Radon-Like Features and their Application to Connectomics

(Article begins on next page)

Radon-Like Features and their Application to Connectomics

Ritwik Kumar[1], Amelio Vázquez-Reina[2], Hanspeter Pfister[1]
[1]School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, USA
[2]Dept. of Computer Science, Tufts University, MA, USA,
{rkkumar, pfister}@seas.harvard.edu, amelio.vazquez-reina@tufts.edu

Abstract

In this paper we present a novel class of so-called Radon-Like features, which allow for aggregation of spatially distributed image statistics into compact feature descriptors. Radon-Like features, which can be efficiently computed, lend themselves for use with both supervised and unsupervised learning methods. Here we describe various instantiations of these features and demonstrate there usefulness in context of neural connectivity analysis, i.e. Connectomics, in electron micrographs. Through various experiments on simulated as well as real data we establish the efficacy of the proposed features in various tasks like cell membrane enhancement, mitochondria segmentation, cell background segmentation, and vesicle cluster detection as compared to various other state-of-the-art techniques.

1. Introduction

Recent breakthroughs in neural imaging has allowed very detailed mapping of the neural connections of the nervous system. Such neuron level maps of brain tissue obtained using Electron Microscopy (EM) are called Connectomes. The identification and analysis of the first complete Connectome, that is, the first full connectivity diagram of the brain of a mammal, is expected to open floodgates to a better understanding of the basic cognitive functions of the human brain and related pathologies [7]. Due to the relatively large number of neurons in the brain (~ 4 million for a mouse), manual annotation of these detailed maps is considered infeasible and automated methods present the only way to obtain meaningful global understanding of these maps. One of the first steps along these lines involves automated segmentation of the constituent structures of the Connectome images, which broadly include – the neuron cell boundaries, the mitochondria, the cell background and the clusters of neurotransmitters (vesicles) (indicated in Fig.1).

The medical imaging and computer vision literature is littered with numerous techniques for automated segmentation of images [9]. Though there are various seminal and landmark pieces of work [10, 4], one of the main reasons
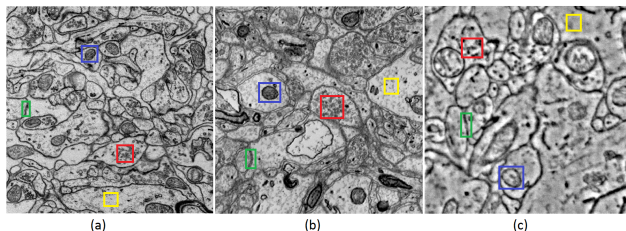

Figure 1. **Connectome Electron Microscopy Images**: The three different kind of images obtained as the staining process in the image acquisition pipeline is changed. Blue boxes show mitochondria, red boxes show vesicles, green boxes show cell boundaries and yellow boxes show cell background.

there are a large number of methods is because they are tailored, and rightly so, to work with a specific kind of images - Photographs, CT Scans, EM images etc. Primarily due to the novel nature of this data, the landscape of the image understanding techniques meant for Connectome data is relatively sparse as compared to the general segmentation literature. Noteworthy work in this field includes techniques focused on cell boundary segmentation [18, 6, 13, 15], mitochondria segmentation [11, 16] and vesicle segmentation [2].

The Connectome mapping process involves precise staining, sectioning and scanning of large volumes of brain tissue (several cubic micrometers in size) with resolutions in the order of a few nanometers per pixel. The unique characteristics of the Connectome image data involve presence of non-trivial amount of scanning noise, staining artifacts - dark blotches, slicing artifacts - grazed vesicles, presence of structures defined by texture (vesicles), structure (neurons) and both (mitochondria). Given this, we have found (also demonstrated later) that a direct off-the-shelf application of existing segmentation and associated methods does not necessarily lead to good results.

Motivated by this nature of the Connectome images, in this paper we propose a novel class of features, called Radon-Like features, that provide a way of aggregating any desired information derived from an image (e.g. image intensities) within structural units (e.g. cell boundaries). These features can then be used with a supervised learn-
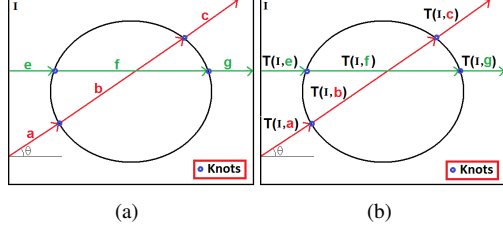
(a)   (b)

Figure 2. **Radon-Like Features**: All the points along the line segments a, b and c on the red line and the segments e, f and g on the green line (part (a)) in the shown image I are assigned value defined by the extraction function **T** (shown in part part (b)).
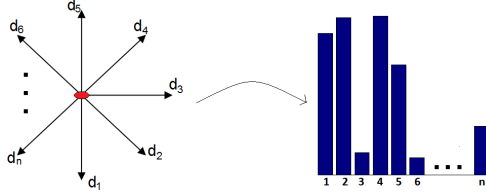

Figure 3. The feature values collected at a pixel (red dot above) along various directions could be used to define a distribution of features (shown on the right).

ing method or simply as an unsupervised way of enhancing structures in an image. We must point out that depending on the staining process, the obtained images vary significantly in appearance. Fig. 1 shows three such images obtained using different staining methods. It can be noted that same structures appear quite different in the three images. This may require supervised segmentation and classification techniques (e.g. [13]) to be retrained for every new acquisition of data.

2. Radon-Like Features

The Radon Transform [1] is an integral transformation which in 2D is defined as

$$\mathbf{R}(m, \tau)[f(x, y)] = \int_{-\infty}^{\infty} f(x, m + \tau y)dx \quad (1)$$

where m and τ define the slope and the intercept of the line, respectively, along which the two dimensional function $f(x, y)$ is integrated. This transform is widely used in tomography e.g. in CT scanning, where the output is, loosely speaking, a Radon Transform image comprising of values obtained using Eq. 1 for some sampling of the slope (m) and the intercept (τ) space. The Radon Back-projection, inverse of the above described transform, is then used to reconstruct the original image

We propose a new class of features, called Radon-Like features, which retains the central idea from the Radon Transform – processing an image (a 2D function, $I(x, y)$) along a line ($\mathbf{l}$, parameterized by t, i.e. $\mathbf{l}(t) = (\mathbf{x}(t), \mathbf{y}(t))$). But instead of collapsing $I(x, y)$ along $\mathbf{l}$ into a scalar value via integration (as in Eq. 1), we distribute some

(a) Input image   (b) Edge map

(c) Features in direction $0°$.   (d) Features in direction $178°$.

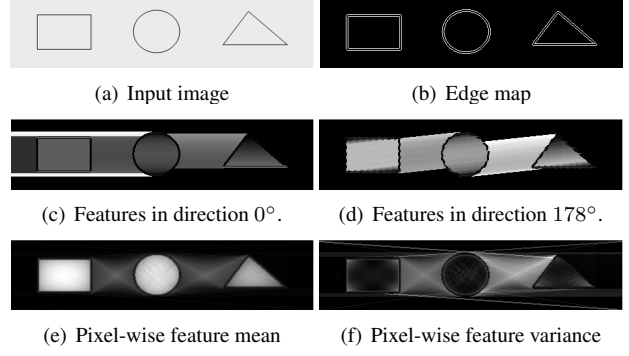(e) Pixel-wise feature mean   (f) Pixel-wise feature variance

Figure 4. Illustration of Radon-Like features when used with the extraction function T_1 (Eq. 3).

desired information derived from $I(x, y)$ among various line segments along $\mathbf{l}$. The line segments are defined by a set of salient points, called *knots*, along $\mathbf{l}$. If the set of knots along $\mathbf{l}$ is given as $(t_1, \dots, t_n)$, value of Radon-Like feature at a point $\mathbf{p}$ along $\mathbf{l}$ between $(\mathbf{x}(t_i), \mathbf{y}(t_i))$ and $(\mathbf{x}(t_{i+1}), \mathbf{y}(t_{i+1}))$ is given by

$$\mathbf{\Psi}(\mathbf{p}, \mathbf{l}, t_i, t_{i+1})[I(x, y)] = T(I, \mathbf{l}(t)), t \in [t_i, t_{i+1}], \quad (2)$$

where T can be any desired function, called the *extraction function*. Note that the information derived along $\mathbf{l}$ need not be the same as $\mathbf{l}$ is traversed in the two opposite directions. In other words, more flexibility can be built into our features if each line also has an associated direction. This can be easily incorporated by associating with each $\mathbf{l}$ an angle $\theta \in (0, 2\pi)$ whose tangent gives its slope. Important point being that θ is considered to be different than $\pi + \theta$.

The output of the above feature detection and description process is yet another image $I'(x, y)$, when the angle θ of $\mathbf{l}$ is fixed and its intercept varies, which is of the same size as the input image $I(x, y)$. This is a significant point of departure from the original Radon Transform where the output in such a case is a 1D function. If the angle, θ, is also varied, we obtain a vector of features at each point or pixel (x, y) in the image.

A simpler version of Radon-Like features is described in Fig. 2 where all the points along a directed line segment in the shown image I are assigned the same value. The red and the green lines in Fig. 2 denote two possible scanning directions. The line segments a, b and c along the red line and the segments e, f and g along the green line are assigned values obtain via the extraction function T. The blue circles represent the chosen knots along the two lines. After the features have been collected, a vector is obtained at each pixel in the image, which, for instance, can be used as a feature distribution (Fig. 3) for a pattern recognition task.

The usefulness of the features described above depends on two critical choices – the knots and the extraction function. For the application at hand, Connectome image analysis, knots derived from the edges in the image provide a

useful guide to the constituent structures of the image. In other words, as a line scans through the input image, its intersections with the image's edge map define the knots and the line segments. Even with the knots fixed in this manner, the choice of the extraction function provides enough flexibility to differentiate among various structures in an image, as we would demonstrate later.

As an example, in Fig. 4 we illustrate Radon-Like features for the following choice of the extraction function:

$$T_1(f, \mathbf{l}(t)) = ||\mathbf{l}(t_{i+1}) - \mathbf{l}(t_i)||_2, t \in [t_i, t_{i+1}]. \quad (3)$$

The extraction function T_1 assigns all the pixels between the knots t_i and t_{i+1} the value equal to the distance between the them. The knots along the line $\mathbf{l}$ were taken to be the points of intersection of $\mathbf{l}$ and the edge map of the image f. For the example in Fig. 4, the intercept parameter of the scan line $\mathbf{l}$ was varied to cover the whole image (**this would always be true in the rest of the paper**) and the angle parameter was varied in steps of $2°$ (from $0°$ to $180°$). Fig. 4(a) and Fig. 4(b) show the input image and its corresponding edge map obtained using Canny filter. Fig. 4(c) and Fig. 4(d) show the features obtained at each pixel when the image is scanned from two different directions (note that for a given scan line $\mathbf{l}$, T_1 is same for both the scan-directions). Once 90 features have been obtained at each pixel, we approximated their distribution by a Gaussian distribution, whose mean and variance is shown in Fig. 4(e) and Fig. 4(f), respectively.

Our Radon-Like features should not be confused with the Fast Ray features of [11]. While they both rely on the casting 2D line segments across the image plane, the Fast Ray features only sample geometric information at each segments's end points, whereas Radon-like features are more general and can both sample or aggregate image statistics (geometric or textural) along a segment.

3. Connectome Image Analysis with Radon-Like Features

In the pervious section we introduced the general class of Radon-Like features. Here we would describe a few specific instantiations of these features which can be useful in Connectome image analysis. We would focus on the four primary structures that are present in Connectome EM images - cell boundaries, mitochondria, cell background and vesicle clusters, as a means of illustrating the above described features. Radon-Like features, like any other features, can be combined with sophisticated machine learning algorithms to desired effect, but the objective of this work is to design features that do most of the job by themselves.

The data that we use in the following experiments is obtained as follows: A volume of brain tissue is first stained with a high-contrast agent. Then an ultra-microtome laser
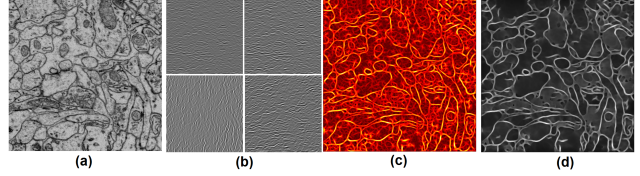

Figure 5. **Cell Boundary Enhancement using Radon-Like features:** Part (a) shows the input and (b) shows the responses of Gaussian-Second-Derivative filters in four directions. Part (c) shows the image obtained by taking the maximum response at each pixel and (d) shows the result of aggregating the information in (c) via Radon-Like features with the extraction function given in Eq. 5.

sections the sample, and the tissue is fed into a high-throughput transmission scanning electron microscope. The result of the imaging process is a stack of 2D, gray-scale (usually 8 bit) images with a resolution in $X - Y$ dimension typically ten times higher than in Z dimension (e.g. 3 to 5 nm in $X - Y$ and 40 nm in Z). A typical 2D section in the obtained stack is $16K \times 16K$ pixels in size. For efficient handling we have diced such sections into smaller 512×512 or 1024×1024 pixels tiles.

3.1. Cell Boundary Enhancement

The primary tasks in Connectome image analysis is segmentation of the data into constituent neuron cells. One of the promising approaches for this involves segmenting the 2D slices of the volume into cell boundaries and then combining them into 3D structures [13, 15]. For this purpose, cell boundary enhancement (and eventual segmentation) has attracted attention in recent years [6, 18]. Focus of the work in this field has been on largely unsupervised diffusion based techniques [17, 12] as well as those based on bulky convolutional networks [5, 6] and graph-cuts methods [18]. In order to emphasize the usefulness of Radon-Like features, we would restrict ourselves to unsupervised methods.

In our scheme for cell boundary enhancement, the knots and the extraction function for Radon-Like features use the following transformation of the input image $I(x, y)$:

$$R(x, y) = \max_{\sigma, \phi} \Delta G(\sigma, \phi) \otimes I(x, y), \quad (4)$$

where σ and ϕ are the scale and the orientation of the boundary enhancing Gaussian-Second-Derivative (GSD) [8] filter $\Delta G(\sigma, \phi)$ ($\otimes$ denotes convolution). The knots for Radon-Like features are defined using an edge map of $R(x, y)$ and the extraction function, T_2, is given as

$$T_2(I, \mathbf{l}(t)) = \frac{\int_{t_i}^{t_{i+1}} R(\mathbf{l}(t))\partial t}{||\mathbf{l}(t_{i+1}) - \mathbf{l}(t_i)||_2}, t \in [t_i, t_{i+1}], \quad (5)$$

where $\mathbf{l}$ is the line along which features are obtained. This extraction function assigns all the pixels between the knots t_i and t_{i+1} along $\mathbf{l}$ the mean value of function R along $\mathbf{l}$ between the same two knots. The transformation $R(x, y)$
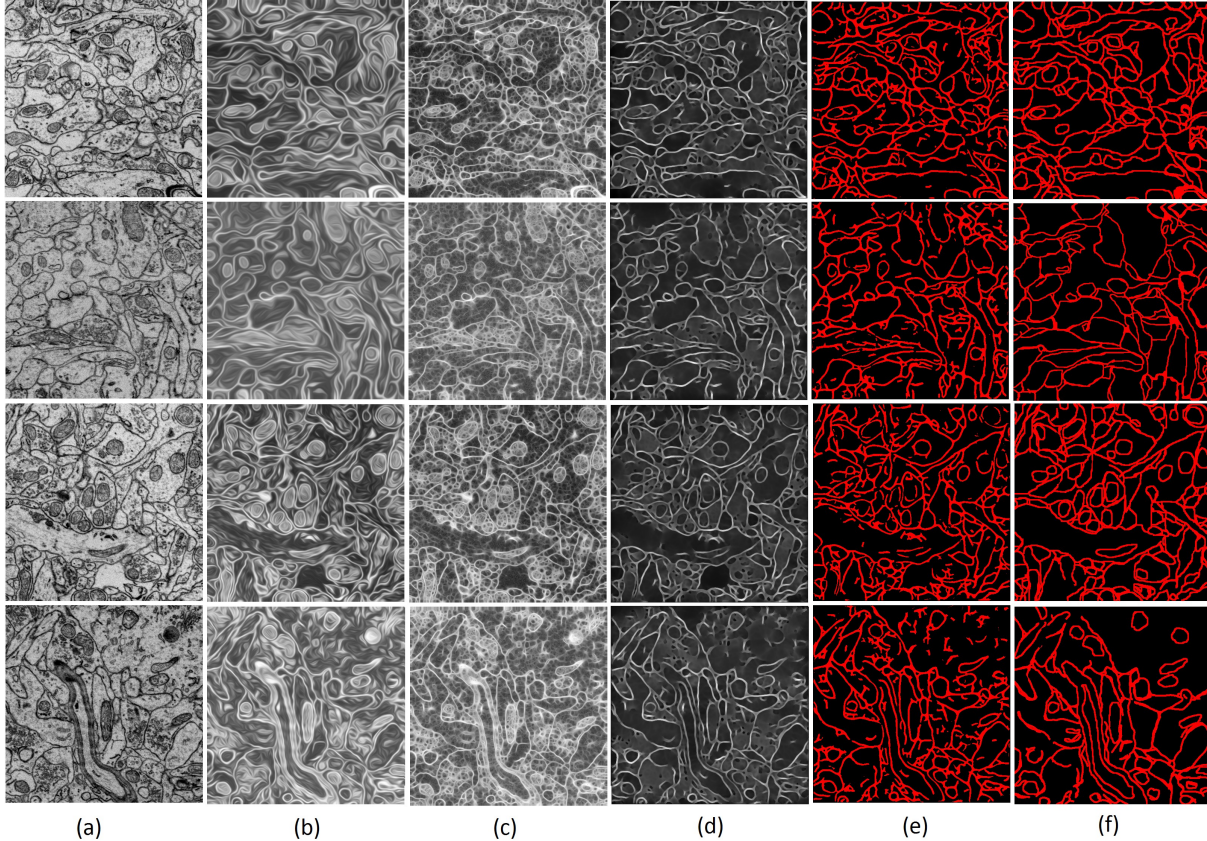
Figure 6. **Cell Boundary Enhancement Results**: Part (a) shows the 4 input images used. Each image is $1K \times 1K$ pixels. Part (b) and (c) presents the results of Coherence Enhancing Diffusion [17] and Hessian Based Diffusion [12], respectively. Part (d) presents the pixel-wise mean of the obtained Radon-Like features (using the extraction function in Eq. 5). Part (e) shows thresholded results for (d) and part (f) show the manually marked ground truth. In part (b) and (c) negative of the images are shown to make edges appear bright.

captures response of the most dominant GSD filter at each pixel, across various σ and ϕ. For illustration, Fig. 5(a) shows an input image and Fig. 5(b) shows response of GSD filters for four different orientations. Fig. 5(c) shows the corresponding $R(x, y)$ and 5(d) is the pixel-wise mean of our Radon-Like features accumulated over 180 uniformly sampled directions. It can be noted that Radon-Like features lead to structured smoothing of the image in Fig. 5(c) and the majority of the spurious non-boundary responses have been suppressed.

Evaluation: Fig. 6 qualitatively compares the performance of the proposed Radon-Like features to existing unsupervised boundary enhancement methods. Fig. 6(a) presents the input images, Fig. 6(b) and Fig. 6(c) show results obtained using Coherence Enhancing Diffusion [17] and Hessian Based Diffusion [12], respectively. Both of these techniques have been applied for cell boundary enhancement [12]. Fig. 6(d) shows results obtained using our method as described in this section. Fig. 6(e) shows a thresholding of the images in Fig. 6(d) and Fig. 6(f) shows the ground truth images. The threshold value used in Fig. 6(e) was chosen heuristically. Note that our features were

able to aggregate the information obtained from Gaussian-Second-Derivative filters to obtain clearly demarcated crisp boundaries.

Next we quantitatively evaluate the cell boundary enhancement methods through a segmentation experiment. Here we used twenty 2D tiles (512 x 512 pixels) with the above mentioned three boundary enhancement methods and the results were thresholded to obtain a binary segmentation of the images. The Receiver Operating Characteristic (ROC) curves in Fig. 7(a) shows the obtained False Positive Rate vs. True Positive Rate plots for pixel-wise segmentation as the threshold was varied. It can be noted that Radon-Like features perform comparatively better than the methods outlined in [17] and [12]. For the sake of completeness, we also plot ROC curves for the cases when the dominant edge response image ($R(x, y)$ in Eq. 4) is directly thresholded, used as input for Coherence Enhancing Diffusion and used as input for Hessian Based Diffusion. These results are better than [17] and [12] directly applied to input images as they benefit from the extraction function designed for our Radon-Like features.

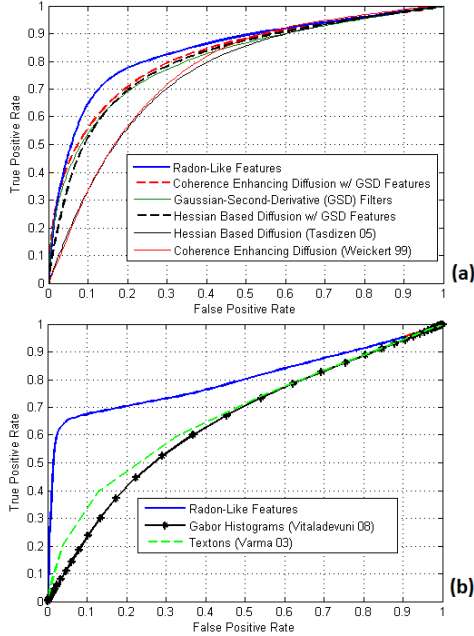Since the intention here is to demonstrate the per-

Figure 7. **ROC Curves**: (a) ROC curves for pixel-wise segmentation of cell boundaries. (b) ROC curves for pixel-wise segmentation of mitochondria.

formance of various features under similar constraints, we have used a simple thresholding based segmentation scheme. Moreover, the unsupervised nature of our method can allow it to work with images obtained through different staining processes (Fig. 1) without any retraining. We must also point out that even though a vector of features was obtained at each pixel using our method, only its mean value was used in the experiments above. In case a more sophisticated classifier is used, the whole feature vector can be provided as input instead of just its central moment.

3.2. Mitochondria Segmentation

Next, we focus on the mitochondria detection and segmentation problem. Mitochondria in Connectome EM images are characterized by their elliptical, but largely irregular shape and dark texture. They are also known to have ribbed structure inside their outer boundaries which can be confused with vesicles (see Fig. 1). Existing methods for mitochondria segmentation include [11], which goes after the shape via boosted ray features, and [16], which uses Gabor features with Boosting for the same task. It should be noted that the objective in [11] was only to localize mitochondria (within bounding boxes) and not to pixel-wise segment them, while [16] attempted pixel-wise segmentation. In addition to these methods, traditional texture segmentation methods [14] can also be used to segment the mitochondria.

The relatively unique geometry and the associated texture of mitochondria in EM image lend itself very well for analysis with Radon-Like features. The closed contours of

the mitochondria boundary (obtained using edge detection) indicate that extraction functions (such as in Eq. 3) that look at the distance between knots (edges derived) might be useful. But we have found that regions with vesicles can also lead to edges with similar spacing between them. Thus we go after the textural and the geometric information simultaneously. We accomplish this by using the edge maps of the intensity images $I(x, y)$ obtained using Canny filters to derive knots, and the mean of the intensity values as the extraction function. More explicitly, the extraction function is defined as

$$T_3(I, \mathbf{l}(t)) = \frac{\int_{t_i}^{t_{i+1}} I(\mathbf{l}(t)) \partial t}{||\mathbf{l}(t_{i+1}) - \mathbf{l}(t_i)||_2}, t \in [t_i, t_{i+1}], \quad (6)$$

where $I(\mathbf{l})$ is the input intensity image evaluated along the line $\mathbf{l}$ and t_i and t_{i+1} indicate two successive knots on $\mathbf{l}$ as before. Note that feature values along line segments with lengths < 10 pixels were explicitly set to zero.

Evaluation: The qualitative results are presented in Fig. 8, where we show results on 4 different tiles of the Connectome EM data. Fig. 8(a) presents the input intensity images while Fig. 8(b) shows the pixel-wise mean of the estimated Radon-Like features. Fig. 8(c) shows results obtained using Textox based texture classification method [14] and Fig. 8(d) shows results obtained using Gabor feature histograms as suggested in [16] ([16] uses boosting with Gabor and intensity histograms). For both of these cases a nearest-neighbor classifier was used to obtain final pixel-wise segmentation. Segmentation results obtained using our Radon-Like features via thresholding are presented in Fig. 8(e) and the ground-truth is presented in Fig. 8(f).

Quantitative comparison of the above-mentioned three methods is presented in Fig. 7(b), where the False Positive Rate for pixel-wise segmentation is plotted against the True Positive Rate for different threshold values. These results are averaged over multiple 512×512 pixel 2D tiles. From both Fig. 8 and Fig. 7(b) the effectiveness of the proposed features can be readily noted.

3.3. Cell Background Segmentation

Background detection and subtraction, besides being a commonly used step in general segmentation, has been explored in particular for EM image analysis [2]. It could also benefit tasks such as cell tracking through a volume of Connectome EM data [13, 15]. The existing techniques for this task include morphological opening [3] and Median filtering (as used in [2]) of an image.

Since Radon-Like features allow us access to spatially varying structural units within an image, we can carry out image-morphology-like operations to a greater effect. In particular, with edge map of the intensity image as a guide for knots and an extraction function defined as,

$$T_4(I, \mathbf{l}(t)) = \min [I(\mathbf{l}(t))], t \in [t_i, t_{i+1}], \quad (7)$$

we can obtain effects superior to morphological opening of an image. The extraction function in Eq. 7 assigns all the
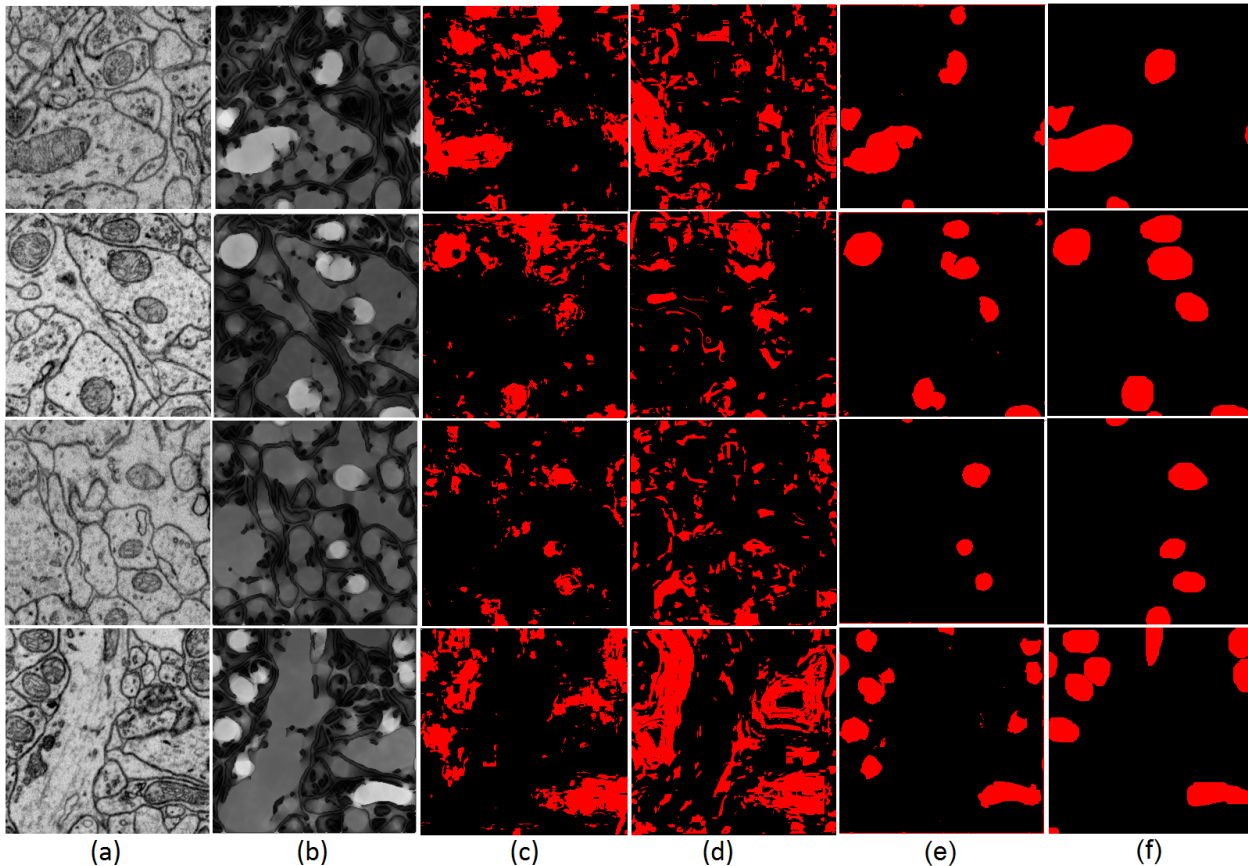
Figure 8. **Mitochondria Segmentation**: Column (a) shows 512×512 input images and column (b) shows the pixel-wise mean of the obtained Radon-Like features (using the extraction function in Eq. 6). Column (c) presents results obtained using the texton-based Texture Classifier [14] and column (d) show results obtained using Gabor histogram features as suggested in [16]. Column (e) shows the segmentation obtained using Radon-Like features (thresholding (b)) and (f) shows the corresponding ground truth.

pixel between the knots the minimum intensity value along l between the same two knots.

Evaluation: In Fig. 9 we present results obtained using the method for background detection described above and compare it with the traditional image morphological operation of opening. Fig. 9(a) shows the input images. For these images we computed the morphological opening using an oriented line segment as the structuring element. The line segment was chosen as the structuring element since it was closest to the scanning procedure used by Radon-Like features. We computed the morphological opening for each of the input image with 180 uniformly varying orientations of the structuring elements and its mean is presented in Fig. 9(b). The mean of the 180 long vector accumulated at each pixel using our method outlined above is presented in Fig. 9(d). Fig. 9(c) and Fig. 9(e) are the results obtained by thresholding the images in Fig. 9(b) and Fig. 9(d) respectively. The thresholds were heuristically chosen for both the methods. It could be noted that our method does a significantly better job at detecting the cell background than the morphological opening operation.

3.4. Vesicle Cluster Enhancement

Finally, we turn our attention to vesicle detection in Connectome EM images. Vesicles appear as very small circular structures that are generally clustered together. Existing methods like [2] go after detecting individual vesicles using morphological operations. It should be noted that the images used in [2], though still EM images, are qualitatively different from the Connectome EM images used in this paper. Detecting vesicles is relatively easier in almost textureless images used in [2] (see Fig. 2 in [2]). Results for vesicle detection in Connectome EM images were shown in [16] but the method was not described in any detail.

Since the Connectome EM images have a large number of structured and textured elements, locating vesicle clusters can help concentrate effort and suppress false positives for any vesicle specific post-processing. Here, using Radon-Like features, we describe a method that can be used to enhance the regions with clusters of vesicles in the EM images. For this task we use two edge maps of the input image $(I(x, y))$. Most edge detectors, like Canny filtering, include a threshold that decides the binary detection of edges. By
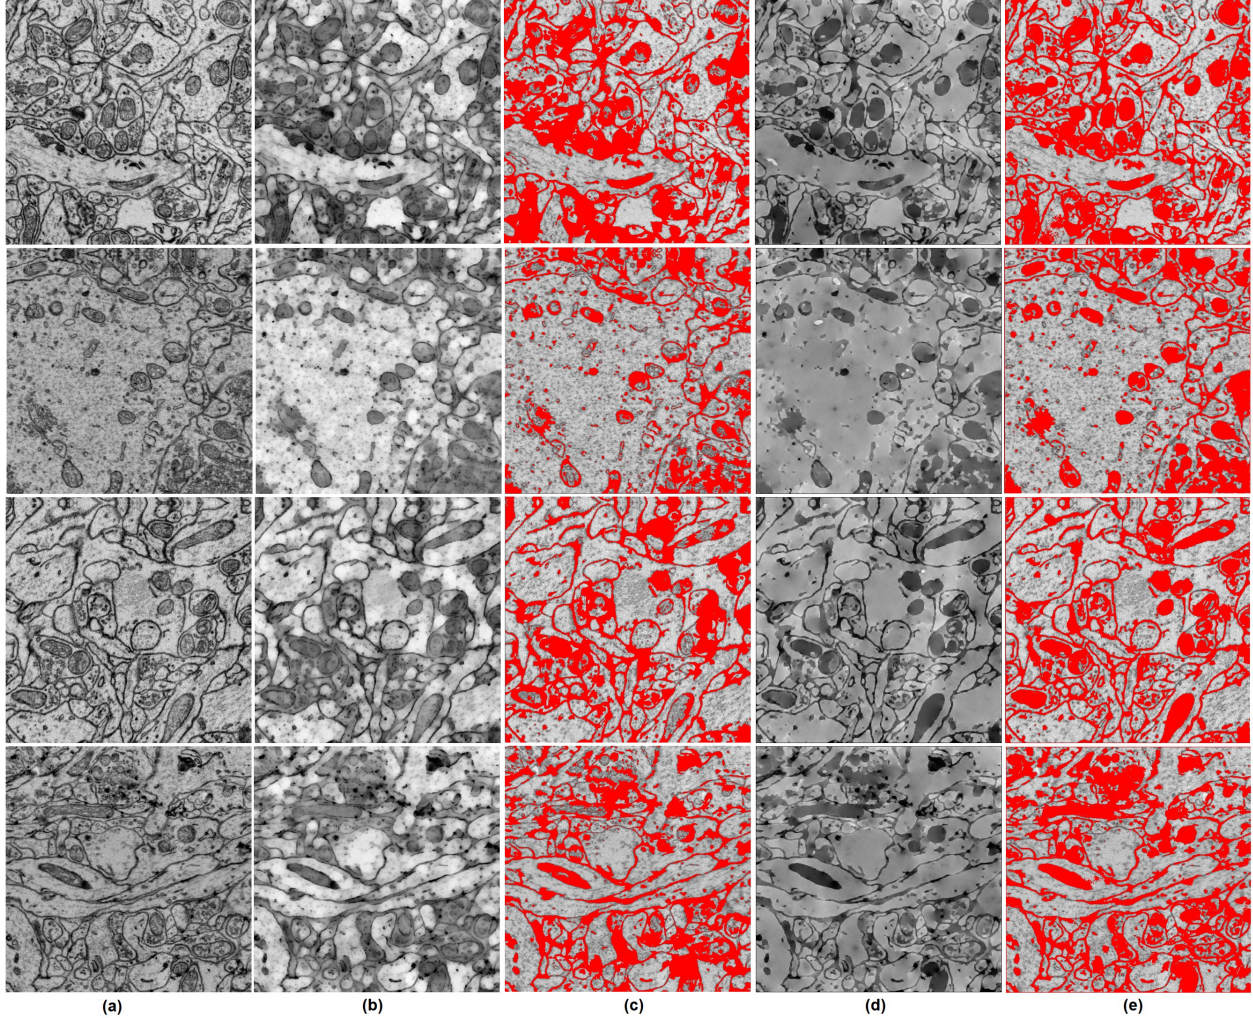
Figure 9. **Cell Background Segmentation**: Part (a) shows the input images. Part (b) is the pixel-wise mean of the results obtained via morphological opening with structuring element oriented along various directions. Part (c) shows (b) thresholded. Part (d) shows pixel-wise mean of Radon-Like features obtained with the extraction function defined in Eq. 7. Part (e) shows the images in (d) thresholded.

changing this parameter, weaker edges can be selected or ignored. One of the edge maps $(E_1(x, y))$ is obtained at the usual threshold used for knot generation and the other $(E_2(x, y))$ is obtained with a lot more of weaker edges included. Since regions with vesicle clusters create a lot of weak, randomly oriented edge, the central idea here is to capture the weak edge within the bounds of knots. We use the following extraction function for this task:

$$T_5(I, \mathbf{l}(t)) = \frac{\int_{t_i}^{t_{i+1}} E_2(\mathbf{l}(t)) \partial t}{||\mathbf{l}(t_{i+1}) - \mathbf{l}(t_i)||_2}, t \in [t_i, t_{i+1}]. \quad (8)$$

Note that $E_2(x, y)$ is a binary image obtained from $I(x, y)$. Note that feature values along line segments with lengths $<$ 10 pixels were explicitly set to zero.

Evaluation: Results obtained using the outlined method are presented in Fig. 10. Fig. 10(a) shows the input images while Fig. 10(b) shows the pixel-wise mean of Radon-Like features obtained using the extraction function T_5. Fig.

10(c) shows manually marked ground truth for comparison. The intention of this task is to locate possible clusters of vesicles within the very large Connectome images for further processing, possibly along the lines outlined in [2]. We must point out that another segmentation of vesicles can be obtained by a trivial combination of the schemes outlined in Sec. 3.1, 3.2 and 3.3.

4. Epilogue

In this paper we have introduced a novel class of features, called Radon-Like features, which can be a useful tool in aggregating information within structural units in an image. This class of features provides flexibility to design features for a specific task by selecting different extraction functions and knot maps. We have demonstrated the effectiveness of these features for various tasks associated with Connectome EM image analysis.
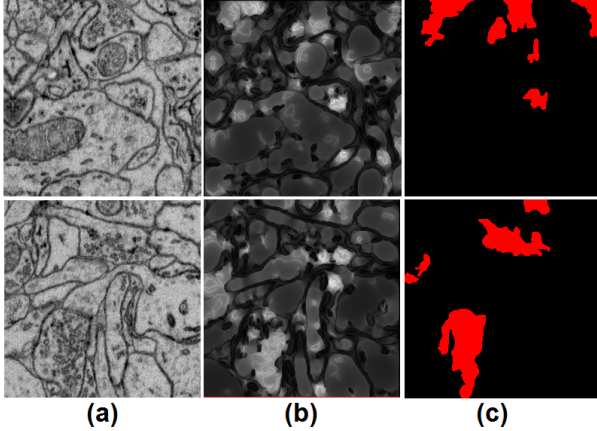
Computational efficiency of these features depends on

Figure 10. **Vesicle Cluster Enhancement**: Part (a) shows in input images. Part (b) shows the pixel-wise mean of the features obtained using Radon-Like features with extraction function as defined in Eq. 8. Part (c) shows the ground truth.

the specific extraction function used but the scanning procedure in general takes only $O(n^2)$ if the image size is $n \times n$ pixels. This is so because each pixel is visited only once during a scan along one direction. Quite importantly, these features lend themselves very well to parallel implementation. Since scans along various lines and directions are independent of each other, they can be performed in parallel with tremendous gains in efficiency. For the task of cell boundary enhancement outlined in Sec. 3.1, on a single CPU core, it takes 496 sec. while on an 64 core cluster, with different scan directions parallelized, it takes only 11 sec. to compute these features for 180 uniformly sampled directions for a 512×512 image. These results were obtained using a MATLAB implementation of these features.

Tunable parameters of Radon-Like features depend on the choice of the knot map and the extraction function. For the instances in this paper, knots were derived using Canny filter based edge detection with the threshold set to 0.45.

We have not overlooked the fact that most of segmentation results presented in this paper used a very simple thresholding based method. More sophisticated machine learning methods like Support Vector Machines and/or schemes like Markov Random Fields can possibly improve the over all results. But we contend that if the classification scheme is provided better features, it would also perform better. Hence, we feel that it is reasonable to use even a simple classifier, as long as it is the same across comparison, when the objective is to evaluate usefulness of features.

As mentioned before, Radon-Like features accumulate multi-dimensional features at each pixel. In all the experiments here, we have used only a scalar summary (mean) of these features at each pixel. We believe that the whole distribution of the features at each pixel can be useful in its classification and that is something we would like to explore in future. The Connectome image analysis, with its

peculiar characteristics, motivated the development of these features, but there is no reason these features could not be used for classification tasks at large e.g. face detection, object recognition etc. and that is also something that we wish to explore further.

5. Acknowledgement

References

[1] S. R. Deans. *The Radon Transform and Some of Its Applications*. John Wiley and Sons, 1983.

[2] E. Diaz, G. Ayala, M. E. Diaz, L.-W. Gong, and D. Toomre. Automatic detection of large dense-core vesicles in secretory cells and statistical analysis of their intracellular distribution. *IEEE TCBB*, 7, 2010.

[3] R. C. Gonzalez and R. E. Woods. *Digital image processing*. Prentice Hall, 2002.

[4] J. V. Grau, A. U. J. Mewes, M. Alcaiz, R. Kikinis, and S. K. Warfield. Improved watershed transform for medical image segmentation using prior information. *TMI*, 2004.

[5] V. Jain, J. F. Murray, F. Roth, S. Turaga, V. Zhigulin, K. L. Briggman, M. N. Helmstaedter, W. Denk, and H. S. Seung. Supervised learning of image restoration with convolutional networks. *ICCV*, 2007.

[6] E. Jurrus, A. Paiva, S. Watanabe, R. Whitaker, E. M. Jorgensen, and T. Tasdizen. Serial neural network classifier for membrane detection using a filter bank. *MIAAB*, 2009.

[7] N. Kasthuri and J. W. Lichtman. Neurocartography. *Neuropsychopharmacology*, 35:342–343, Jan 2010.

[8] T. Leung and J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textons. *IJCV*, 43(1), 2001.

[9] N. R. Pal and S. K. Pal. A review on image segmentation techniques. *Pattern Recognition (PR)*, 1993.

[10] J. Shi and J. Malik. Normalized cuts and image segmentation. *PAMI*, 2000.

[11] K. Smith, A. Carleton, and V. Lepetit. Fast ray features for learning irregular shapes. *ICCV*, 2009.

[12] T. Tasdizen, R. Whitaker, R. Marc, and B. Jones. Enhancement of cell boundaries in transmission electron microscopy image. *ICIP*, 2005.

[13] S. C. Turaga, J. F. Murray, V. Jain, F. Roth, M. Helmstaedter, K. Briggman, W. Denk, and H. S. Seung. Convolutional networks can learn to generate affinity graphs for image segmentation. *Neural Computation*, 22, 2010.

[14] M. Varma and A. Zisserman. Texture classification: Are filter banks necessary? *CVPR*, 2003.

[15] A. Vazquez-Reina, E. Miller, and H. Pfister. Multiphase geometric coupling for segmentation of neural processes. *CVPR*, 2009.

[16] S. N. Vitaladevuni, N. J. Sofroniew, Y. Mishchenko, A. Genkin, D. B. Chklovskii, and K. Harris. Automatic mitochondria detection in electron micrographs. *MIAAB*, 2008.

[17] J. Weickert. Coherence-enhancing diffusion filtering. *IJCV*, 31, 1999.

[18] H.-F. Yang and Y. Choe. Cell tracking and segmentation in electron microscopy images using graph cuts. *ISBI*, 2009.