the

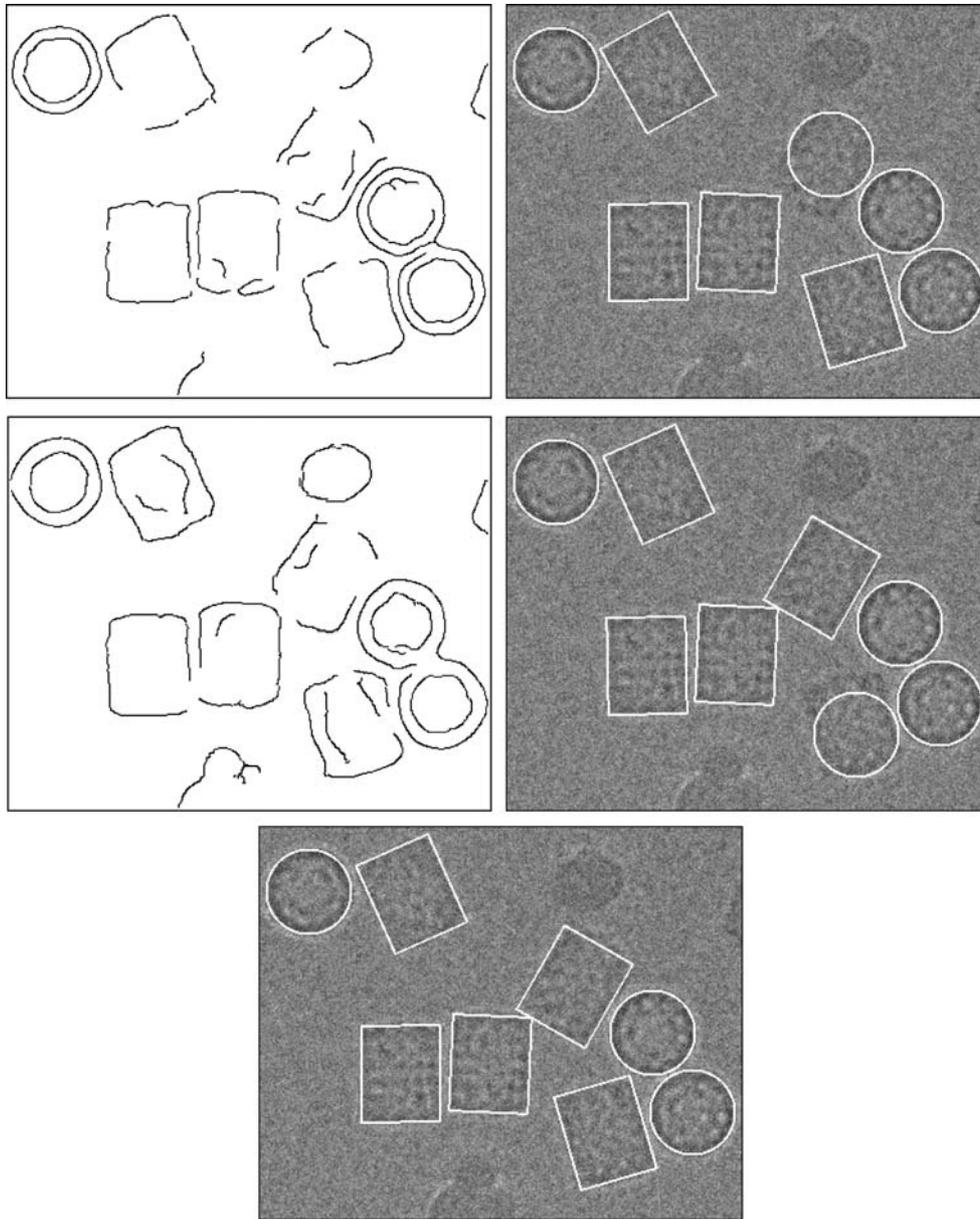


Fig. 9. Particle detection can be enhanced by multiple sets of parameters in Canny edge detection. Top row: results (edges and particles detected) using one set of parameters ($\sigma = 4$; $t_{low} = 0.9$; $t_{high} = 0.8$). Middle row: results using another set of parameters ($\sigma = 6$; $t_{low} = 0.7$; $t_{high} = 0.8$). In both cases, there is one rectangle that is wrongly recognized as circle. Bottom row: combined results using the above two sets of parameters. All edges are cleaned by removing short edges with length below 20 pixels.

weighting function. A partial differential equation (PDE) based technique, known as anisotropic heat diffusion, has also been studied (Perona and Malik, 1990; Weickert, 1998) and a fundamental relationship was discussed between this technique and the bilateral filtering (Barash, 2002). Another popular technique for anisotropic filtering is by wavelet transformation (Xu et al., 1994). The basic idea is to identify and zero out wavelet coefficients of a signal that likely correspond to image noise. Finally, the development of nonlinear median-based filters in recent years has also resulted in

promising results. One of those filters is called mean-median (MEM) filter (Hamza et al., 1999).

We have tested several methods for noise reduction, including Perona–Malik model (Perona and Malik, 1990), bilateral filtering (Tomasi and Manduchi, 1998), and anisotropic diffusion (Bajaj et al., 2003). Fig. 11 shows the difference between the results on the original image and the results enhanced by bilateral filter (Tomasi and Manduchi, 1998) and by anisotropic diffusion (Bajaj et al., 2003). The particle detection on the filtered data gives more accurate results due to the

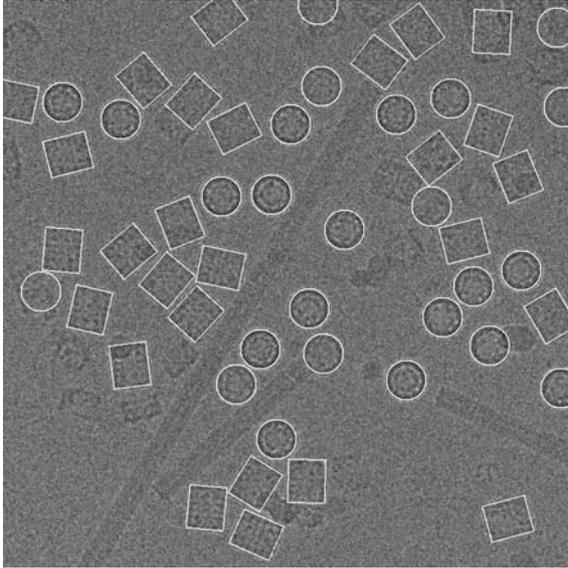


Fig. 10. Particle detection is enhanced by three sets of parameters in Canny edge detection. Set I: $\sigma = 4$; $tlow = 0.9$; $thigh = 0.8$. Set II: $\sigma = 4$; $tlow = 0.7$; $thigh = 0.8$. Set III: $\sigma = 4$; $tlow = 0.85$; $thigh = 0.75$. Edges are cleaned by removing short edges with length below 20 pixels.

anisotropic filtering that smoothes the noise while preserving sharp edges.

5.3. Particles with other shapes

The particles we have been working on so far are restricted to circular and rectangular shapes. However, a similar pipeline also works for particles with other types of shapes if some kind of geometric features can be derived from the shapes (e.g., icosahedral viruses). The general pipeline is described as follows:

1. From the given electron micrographs, determine the template shapes (e.g., radius for circle, width/length for rectangle, and so on).
2. Smooth original images using anisotropic filtering.
3. Apply Canny edge detection followed by edge cleaning.
4. Compute distance transform and Voronoi diagram of the edges.
5. Detect geometric features from the distance transform and/or Voronoi diagram. In case of circular/rectangular particles, the geometric features are centers/corners, respectively.
6. From the geometric features, guess the initial particles (including the centers and, if applicable, the orientations). Every candidate particle is attached with a “survival” factor, which is defined by the inverse of the average distance value along the particle’s contour.
7. Any two candidate particles, if they intersect, should compete with each other. The “winner” in a competi-

tion is the one that has a larger “survival” factor than its competitor. As described in Section 3, particle refinement may be necessary to improve the detection accuracy. In addition, multiple sets of parameters for Canny edge detection can be helpful.

8. The output is the final “winners” of all the particle competitions.

Detecting particles of arbitrary shapes is much more challenging. However, particle detection in the single particle analysis does not require accurate boundary segmentation. Edge detection, coupled with anisotropic filtering and edge-linking based on distance transform, may still be a feasible way in such cases.

6. Conclusions and future work

In this paper we described an automatic method for particle detection from electron micrographs. Our method is based on shape matching between the edge map calculated from the original/filtered image and the template shapes of the true particles (what we call “virtual” templates). Unlike the conventional template-matching method that is based on cross-correlation of intensities, our method only needs the geometric features of the particles and hence avoids the manual picking of a number of particles and the tedious generation of the reference images (e.g., by projections from a 3D reference map). Multiple views of the particles are allowed by simply applying different geometric features to each view (e.g., the circular/rectangular examples as we considered in Section 3). In addition, with the help of the distance transform and Voronoi diagram, we can quickly identify the geometric features from the edges and make a good guess on the initial positions/orientations of the candidate particles. Good initialization is a crucial part of our algorithm in order to reduce the computational time as well as improve the detection accuracy. Since our method is based on boundaries (or edges), which are said to be the most important features of an image, other irrelevant details (or noise) have a very small influence on the detection of a particle from the original image. This “weighted” matching approach differs a lot from the conventional template-matching methods where all intensities around a pixel under consideration evenly influence the matching score.

Part of our future work is *flexible shape matching*, meaning that, in particular, the radius of the template circle and the width/length of the template rectangle could be flexible. We expect that this could further reduce the sensitivity of our particle detection algorithm to the determination of template shapes. Another interesting work is that the geometric features of each particle may also help to enhance the *particle alignment* and *particle classification* in the single particle analysis.

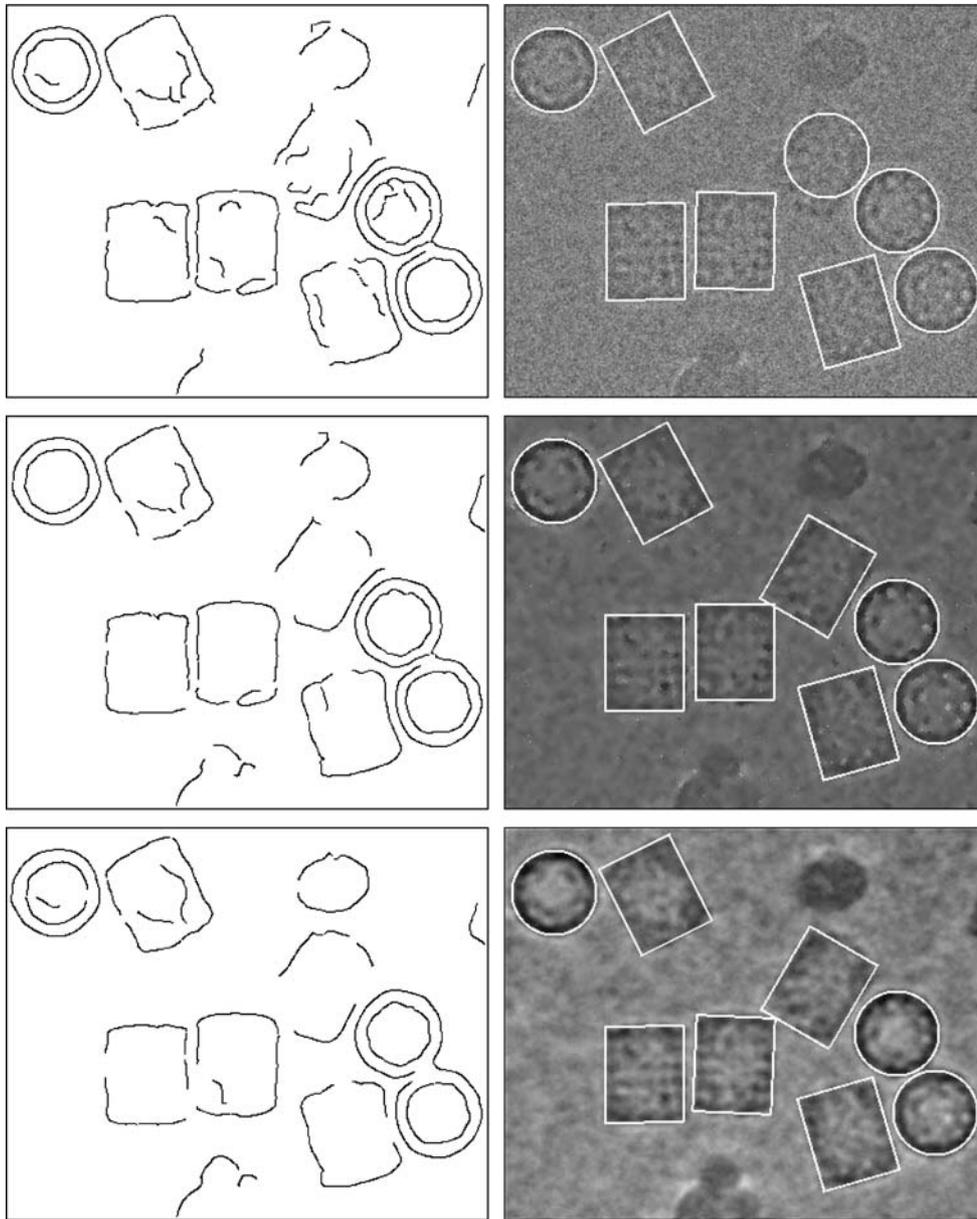


Fig. 11. Particle detection can be enhanced by anisotropic filtering. Top row: results (edges and particles detected) without pre-filtering. Middle row: results (edges and particles) enhanced by bilateral filtering. Bottom row: results (edges and particles) enhanced by anisotropic diffusion filtering. In all three cases, the parameters for edge detection are chosen as: $\sigma = 4$; $t_{low} = 0.8$; $t_{high} = 0.8$. All edges are cleaned by removing short edges with length below 20 pixels.

Acknowledgments

This research is supported in part by NSF Grants ACI-9982297, CCR-9988357, ITR:ACI-0220307, ITR: EIA-0325550, and from Grant UCSD 1018140 as part of NSF-NPACI, Interactive Environments Thrust. Thanks are also due to Qiu Wu of CCV for implementations of the bilateral and anisotropic diffusion filtering, and to Yuanxin Zhu, Bridget Carragher at The Scripps Institute for providing the KLH datasets. Some of the data used in the work presented here was provided by the National Resource for Automated Molecular

Microscopy (NRAMM) which is supported by the National Institutes of Health through the National Center for Research Resources' P41 program (RR17573).

References

- Andrews, D.W., Yu, A.H.C., Ottensmeyer, F.P., 1986. Automatic selection of molecular images from dark field electron micrographs. *Ultramicroscopy* 19, 1–14.
- Bajaj, C., Wu, Q., Xu, G., 2003. Level-set based volumetric anisotropic diffusion for 3D image denoising. ICES Technical Report #301, University of Texas at Austin.

- Baker, T.S., Olson, N.H., Fuller, S.D., 1999. Adding the third dimension to virus life cycles: three-dimensional reconstruction of icosahedral viruses from cryo-electron micrographs. *Microbiol. Mol. Biol. Rev.* 63 (4), 862–922.
- Barash, D., 2002. A fundamental relationship between bilateral filtering, adaptive smoothing and the nonlinear diffusion equation. *IEEE Trans. Pattern Anal. Machine Intelligence* 24 (6), 844–847.
- Canny, J., 1986. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Machine Intelligence* 8, 679–698.
- Carragher, B. et al., 2000. Leginon: an automated system for acquisition of images from vitreous ice specimens. *J. Struct. Biol.* 132, 33–45.
- Durand, F., Dorsey, J., 2002. Fast bilateral filtering for the display of high-dynamic-range images. In: *Proceedings of the ACM Conference on Computer Graphics (SIGGRAPH)*, pp. 257–266.
- Elad, M., 2002. On the bilateral filter and ways to improve it. *IEEE Trans. Image Processing* 11 (10), 1141–1151.
- Frank, J., 1996. *Three-dimensional electron microscope of macromolecular assemblies*. Academic Press, San Diego.
- Frank, J., Wagenknecht, T., 1984. Automatic selection of molecular images from electron micrographs. *Ultramicroscopy* 12, 169–176.
- Gavrila, D.M., 1998. Multi-feature hierarchical template matching using distance transforms. In: *Proceedings of the 1998 International Conference on Pattern Recognition*, pp. 439–444.
- Gonzalez, R.C., Woods, R.E., 1992. *Digital Image Processing*. Addison-Wesley.
- Guan, W., Ma, S., 1998. A list-processing approach to compute Voronoi diagram and Euclidean distance transform. *IEEE Trans. Pattern Anal. Machine Intelligence* 20 (7), 757–761.
- Hamza, A.B. et al., 1999. Removing noise and preserving details with relaxed median filters. *J. Math. Imaging and Vis.* 11 (2), 161–177.
- Harauz, G., Fong-Lochovsky, A., 1989. Automatic selection of macromolecules from electron micrographs by component labeling and symbolic processing. *Ultramicroscopy* 31 (4), 333–344.
- Heath, M.D., 1996. *A robust visual method for assessing the relative performance of edge-detection algorithms* (Master's Thesis). Department of Computer Science, University of South Florida, Tampa, Florida, p. 145.
- Heel, M.v., 1982. Detection of objects in quantum-noise-limited images. *Ultramicroscopy* 7, 331–342.
- Heel, M.v. et al., 2000. Single-particle electron cryo-microscopy: towards atomic resolution. *Q. Rev. Biophys.* 33 (4), 307–369.
- Henderson, R., 1995. The potential and limitations of neutrons, electrons and X-rays for atomic resolution microscopy of unstained biological macromolecules. *Quart. Rev. Biophys.* 28, 171–193.
- Huttenlocher, D.P., Klanderman, G.A., Rucklidge, W.J., 1993. Comparing images using the Hausdorff distance. *IEEE Trans. Pattern Anal. Machine Intelligence* 15 (9), 850–863.
- Kivioja, T. et al., 2000. Local average intensity-based method for identifying spherical particles in electron micrographs. *J. Struct. Biol.* 131 (2), 126–134.
- Lata, K.R., Renczek, P., Frank, J., 1995. Automatic particle picking from electronic micrographs. *Ultramicroscopy* 58, 381–391.
- Ludtke, S.J., Baldwin, P.R., Chiu, W., 1999. EMAN: semiautomated software for high-resolution single-particle reconstructions. *J. Struct. Biol.* 128 (1), 82–97.
- Martin, I.M.B. et al., 1997. Identification of spherical virus particles in digitized images of entire electron micrographs. *J. Struct. Biol.* 120 (2), 146–157.
- Nicholson, W.V., Glaeser, R.M., 2001. Review: automatic particle detection in electron microscopy. *J. Struct. Biol.* 133 (2-3), 90–101.
- Ogura, T., Sato, C., 2001. An automatic particle pickup method using a neural network applicable to low-contrast electron micrographs. *J. Struct. Biol.* 136, 227–238.
- Perona, P., Malik, J., 1990. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Machine Intelligence* 12 (7), 629–639.
- Potter, C.S. et al., 1999. Leginon: A system for fully automated acquisition of 1000 micrographs a day. *Ultramicroscopy* 77, 153–161.
- Roseman, A.M., 2003. Particle finding in electron micrographs using a fast local correlation algorithm. *Ultramicroscopy* 94, 225–236.
- Rucklidge, W.J., 1995. Locating objects using the Hausdorff distance. In: *Proceedings of the 1995 IEEE Conference on Computer Vision*, pp. 457–464.
- Saad, A., Chiu, W., Thuman-Commike, P., 1998. Multiresolution approach to automatic detection of spherical particles from electron cryomicroscopy images. In: *Proceedings of the IEEE International Conference on Image Processing*, pp. 8–10.
- Stoschek, A., Hegerl, R., 1997. Automated detection of macromolecules from electron micrographs using advanced filter techniques. *J. Microsc. (Oxford)* 185, 76–94.
- Thuman-Commike, P., Chiu, W., 1995. Automatic detection of spherical particles from spot-scan electron cryomicroscopy images. *J. Microsc. Soc. Am.* 1, 191–201.
- Thuman-Commike, P., Chiu, W., 1996. PTOOL: a software package for the selection of particles from electron cryomicroscopy spot-scan images. *J. Struct. Biol.* 116 (1), 41–47.
- Tomasi, C., Manduchi, R., 1998. Bilateral filtering for gray and color images. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 836–846.
- Weickert, J., 1998. *Anisotropic Diffusion In Image Processing*. ECMI Series, Teubner, Stuttgart, ISBN 3-519-02606-6.
- Xu, Y. et al., 1994. Wavelet Transform Domain Filters: a spatially selective noise filtration technique. *IEEE Trans. Image Processing* 3 (6), 747–758.
- Zhu, Y. et al., 2001. Automated identification of filaments in cryo-electron microscopy images. *J. Struct. Biol.* 135 (3), 302–312.
- Zhu, Y., Carragher, B., Potter, C.S., 2002. Fast detection of generic biological particles in cryo-EM images through efficient Hough transforms. In: *Proceedings of the 2002 IEEE International Symposium for Biomedical Imaging*, pp. 205–208.