Efficient Scanning for EM Based Target Localization

Raphael Sznitman, Aurelien Lucchi, Natasa Pjescic-Emedji,
Graham Knott, and Pascal Fua

Computer Vision Lab,
École Polytechnique Fédérale de Lausanne, Switzerland
`firstname.lastname@epfl.ch`

Abstract. For biologists studying the morphology of cells, Electron Microscopy (EM) is the method of choice with its *nm* resolution. However, the time necessary to acquire EM image series is long and often limits both the number and size of samples imaged. This paper presents a strategy for fast imaging and automated selection of regions of interest that significantly accelerates this process. In particular, this strategy involves scanning a tissue sample once, finding regions of interest in which target structures might be located, scanning these regions once again, and iterating the process until only relevant regions of the block face have been scanned repeatedly. For mitochondria and synapses, this approach is shown to produce near equal localization results to current state-of-the-art techniques, and does so in almost a tenth of the time.

1 Introduction

Focused Ion Beam Scanning Electron Microscopes (FIB-SEM) and their ability to image with isotropic resolution of up to $4nm$ per pixel are becoming invaluable tools in studying cell ultrastructure and model organelles, such as mitochondria, synapses, and vesicles. Acquiring images such as those depicted in Fig. 1 involves repeatedly milling a few *nm* from the surface of a tissue block using a focused gallium ion beam, scanning each line of a rectangular region of the block face several times, averaging the results, and milling again.

The resulting images have already yielded many new insights in the structure and functioning of cells [1, 2], but the acquisition process is desperately slow. For example, imaging the $10 \times 10 \times 10\mu m$ tissue block of Fig. 1 at full resolution took approximately 50 hours. Such lengthy processing times are limiting because neuroscientists now require larger volumes to enable multiple cells, and even entire tissue samples, to be analyzed, which would currently be prohibitively slow. Furthermore, because thermal changes can cause the block face to drift and produce misaligned image series, considerable precision is needed to maintain consistent imaging. This is difficult to achieve over extremely long periods and limits the size of the images that can be captured.

In short, a pressing need exists to reduce scanning time without compromising the usefulness of the resulting images. Some research has already gone into

Fig. 1. FIB-SEM Scanning Process. The microscope repeatedly scans the surface of a sample block until a clear image is produced.

achieving such a result. For example, synthetic and real data have been used to show that Sparse Sampling techniques could potentially help [3]. However, the current generation of microscopes simply cannot perform the required random sampling because imaging extremely small and random locations on the block face would be incredibly time consuming.

In contrast, this paper uses a combination of real and synthetic data to demonstrate an approach to achieving much faster scanning using existing technology [4] when the images are intended for automated detection, counting, and modeling of organelles. The strategy involves scanning each image line of the block surface *once* instead of several times, finding rectangular regions of interest in which target structures might be located, scanning these regions once again, and iterating the process until only relevant regions of the block face have been scanned repeatedly. This process will be shown to result in much reduced imaging time at almost no performance loss.

The remainder of this paper is organized as follows. The paper begins by specifying our approach in Sec. 2, and provides a problem formulation in Sec. 2.1 and an algorithm description in Sec. 2.2. In Sec. 3, the method is validated experimentally. The paper concludes with final remarks in Sec. 4.

2 Sequential Region Cascades

The following observations are used as the starting point:

1. If the goal is to count or model organelles that occupy only a small fraction of a tissue block, precisely imaging the entire tissue block is a waste of time. This concept has long been known and exploited by Computer Vision researchers [5], [6] but, to the best of our knowledge, has been neglected by microscopists.
2. Once an organelle has been found in one slice, it will be seen with very high probability at similar locations in subsequent slices, *e.i.* organelles are 3D structures. Similarly, a slice must exist in which an organelle first appears, as well as a final one.

3. Current microscopes are good at quickly scanning rectangular regions, and new scanning engines that will make this process entirely programmable on a slice-by-slice basis will soon be commercially available [4]. In contrast, microscopes that can quickly and randomly access specific image locations are likely to remain beyond the state-of-the-art for the foreseeable future, thus precluding the use of standard Sparse Sampling techniques [3].
4. The microscopes are optimized to scan lines at one particular speed. Total scanning time, and image quality, are controlled by the number of times each line of the rectangular regions of interest is scanned. Typically, the final result is the average of these scans.

Therefore, we propose our Sequential Region Cascades (SRC) approach for accelerating the scanning process by building a cascade of region classifiers and sequentially evaluating regions that appear likely to contain a particular intracellular structure. While most cascade systems increase classification accuracy over levels, given the technical constraints discussed above, classification is instead performed on images acquired with varying scan counts. To take advantage of what was found in one slice and accelerate the scanning in the next, target locations are directly imaged using locations from previous slices with the largest possible numbers of scans, and new targets are searched for in the remainder of the slice using the cascade.

In the remainder of this section, the SRC approach is first formalized, and then its implementation is discussed. Section 3 shows that it allows counting of both mitochondria and synapses with much reduced acquisition times and at almost no loss in accuracy.

2.1 Problem Formulation

We formalize our problem as follows. Let the volume to image be denoted by $\mathcal{V} = \{S^1, \dots, S^T\}$, where S^t corresponds to a slice of the volume. When using a scanning EM microscope, we consider two sets of parameters when acquiring images. First, we define a rectangular region to scan, $R = (r_1, r_2)$, where r_1 and r_2 are the upper left and bottom right pixel coordinates, respectively. Second, the *scan count* is defined as the number of times the electron beam images one pixel and denote this value as $C = \{1, \dots, C_{\max}\}$.

Given these two parameters, (R, C) , the process of acquiring an image by scanning a region of a slice can be described by the function f , where $f : S \times R \times C \mapsto I$. That is, evaluating the function $f(S, R, C)$ provides an image I_C^R of size R and corresponds to the average of C independently scanned samples. Typically, the time cost associated with evaluating $f(S, R, C)$ is $C \times \text{area}(R)$.

As in [7–9], we may train a classifier to verify if a pixel in I_C^R belongs to a target organelle. As can be seen in Fig. 1, images acquired using different scan counts exhibit different statistics and we may train different classifiers for each. We therefore take a family of classifier to be $\mathcal{H} = \{h_1, \dots, h_{C_{\max}}\}$, where $h_c : I_C^R \mapsto \{0, 1\}^R$. Note that these classifiers return binary images.

Algorithm 1. Sequential Region Cascades (SRC)

```

1:  $\mathcal{P} \leftarrow$  empty queue,  $\mathcal{Q} \leftarrow$  empty queue.
2: for  $t = 1, \dots, T$  do
3:   Push( $([0, 0], [M, N]), 1$ ) into  $\mathcal{Q}$ 
4:   RemoveOverlaps( $\mathcal{Q}, \mathcal{P}$ )
5:   while  $|\mathcal{Q}| > 0$  do
6:      $[R, C] \leftarrow \text{Top}(\mathcal{Q})$ 
7:      $I_C^R \leftarrow f(S^t, R, C)$ 
8:      $\mathcal{R} \leftarrow \text{ExtractRegions}(h_C(I_C^R))$ 
9:     if  $C_{\max} == C$  then
10:      Push( $\mathcal{R}, C$ ) into  $\mathcal{P}$ 
11:     else
12:      Push( $\mathcal{R}, C + 1$ ) into  $\mathcal{Q}$ 
13:     end if
14:   end while
15: end for

```

Finally, let P^t to be the set of pixels corresponding to the location of target structures in slice S^t . Our goal is to discover these sets $\{P^t\}_{t=1}^{t=T}$ for all slices as efficiently and as quickly as possible.

2.2 Algorithm and Implementation

An outline of the algorithm is shown in Alg. 1. To begin, the user provides the set of classifiers and the volume to image, $\mathcal{H} = \{h_1, \dots, h_{C_{\max}}\}$ and $\mathcal{V} = \{S^1, \dots, S^T\}$, respectively. The algorithm begins by forming two queues that will maintain tuples of regions and scan counts, *i.e.* (R, C) . The first, $\mathcal{P}$, maintains a set of regions deemed the target structure on a given slice. The second, $\mathcal{Q}$, maintains an intermediary list of candidate regions that appear likely to contain target structures within them. Initially, both queues are empty.

For each new slice, we begin by pushing the entire observable domain as a candidate region using the smallest scan count. The following sequence of steps is then looped, which is called the *refining stage* and is depicted by Fig. 2 (lines 5 to 14): A candidate region and scan count index from the queue $\mathcal{Q}$ is retrieved. The associated image region, I_C^R , is then acquired and the corresponding classifier is evaluated by computing $h_C(I_C^R)$. At this point, the binary classification image is searched for disjoint sets of rectangular regions that indicate potential target locations. If newly extracted regions were acquired using the highest possible scan count, these are pushed into the target region queue, $\mathcal{P}$, otherwise, they are pushed into the candidate region queue, $\mathcal{Q}$. This process iterates until no candidate region remains.

In addition, before starting the *refining stage*, the overlapping region that coexists in both $\mathcal{Q}$ and $\mathcal{P}$ are removed. This effectively creates a new set of regions in $\mathcal{Q}$ that are disjoint of $\mathcal{P}$ and reduces the direct need for searching targets likely to have stayed in the same location. Doing so is one way of encoding

Fig. 2. SRC Refining Stage: For any slice, the queue Q contains a list of tuples (R, C) that specify a region and scan count with which the microscope should image. Once an image region acquired, I_C^R , a dedicated classifier assigns a binary label to each location in the image for the presence of targets. New tuples are then formed and inserted into Q or into the permanent queue P .

3D information for location of targets in a volume. Obviously, during the first slice of the tissue block, this step is irrelevant because P is empty.

3 Experiments and Results

We used a Zeiss NVision40 FIB-SEM microscope to mill and scan a rodent brain sample of $10 \times 10 \times 1 \mu m$, which produced 165, 1024×1536 images. We also collected a second stack of 377, 655×429 images. In both cases, each line was scanned $C_{\max} = 44$ times. This took 5 and 12 hours, respectively. We evaluated our algorithm for the tasks of localizing two types of organelle: mitochondria and synapses. Here, we show how the SRC strategy could be used to divide the scanning time by a factor of 10 to 15 depending on the target type.

Test Data: Given that the scanning engine of our NVision40 microscope was not designed for this, implementing our approach on it would be very difficult. However, this will soon change when newer scanning engines come to market [4]. To demonstrate our approach in the meantime, we therefore proceed as in [3] and synthesized the scans we would have gotten using values of $C < C_{\max}$ by appropriately degrading the higher-quality ones. To ensure realism we proceeded as follows.

Before scanning, and thus destroying, the whole first block, we collected six independent images $\{I_1, \dots, I_6\}$ of the first slice using a single scanning pass for each and a single one, which we denote as $\hat{I}$, using 6 passes. From this, we first verified that for any pixel location u , the gray-level $I_i(u)$ in any of the 6 single-scan images is well approximated by a Gaussian of mean $\hat{I}(u)$ and standard deviation $\sigma_u = m\hat{I}(u) + b$, where m and b are linear regression parameters, as

Fig. 3. Reconstructing small scan count images. (left) Relation between Average intensity and standard deviation for pixels acquired with low scan counts. (right) Top rows shows real images acquired under 1, 4, 12 and 44 scan counts, and the bottom shows the corresponding reconstructed images.

illustrated by Fig. 3 where we plot μ_u against σ_u for 10'000 randomly selected pixels. In other words, the gray level variance is directly proportional to the gray level value.

From this, using an image acquired with a large number of scans, we can simulate acquiring an image from a smaller number of scans. For example, a pixel with n scans can be reconstructed by sampling the Gaussian $\mathcal{G}(\hat{I}(u), m\hat{I}(u) + b)$ n times and averaging the samples. Fig. 3 shows the true (top row) and reconstructed (bottom row) images using this process with 1, 4, 12 and 44 scan counts.

Experimental Setup: We tested four scan count sequences— $Cs = \{12, 44\}$, $Cs = \{6, 12, 44\}$, $Cs = \{1, 6, 12, 44\}$, $Cs = \{1, 4, 6, 44\}$ —for synapse and mitochondria detection purposes. Each classifiers h_c were built as in [7]. That is, we first extracted regularly spaced superpixels from which we computed both intensity histograms and steerable features. We then used 15 training images for each scan count C to train a different Support Vector Machine (SVM) classifier with an RBF kernel. In the $C = C_{\max}$ case, we used the original images and for all other $C < C_{\max}$, we used synthesized images obtained as discussed above.

Fig. 4 depicts the algorithm's behavior when attempting to locate mitochondria using the scan sequence $Cs = \{6, 12, 44\}$. From left to right, we show the complete set of regions evaluated with each scan count on the initial slice, S^1 . In red, we show $C = 6$, in green $C = 12$ and in blue $C = 44$. The ground truth regions are also shown. Additionally, for $C = 12$ and 44, we also display what regions from the previous step were evaluated.

Evaluation: Recall that our goal is to image at full resolution all the regions containing target structures—mitochondria and synapses in our experiments—while spending as little time as possible scanning the block faces. In this context,

Fig. 4. SRC Examples. Images depict the algorithm’s behaviour when attempting to locate mitochondria using the scan sequence $Cs = \{6, 12, 44\}$. From left to right, we show the complete set of regions evaluated with each scan count on the initial slice, S^1 . In red, we show $C = 6$, in green $C = 12$ and in blue $C = 44$. The ground truth regions are also shown. Additionally, for $C = 12$ and 44, we also display what regions from the previous step were evaluated.

the proper measure of success is the True Positive Rate (TPR) as a function of scanning time, which is plotted in Fig 5. The false positive rate is less relevant as false positives only cause irrelevant parts of the block being scanned, which implies no loss of information but a time penalty that the increased scanning time already reflects.

For the purpose of this evaluation, we consider the time cost of one imaging strategy to be the sum of the number of times each individual pixel is scanned. For simplicity’s sake, we normalize these numbers by the corresponding count when scanning the whole block at the maximum scan count. As a result, the times that appear in Fig 5 are numbers between 0 and 1.

From these results, we can see all the scan count sequences we tested provide a significant speed increase, mostly at a very small loss in TPR. By choosing the appropriate sequence we can establish more than five-fold speedups for an insignificant TPR loss. In practice, this means a neuro-scientist could examine and gather statistics for five times as many synapses in the same scanning time.

Note that the choice of which specific scan count sequence to use is not innocuous as it implies training different classifiers for each scan count value, some of which might be more appropriate than others. For example, the $\{1, 4, 6, 44\}$ sequence appears to outperform the $\{1, 6, 12, 44\}$ one, which may imply that additional research into optimizing these sequences might lead to further gains.

Fig. 5. Mean relative time and TPR for localizing Mitochondria and Synapses by tested methods. Each point indicates a specific scanning strategy. See text for details.

Also, the gain in time achieved by our approach is highly dependent on the type of organelle that must be found. In particular, targets that cover smaller surface areas, such as synapses (roughly 0.05% of the surface) allow higher time gains than for mitochondria (3 to 5% surface covered).

4 Conclusion

We presented an approach for speeding-up image acquisition when tasked with localizing specific structures in FIB-SEM imagery. It exploits the fact that low-quality images can be acquired faster than higher-quality ones and yet be sufficient for inference purposes. We have demonstrated greater than five-fold speed-ups at very little loss in accuracy in the context of mitochondria and synapse detection. Furthermore, the algorithm we propose is generic and applicable to many imaging modalities that allow trading quality for speed.

Acknowledgements. This work was funded in part by the ERC MicroNano Grant.

References

1. Knott, G., Marchman, H., Wall, D., Lich, B.: Serial Section Scanning Electron Microscopy of Adult Brain Tissue Using Focused Ion Beam Milling. *Journal of Neuroscience* 28(12), 2959–2964 (2008)
2. Kreshuk, A., Straehle, C., Sommer, C., Koethe, U., Knott, G., Hamprecht, F.: Automated segmentation of synapses in 3D EM data. In: 2011 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, March 30–April 2, pp. 220–223 (2011)
3. Veeraraghavan, A., Genkin, A., Vitaladevuni, S., Scheffer, L., Xu, S., Hess, H., Fetter, R., Cantoni, M., Knott, G., Chklovskii, D.: Increasing depth resolution of electron microscopy of neural circuits using sparse tomographic reconstruction. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 1767–1774 (June 2010)
4. Fibics, <http://www.fibics.com/>
5. Viola, P., Jones, M.: Robust Real-Time Face Detection 57(2), 137–154 (2004)
6. Sznitman, R., Jedynak, B.: Active Testing for Face Detection and Localization 32(10), 1914–1920 (2010)
7. Lucchi, A., Smith, K., Achanta, R., Lepetit, V., Fua, P.: A Fully Automated Approach to Segmentation of Irregularly Shaped Cellular Structures in EM Images. In: Jiang, T., Navab, N., Pluim, J.P.W., Viergever, M.A. (eds.) MICCAI 2010, Part II. LNCS, vol. 6362, pp. 463–471. Springer, Heidelberg (2010)
8. Kaynig, V., Fuchs, T., Buhmann, J.: Neuron geometry extraction by perceptual grouping in sstem images. In: IEEE Conference on Computer Vision and Pattern Recognition, pp. 2902–2909 (June 2010)
9. Kumar, R., Vazquez-Reina, A., Pfister, H.: Radon-like features and their application to connectomics. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, pp. 186–193 (June 2010)