Automated Membrane Detection in Electron Microscopy using Convolutional Neural Networks

Joris Roels   JORIS.ROELS@TELIN.UGENT.BE

Ghent University, TELIN, Sint-Pietersnieuwstraat 41, B-9000, Ghent, Belgium

VIB, IRC, Technologiepark 927, B-9052 Ghent (Zwijnaarde), Belgium

Yvan Saeys   YVAN.SAEYS@IRC.VIB-UGENT.BE

VIB, IRC, Technologiepark 927, B-9052 Ghent (Zwijnaarde), Belgium

Jonas De Vylder   JONAS.DEVYLDER@TELIN.UGENT.BE
Jan Aelterman   JAN.AELTERMAN@TELIN.UGENT.BE
Wilfried Philips   PHILIPS@TELIN.UGENT.BE

Ghent University, TELIN, Sint-Pietersnieuwstraat 41, B-9000, Ghent, Belgium

Abstract

A well-known problem in order to analyze electron microscopy images is that the corresponding segment extraction typically is a tedious and time-consuming process. This is due to the currently used image processing techniques which are hampered by large time or memory consumption or simplistic underlying models. Advanced techniques that obtain high segmentation quality exist, but suffer from practical limitations such as difficult algorithm interpretability or parameter dependencies. We therefore propose a trainable segmentation technique that is based on convolutional neural networks and extracts mitochondrial membranes from EM images. The training phase depends on a limited number of manually annotated samples. Once the model is trained, segmentation becomes completely parameter-independent.

1. Introduction

One of the most popular subjects of biological cell research are mitochondria because of their importance in the cell's energy maintenance. The study of energy

overflow (respectively, underflow) resulting into superfluous cell division (respectively, cell death) requires high throughput mitochondria analysis. Even though there has been a significant amount of research in the field of automated mitochondria segmentation (Seyedhosseini et al., 2013; Jorstad & Fua, 2015), biological researchers have not yet adopted these algorithms because of their complex interpretability and/or parameter dependencies. We propose a membrane detection method based on learned convolutional neural networks (CNN), specifically designed for electron microscopy (EM) images. Using these membrane segments, automated mitochondria segmentation becomes significantly easier. The algorithm requires a one-time limited training in order to determine its parameters. Once the model is trained, segmentation becomes completely parameter-independent.

2. CNN based segmentation

Our segmentation algorithm performs pixel-wise binary classification based on a 6-layer CNN. For each pixel, a local $N \times N$ window (we chose $N = 32$ as this provides a good trade-off between local information and memory usage) is extracted and fed through a variant of a network that has been optimized for EM segmentation (Ciresan et al., 2012) (see Figure 1). In contrast to (Ciresan et al., 2012), our network incorporates dropout layers, which are well-known to counteract overtraining (Srivastava et al., 2014). Additionally, the kernel size and number of output maps

Figure 1. Proposed network architecture: 4 convolutional layers with intermediate rectified linear units and max-pooling layers are followed by 2 fully connected layers. The first fully connected layer is connected to a dropout layer in order to avert overtraining. Layer output dimensions are shown below, the convolution kernel size is given between brackets.


Figure 2. (a) Input image, (b) extracted membrane probability map and (c) segmentation result after thresholding.

in the first convolutional layers is larger. This way, the network can capture more variability in the low-level features. More specifically, the network consists of 4 consecutive convolutional layers, followed by 2 fully connected layers, resulting in a binary output that represents the probability of the reference pixel being part of mitochondrial membrane or not. Each layer is followed by a rectified linear unit. In order to prevent overtraining, each convolutional layer is followed by a max pooling unit that samples down by a factor of 2 and the second-to-last fully connected layer is followed by a 50% dropout layer. The network was trained by 6-fold cross-validation and the resulting output probability maps were averaged in order to get the final mapping.

3. Results

The algorithm was implemented using the Caffe deep learning framework (Jia et al., 2014) which exploits GPU processing power. Note the number of training samples equals the number of pixels to be annotated by the user. Therefore, this should be kept as low as possible. However, convolutional networks typically require a significant amount of training samples. The solution for this is augmentation (we used random rotation). For training, we employed $300,000$ annotated samples (which is a fraction of five fully segmented mitochondria). After augmenting, we obtained $600,000$ samples for training. The derived membrane probability map and segmentation result is shown in Figure 2. The mitochondria membranes are separated relatively well. However, the results may be improved more by applying proper pre-processing techniques such as image restoration in order to improve the input data quality. Additionally, alternative network architectures should be studied in order to keep the required amount of input samples as low as possible without introducing overtraining.

References

Ciresan, D., Giusti, A., Gambardella, L., & Schmid-huber, J. (2012). Deep Neural Networks Segment Neuronal Membranes in Electron Microscopy Images. *Nips*, 1–9.

Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., & Darrell, T. (2014). Caffe: Convolutional Architecture for Fast Feature Embedding. *Proceedings of the ACM International Conference on Multimedia*, 675–678.

Jorstad, A., & Fua, P. (2015). Refining mitochondria segmentation in electron microscopy imagery with active surfaces. *Proc. European Conference on Computer Vision* (pp. 367–379).

Seyedhosseini, M., Ellisman, M. H., & Tasdizen, T. (2013). Segmentation of mitochondria in electron microscopy images using algebraic curves. *IEEE International Symposium on Biomedical Imaging* (pp. 860–863).

Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. (2014). Dropout : A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research (JMLR)*, *15*, 1929–1958.