

Segmentating Nucleus Membranes in SBFSEM Volume Data with Deep Neural Networks

Yassar Almutairi, Tim Cootes, and Karl Kadler

The University of Manchester, Manchester, UK,
yassar.almutairi@postgrad.manchester.ac.uk
{t.cootes, karl.kadler}@manchester.ac.uk

Abstract. We describe and evaluate a method for segmenting cell nuclei from serial block-face scanning electron microscope volumes. The nucleus is a roughly ellipsoidal structure near the centre of each cell, appearing as an irregular ellipse in each image slice. It is common to segment it manually, which is very time-consuming. We use a Convolutional Neural Network to locate the boundary of the nuclei in each image slice. Geometric constraints are used to discard false matches. The full 3D shape of each nucleus is reconstructed by linking the boundaries in neighbouring slices. We demonstrate and evaluate the system on several large image volumes.

1 Introduction

Serial block-face scanning electron microscopy (SBFSEM) creates detailed image stacks covering a volume of tissue [1]. The sample is embedded in a resin block, and a scanning electron microscope is used to take an image of one face of the block. An ultramicrotome shaves off the surface of the block, and a new image is taken. Each individual image may be up to 4096×4096 and there may be over 1000 individual slices in a volume. Thus the volume may contain many individual structures and it is very time consuming to manually examine each image. Techniques for automatically analysing such volumes are important for bio-science.




Fig. 1. Stack of images showing one nucleus (left) and 3D reconstruction of the nucleus (right)

This work aims to automate the process of segmenting and reconstructing cell nuclei from SBFSEM volumes, to enable biologists to studying the structure of the nucleus in 3D, qualifying nuclei and measuring their volume, radii and length (see Figure 1).

Studies on nuclei analysis can be dated back six decades [2]. Many methods have been proposed for segmenting cells for a wide range of applications, including cell counting, cell typing, and studies of cell migration and interaction. In general, there is a significant complexity and variability in cell image data, which makes segmenting cells from other objects more challenging. However, most of the work on nuclei detection is on images with much lower resolution than the SEM images we work on.

The images we work with are very high resolution, so that details of the nucleus can be resolved - the nucleus can be hundreds of pixels across (see Figure 1). We seek to find the membrane around the nucleus so that we can reconstruct the 3D shape. Automated methods for this do not appear to have been proposed before.

Our method uses a CNN to identify pixels belonging to the membrane of each nuclei, then post-processing to identify complete membranes. We use these to propose candidates for nuclei in every image, then link the candidates across neighbouring images to form extended 3D structures.

In the following we describe the approach in detail, and present results evaluating the technique.

2 Related Work

Several works on nucleus segmentation have been published [3][4]. It should be noted, however, that although these approaches might work efficiently on the authors data sets, their wider application has not been demonstrated. In [4], a smoothing is applied to strengthen the nucleus border and maintain the low contrast between background and cytoplasm. The edges are then enhanced by applying a gradient direction edge enhancement operator to obtain a better performance. The nuclear membrane is refined using a nucleus intensity threshold. Their approach was designed for isolated cells, however, not for multiple cells.

Convolution neural networks (CNN) have been shown to be very effective for a wide range of tasks including object detection [5][6] and image classification [7][8][9]. A limitation is that to get good performance CNNs need to be trained with large numbers of training examples.

Most of the recent studies on EM images are focused on segmenting and reconstruction neurones. Jian et al. [10] trained CNN a pixel classifier to classify each pixel into two classes (membrane or non-membrane). Ciresan et al. applied Deep neural networks (DNN) to detect membrane neuronal [11] and mitosis detection in breast cancer [12]. Both methods use a pixel classifier where the classifier is trained on patches centred on the membrane itself, and their work achieves good results.

For nucleus detection, most of the existing work in literature on EM images was performed to detect nuclei centroids in histopathology images where details of the nuclei cannot be resolved. In [13] Xu et al. used stacked sparse auto-encoder on high-resolution histopathological images for nuclei detection by learning high level features of nuclei using pixel intensity. Recently a structural regression CNN [14] has been proposed which learn a proximity map of nucleus. They modify the typical CNN by replacing the last layer with structured regression layer to encode topological information.

Because we only have relatively small numbers of individual nuclei, we do not have enough to train a general "whole nuclei" detector. Instead we train a CNN to identify the boundaries, and use post-processing techniques to identify the whole boundary and reconstruct the 3D shape. This has the advantage that we can easily deal with the partial nuclei which intersect the edge of the image volume.

3 Method

SBFSEM images are time-consuming to collect (involving extensive sample preparation and many hours to image each sample), we only have small numbers of volumes. It is thus impractical to annotate large numbers of complete nuclei, which one might like to do to build a comprehensive detector. Instead we work with a small number of annotated nuclei, each containing many labelled pixels (both on and off the target boundary). We train a CNN classifier to identify whether a pixel belongs to a nucleus membrane or not. This can identify the nucleus, but also responds on similar looking membranes such as mitochondria and the Golgi apparatus, leading to many false positives. These objects' membranes share some features of nuclei membranes and need to be eliminated. We identify every connected component and perform morphological operations to remove noise and discard false matches.

In the final stage, we reconstruct every nuclei through the image stack. The nuclei radii in SBFSEM range from 200 pixels to 900 pixels. Finding the centre of gravity of each nucleus and matching it its correspondence in the next frame is sufficient to link them.

3.1 Convolutional Neural Network Architecture

Our segmentation algorithm is based on training a CNN to perform pixel binary classification. The architecture of the network consists of six learning layers, four convolutional layers, and two fully connected layers, as shown in Figure 2. The output of the network is the probability that the central pixel belongs to a nuclear membrane or not. The CNN is trained on $N \times N$ patches (we choose $N = 64$ to provide a sufficient trade-off between the local information to be extracted and memory usage) centred on boundary pixels (positives) and background pixels (negatives).

4 Y. Almutairi et al.

Our architecture is similar to [11], but our network includes a dropout layer to prevent overfitting [15]. The kernel size and number of output maps is larger in the first convolutional layers to capture more information about variability in the low-level features. The output of the last fully connected layer is fed into a binary softmax to produce the probability of class label.




Fig. 2. CNN architecture consisting of four convolutional layers and two fully connected layers where each convolutional layer is followed by Rectified Linear Units (ReLU) and Max-pooling.

Each layer of convolution and fully connected output is followed by Rectified Linear Units (ReLU). Max pooling is applied to the output of each convolutional layer to prevent overfitting and to down-sample by a factor of 2. The second-to-last fully connected layer is followed by a drop out layer.

3.2 Training Examples

The training examples for nuclei segmentation were collected from a stack of images consisting of 157 slices with a resolution of 966×876 . We selected 16 representative slices, and annotated the nucleus membranes by marking the middle of the nucleus membrane. Membranes around other structures, such as Golgi and mitochondrial cytoplasm membranes, were annotated in the same way, to provide negative examples. We then extracted patches with a size of 64×64 centred on the annotated points.

We augmented our samples by applying eight different rotations for each, and different scaling and reflection, giving a training set of 1.2M samples (50% positive).

3.3 Removing False Positives

The segmentation results contain many false positives (see Figure3.b) due to similarities in the appearance of the membranes of different structures. How-

ever, since the nuclei are usually the largest structures in the image, we can use straightforward processing steps to pick them out. We apply erosion to the binary image to eliminate thin false positives (Figure 3.c), with a radii chosen based on the known width range of the nuclei membrane. We then identify all connected structures and remove those with areas below a threshold, again chosen based on the known size ranges of the nuclei. This enables us to identify the nuclei membranes (see Figure 3.d).




Fig. 3. (a) Original image. (b) CNN pixel classification (c) Result after applying erosion. (d) Result after eliminating the small connected components.

4 Nuclei Detection and Reconstruction

The steps described above enable us to identify the parts of the nuclear membranes present in each 2D image slice through the volume. To find the nuclear membrane we find the contour of every connected component. Figure 4 shows an original image with the found contour following a nuclear membrane. Not all nuclear membranes contours are continuous, however, and some have broken links

6 Y. Almutairi et al.

between them. This is because (a) during tissue preparation errors can occur, especially when slicing the sample, and (b) the classifier can fail to detect the nuclear membranes during the segmentation process. To address this issue, we merged three neighbouring slices together by averaging them as the movement of the nuclei between slices is small. Thus, if a nuclear membrane is broken, it will be fixed by merging the previous and subsequent slice (see Figure 5). We then find the external contour for each nucleus.

As each nucleus in our dataset had little drifting between slices, and the distance is relatively large between nuclei, we used a simple approach for tracking nuclei and linking each nucleus to its match in the next slice. We fitted a disk based on the set of points we obtained from the nuclei contours, similar to the method in [16].

We assumed that all candidate disks in the first slice were the start of a nucleus, creating n_1 nuclei each containing a single disk. We then processed the subsequent frames one at a time, using the candidate disks to either extend an existing nucleus, or create a new nucleus if there was no match in the previous frame. This identified a large number of nuclei. By comparing the disks at the ends of the nuclei we could identify and correct the small gaps that were caused by detection failures. Short structures (those visible in fewer than five frames) were rejected as they were unlikely to be nuclei.




Fig. 4. (a) Original Image (b) Located contour around the nucleus membrane.

5 Experimental Results for Nuclei Segmentation

We applied our model to four different datasets, each consisting of hundreds of images that were acquired using the SBFSEM technique (see table 1). The images contained nuclei, mitochondria, and Golgi apparatus. Most of the nuclei had a roughly circular or elliptical cross-sections, whereas the mitochondria had a mostly circular shape. Golgi and mitochondrial membranes share several features and they look the same, but the mitochondrial inner body has parallel lines due

Segmenting Nucleus Membranes

7




Fig. 5. A sequence of three slices. slice ($z-1$), slice(z), slice($z+1$), and all combined

to the folding of the inner membrane. The outer membrane of mitochondria, however, sometimes appear as thick as in nuclei, and share the same appearance as nuclei. This led our classifier to classify some mitochondrial outer membranes as nuclei.

Dataset	Number of Slices	Image dimensions
Dataset 1	1000	3000×3000
Dataset 2	700	4096×4096
Dataset 3	603	4096×4096
Dataset 4	157	966×876

Table 1. Details of the four image volumes

5.1 Evaluation Criteria

For quantitative assessment, the ground truth and segmenting results are given by filled closed contours. The foreground objects (nuclei) are shown as white pixels and the background is shown as black pixels. To evaluate a binary segmentation, we selected the per-pixel Precision, Recall, Accuracy, and Dice coefficients [17] as performance measures:

$$Precision = \frac{TP}{TP + FP} \quad (1)$$

$$Recall = \frac{TP}{TP + FN} \quad (2)$$

$$Accuracy = \frac{TP + TN}{TP + FP + FN + TN} \quad (3)$$

$$Dice = \frac{2 \cdot TP}{TP + FP + TP + FN} \quad (4)$$

Where (TP) is true positives, (TN) is true negatives, (FP) is false positives,

8 Y. Almutairi et al.

and (FN) is false negatives of the classified pixels. A perfect segmentation is obtained when Precision and Recall are 1. When Precision is low, this indicates over-segmentation, and when Recall is low, it indicates under-segmentation.

Accuracy is the proportion of the true positives and true negatives among the total number of pixels. To assess the accuracy of image segmentation, the Dice coefficient is used to measure the spatial overlap. The Dice coefficient value ranges from 0 to 1, where 0 indicates no spatial overlap between two binary images and 1 indicates a complete overlap [17].

Selected qualitative segmentation results are shown in Figure 8 from the four datasets. This figure shows the original image, the manually annotated nuclei, and the segmentation results.

We annotated a total of 120 images manually, selecting 40 from datasets 1 and 4, and 20 from data sets 2 and 3. Images from the first three datasets each image roughly 10 nuclei, those from dataset 4 contain only two nuclei.

Figure 6 summarises the results. The average accuracy is above 96% for dataset 1, 96.5% for dataset 2, 97% for dataset 3, and 95.5% for dataset 4.

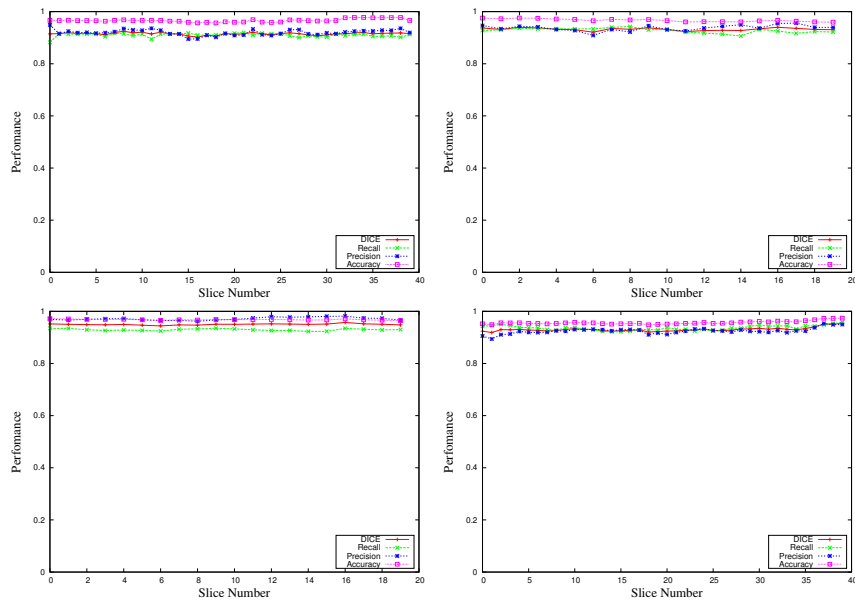


Fig. 6. Quantitative nuclei segmentation results for the four datasets.

Segmenting Nucleus Membranes 9

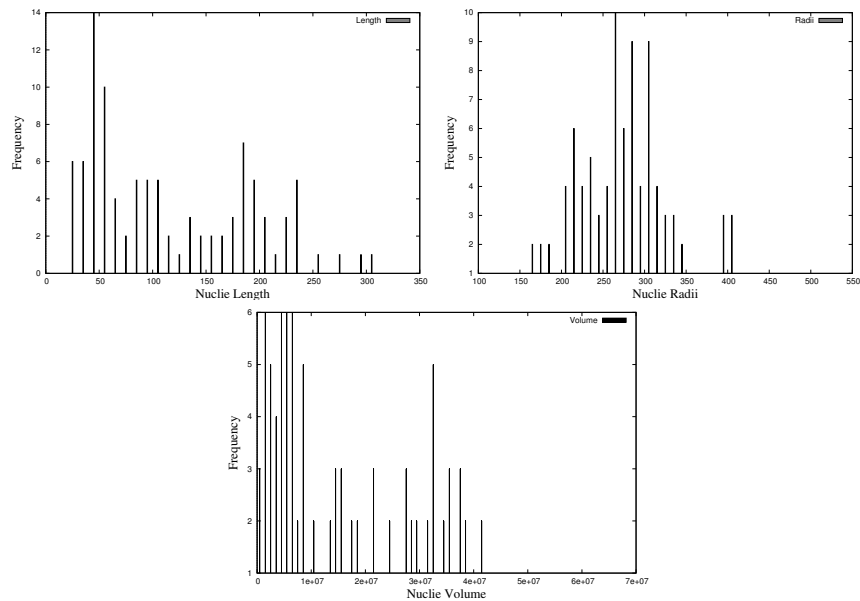


Fig. 7. Quantitative analysis of nuclei length, radii and volume for Dataset 1.

Figure 9 and 10 shows the full reconstruction of nuclei from Dataset 1 and 2. The utility of this visualisation is done using IMOD software and random colour was given to each nucleus.

5.2 Quantifying the Nuclei

Biologists are interested in quantifying nuclei based on shape information such as length, radii and volume. We show histograms of the length, average radii and volume of each nucleus for dataset 1 in Figure 7 (where the length is the number of slices in which the structure appears). To measure the radii of each nuclear we fit a circle based on the nuclear membrane contours points for each slice and then take the mean radii. The volume is calculated from the total number of pixels inside nuclear membrane for each slice.

6 Discussion and Conclusions

We describe a novel method for segmenting and reconstructing nuclei by extracting the 2D boundaries of each nucleus from SBSEM image stack. The main strength of our method lies in that a CNN that trained on small patches gives a powerful classifier. The method provides accurate results, and will enable automatic analysis of large image volumes from SEM images.




Fig. 8. Random image were selected from each dataset. First column shows the original image, second column shows the ground truth and the third column shows our segmentation result. First row from Dataset 1, second row from Dataset 2, third row from Dataset 3 and the fourth row from Dataset 4.




Fig. 9. Visualisation of all nuclei of Dataset 1 found by the algorithm.




Fig. 10. Visualisation of all nuclei of Dataset 3 found by the algorithm.

References

1. Winfried Denk and Heinz Horstmann. Serial Block-Face Scanning Electron Microscopy to Reconstruct Three-Dimensional Tissue Nanostructure. *PLoS Biol*, 2(11):e329, 2004.
2. Erik Meijering. Cell Segmentation: 50 Years Down the Road. *IEEE Signal Processing Magazine*, 2012.
3. H Irshad, A Veillard, L Roux, and D Racoceanu. Methods for Nuclei Detection, Segmentation and Classification in Digital Histopathology: A Review. Current Status and Future Potential. *Biomedical Engineering, IEEE Reviews in*, 2013.

12 Y. Almutairi et al.

4. Shys Fan Yang-Mao, Yung Kuan Chan, and Yen Ping Chu. Edge enhancement nucleus and cytoplasm contour detector of cervical smear images. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 2008.
5. Christian Szegedy, a Toshev, and D Erhan. Deep Neural Networks for Object Detection. *Advances in Neural Information ...*, 2013.
6. Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2014.
7. A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. *NIPS'12 Proceedings of the 25th International Conference on Neural Information Processing Systems*, 2012.
8. Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2015.
9. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015.
10. V Jain, J F Murray, F Roth, S Turaga, V Zhigulin, K L Briggman, M N Helmsstaedter, W Denk, and H S Seung. Supervised Learning of Image Restoration with Convolutional Networks. *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, 2007.
11. Dan C. Ciresan, Alessandro Giusti, Luca M. Gambardella, and Jurgen Schmidhuber. Deep Neural Networks Segment Neuronal Membranes in Electron Microscopy Images. *Nips*, 2012.
12. D C Ciresan, A Giusti, L M Gambardella, and J Schmidhuber. Mitosis Detection in Breast Cancer Histology Images using Deep Neural Networks. *Proc Medical Image Computing Computer Assisted Intervention (MICCAI)*, 2013.
13. Jun Xu, Lei Xiang, Qingshan Liu, Hannah Gilmore, Jianzhong Wu, Jinghai Tang, and Anant Madabhushi. Stacked sparse autoencoder (SSAE) for nuclei detection on breast cancer histopathology images. *IEEE Transactions on Medical Imaging*, 2016.
14. Y. Xie, F. Xing, X. Kong, H. Su, and L. Yang. *Beyond classification: Structured regression for robust cell detection using convolutional neural network*. 2015.
15. Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 2014.
16. Emo Welzl. Smallest enclosing disks (balls and ellipsoids). In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 1991.
17. Kelly H. Zou, Simon K. Warfield, Aditya Bharatha, Clare M.C. Tempany, Michael R. Kaus, Steven J. Haker, William M. Wells, Ferenc A. Jolesz, and Ron Kikinis. Statistical Validation of Image Segmentation Quality Based on a Spatial Overlap Index. *Academic Radiology*, 2004.