

Nuclei Segmentation of Whole Slide Images in Digital Pathology

Dennis Ai

Department of Electrical Engineering
Stanford University
Stanford, CA
dennisai@stanford.edu

Abstract—Pathology is the study of the cause and effect of disease from tissue samples on microscope slides. The assessment of a tissue sample, which is inspected under a microscope, is inherently subjective, and can cause disagreement amongst pathologists. Therefore, tissue analysis is an area where objective quantification can be critically helpful in determining the right diagnosis or prognosis. Here, we evaluate a set of algorithms for nuclei detection and segmentation in hematoxylin and eosin (H&E) stained tissue samples.

Keywords—digital pathology, detection, segmentation, nuclei, hematoxylin and eosin stains

I. INTRODUCTION

Tissue samples, usually extracted from biopsies, surgeries, or excisions, are used for clinical diagnosis and prognosis for millions of patients each year. In 2014, Medicare, which covers approximately 30% of the United States population, reimbursed microscopic analysis of pathology samples nearly 20 million times, amounting to \$964 million [1].

This procedure usually involves a pathologist, who inspects a tissue sample under a microscope. From the cell morphology and tissue structure, he or she aims to understand the learn about the patient's disease and draw conclusions that lead to a diagnosis or prognosis.

However, this process is inherently subjective. As a result, a pathologist often calls upon his or her colleagues for a second opinion. In many cases, when that second pathologist views the same microscope slide, he or she may draw a different set of conclusions.

In parallel, the advent of the digital age has ushered in a new generation of digital pathology scanners, which are capable of whole slide imaging. That is, these automated scanners can provide a high-resolution image of an entire tissue sample in just a few minutes. There are two primary benefits of this technology. First, sharing a digital slide is much easier than sharing a physical one, and it is easy for

two pathologists in remote locations to compare analyses with one another.

Second, a digital slide paves the way for automated image analysis. Learning models and image processing algorithms can provide quantitative, objective determinations about a given tissue sample, such as the average size of the nucleus in a specific region of interest. This is especially important in the diagnosis of various forms of cancer, as tumor cells tend to be characterized by a large nucleus with an irregular size and shape [2].

While there are many modes of microscopy, such as brightfield, fluorescence, and phase contrast, these types of images all share a common need to begin analysis by performing detection and segmentation of individual cells. However, many of the most commonly used stains do not specifically stain the cell membrane – instead, stains such as hematoxylin (brightfield) or DAPI (fluorescence) dye the nuclei. Therefore, we begin with nuclei detection and segmentation as the initial and most fundamental step in digital pathology image analysis.

II. PREVIOUS WORK

There are a number of algorithms that have been utilized for the problem of nuclei detection. That is, given an image of a tissue sample, we would like to define a single point (x, y) that is located within the nucleus for each cell in the image.

The first class of algorithms utilizes the distance transform, coupled with the watershed transform, to detect the centers of nuclei. For example, Adiga *et al.* used a set of preprocessing filters such as directional diffusion before applying a distance transform in multispectral images of breast cancer tissue with the goal of analyzing genetic changes in cell nuclei that are far aware from the primary tumor [3].

A second class of algorithms relies on morphological operators, specifically the erosion operation. Yang *et al.* approach this problem by utilizing two sets of structuring elements that are 7×7 and 3×3 [4]. They iteratively apply erosions with the 7×7 coarse filters until a specified

threshold is reached, and then apply the 3x3 fine filters until another threshold is reached. The resulting markers then represent the nuclei in the image. What makes these techniques more complex to implement is that one needs to start with a binary image, so the original pathology slide needs to be binarized in some fashion.




Figure 1: The fine and coarse structuring elements used by Yang [4].

The H-minima and maxima transforms are two alternative methods that build upon morphology operations. These two transforms suppress regional minima / maxima that whose depth / height is smaller than a parameter h , respectively. The results from this operation are highly dependent on the value chosen for h , which can be iteratively determined. For instance, Cheng *et al.* continuously increase the value of h until region merging between two previously distinct nuclei occurs [5].

Shifting gears, the Laplacian of Gaussian (LoG) filter has also been used as a blob detector. There are numerous variations of this technique, such as the multi-scale version proposed by Al-Kofahi *et al.* to make the detector scale-invariant [6], or the elliptical Gaussian kernel put forward by Kong *et al.* that can detect rotationally asymmetric nuclei [7].

$$G(x, y) = Z \cdot e^{-(ax^2+2bxy+cy^2)}$$

Equation 1: A Gaussian kernel whose shape and orientation can be tuned by the parameters a , b , and c [7].

The final class of algorithms is the Hough transform, which was generalized to detect arbitrary shapes by Ballard [8]. This is especially useful for detecting nuclei, which are often circular or elliptical in pathology images. The Hough transform is usually used in conjunction with an edge detector, such as the Canny detector or Sobel filters. However, the downside is that the results of the

Hough transform can be noisy, and additional noise suppression steps are required.

III. METHOD

A. The Dataset

Our dataset consists of histology images of a normal pancreas from the Iowa Virtual Slidebox [9]. We extracted several regions of interest from a set of slides that were obtained at a 40x magnification.




Figure 2: A region of interest in full color from the digital slide of pancreas tissue.

B. Linear, Color-Based Pixel Classifier

We began by transforming the full-color pathology slide into a grayscale image, where each pixel reflected the likelihood of being a nuclei pixel.

To do this, we asked the user to select a rectangular region of interest within the image that consisted entirely of nuclei pixels, and another region of interest that consisted of non-nuclei pixels (e.g. cytoplasm). We then compute the RGB, HSV, and LAB color values for both the nuclei and non-nuclei pixels. These colors are used as the feature matrix in our linear regression classifier.

We then use our classifier to predict the likelihood that any pixel in the original image belongs to a nucleus of a cell. This process results in an image such as Figure 3.




Figure 3: The transformed nucleus "map" as predicted by the color-based pixel classifier.

Inspection of Figure 3 validates our hypothesis. The nuclei have been transformed into brighter regions, while the cytoplasm and other non-nuclear regions are black. However, while approximately three-quarters of the nuclei are distinct, the other one-fourth are overlapping, or are joined together by artifacts from our classifier.

If we analyze the original image, we will see that these artifacts are the byproduct of the hematoxylin stain which causes a purplish smear in the regions between nuclei. This is due to the fact that hematoxylin will bind to RNA in the ribosomes, which are located in the rough endoplasmic reticulum. Consequently, we need to apply additional image processing techniques to adjust for this biological reaction.

C. Blob Detection

Our next step is to utilize cell morphology – specifically the fact that nuclei are often circular or elliptical in nature. We first apply a Gaussian filter with $\sigma = 2$ to smooth the image since the nuclei have a grainy texture. Then, we apply a morphological opening operator on our grayscale image to detect circular blobs, which can be visualized in Figure 4.




Figure 4: The resulting blobs from the a morphological opening (left) and the results from finding the regional maxima (right).

However, at this point, the morphological opening picks up some of the smear discussed in the previous section. We can remove these false positives by identifying the regional maxima. This takes advantage of the property that nuclei are almost always brighter in the grayscale image than the smear that joins them together.

A connected components analysis, followed by the computation of the centroids, leads to a map of the “seed points” of the nuclei (Figure 5). While a more detailed analysis will be given in the Results section, the largest source of false negatives (i.e. nuclei that were not detected) are nuclei that overlap with one another. In these cases, only one of the overlapping nuclei are detected. This is the result of the application of the regional maxima algorithm. If there are multiple connected nuclei, the regional maxima algorithm will preserve the one with the highest intensity value.




Figure 5: The seed points of the nuclei (pink) overlaid on the original tissue sample (green).

D. Morphological Opening and Closing by Reconstruction

With a developed algorithm for nuclei detection, we pivoted to the problem of segmentation. To do this, we used a different set of morphological operators, namely the morphological reconstruction.

In contrast to a standard morphological opening or closing, morphological reconstruction:

1. Utilizes two images, a marker and a mask. In an opening-by-reconstruction, the marker is an eroded image, and the mask is the original image. In a closing-by-reconstruction, the complement of a dilated image is the marker, and the complement of the original image is the mask.
2. Instead of using a structuring element, the reconstruction is based on the connectivity of the pixels.
3. Dilations are repeated on the marker, until the contour of the marker fits inside the mask. Once the modified marker image no longer changes, the processing stops.




Figure 6: The nucleus map, after an opening and closing by reconstruction.

When we apply an opening-by-reconstruction followed by morphological closing-by-reconstruction, we get the result in Figure 6. The smearing of the hematoxylin dye is once again present in this image. In order to alleviate this, we can once again find the regional maxima. If we then overlay this an eroded version of the regional maxima (pink) on top of the gradient of the image (green) in Figure 7, we notice that while most of the cells are detected, a few of the cells are joined together.




Figure 7: The regional maxima overlaid on the gradient of the nucleus map.

To separate these cells, we will need use these maxima as the foreground markers for a watershed segmentation.

E. Segmentation using Marker-Controlled Watershed Transform

With foreground markers in hand, the other piece of the puzzle is the set of background markers, which we can compute using the following algorithm.

1. Binarize the image that has been morphologically reconstructed.
2. Calculate the distance map based on the inverse of this binary image, and negate the result (Figure 8).
3. Apply the watershed transform, and preserve the ridge lines, which are the pixels that have a value of zero. These are our background markers.




Figure 8: The distance map of the binarized image.

The final two steps of our algorithm include another morphological reconstruction that imposes regional minima where are foreground and background markers are, and a watershed transform on the gradient of the image, resulting in Figure 9.




Figure 9: The watershed segmentation of the pancreas tissue.

While the overall recall of the nucleus segmentation algorithm is not as high as one would like, the shapes of the segmentation appear to be relatively accurate.

IV. RESULTS AND DISCUSSION

A. Nuclei Detection

We performed a manual count of the detected nuclei in a randomly sampled region of the pancreas tissue slide, and obtained the results found in Table 1.

123 True Positives	20 False Negatives
9 False Positives	

Table 1: Confusion matrix for the nuclei detection algorithm.

In other words, 123 of our seed points were inside actual nuclei, 9 of the seed points were outside the nuclei, and we failed to detect 20 of the nuclei altogether. Therefore, the precision of our detection algorithm was 93.2%, the recall was 86.0%, and the overall accuracy was 80.9%.

Upon inspection of the opening-based blob detector (Figure 4, left), we can see that almost every nucleus was still present in the image. However, once we determined the regional maxima, we “lost” several of the nuclei. The reason for this is that the regional maxima preserves only

the pixels with the highest intensity value in a connected component. If several nuclei are connected to one another, then only one will be preserved.

One approach for improving the recall is to apply an iterative regional maxima operation. That is, after we find our initial set of regional maxima, we can subtract them from the blob image (Figure 4, left), and detect a new set of regional maxima based on this image. For each additional iteration, we will be able to find an additional nucleus in every group of overlapping nuclei. However, this will add the risk of finding additional false negatives, so the number of iterations must be kept at a limit.

On the other hand, there were a few false positives as well. Upon analysis of the false positives, we can see that they correlated to connected components with low intensity values. We can solve this problem by zeroing out any connected component whose average value falls below a certain threshold (e.g. 0.01).

We would expect that application of both these techniques would lead to a strong improvement in the precision, recall, and overall accuracy of the nuclei detection algorithm.

B. Nuclei Segmentation

Our nuclei segmentation algorithm consists of three input algorithms: foreground marker detection, background marker detection, and gradient computation. As the latter is fairly straightforward, we need to investigate the effectiveness of our foreground and background marker algorithms. Figure 10 illustrates this, showing both the foreground markers (black regions in the center of each cell), and the background markers (the black lines).




Figure 10: The gradient image, with the foreground markers and background markers (ridge lines) superimposed.

Our foreground markers appear to be largely accurate, as they are always found in the interior of the white edges of the gradient image. However, we note that our watershed ridge lines cut across the center of many of our

nuclei. This is the likely cause of the low recall that is evident in our final watershed segmentation.

One way to verify this is to modify the algorithm for calculating the distance map that is used to compute the watershed ridge lines. In the algorithm described in the Method section above, we negate the distance map of the inverse of the binarized image. If we simply compute the distance map of the binarized image, we get the result found in Figure 11.




Figure 11: Watershed segmentation using an alternative approach to calculating the distance map.

The recall of the algorithm is much improved, but the segmentation is unable to separate overlapping cells. This is the inherent tradeoff of this alternative approach to calculating the distance map.

As a result, we propose that further pre-processing or post-processing needs to be done on the distance map before using it to calculate the watershed ridge lines.

V. CONCLUSION

In this report, we have investigated the use of a supervised linear classifier to predict whether a pixel belongs to a nucleus based on its color. This is useful for a number of reasons.

1. It allows flexibility when dealing with different tissue types and staining protocols. For example, if a user is imaging breast or head and neck tissue, the same classifier would be able to adapt to changes in illumination, color schemes, and digital pathology scanners.
2. Through additional user feedback (e.g. the user identifies false positives and false negatives in the final segmentation), it is possible to refine the entire segmentation algorithm.

The linear classifier could be improved by a wider selection of image features, such as neighborhood metrics (e.g. the average color values of a 3x3 neighborhood). Additionally, it could be made more flexible if the user was able to select a non-rectangular region of interest.

We have also implemented and investigated algorithms for nuclei detection and segmentation and discussed areas for improvement.

While the above does not yet constitute a clinical-ready set of algorithms for nuclei detection and segmentation, there continues to be significant potential for digital image analysis to improve the practice of pathology, and the diagnosis and treatment of patients across the world.

REFERENCES

- [1] "Part B National Summary Data File," *Centers for Medicare & Medicaid Services*. [Online]. Available: <https://www.cms.gov/Research-Statistics-Data-and-Systems/Downloadable-Public-Use-Files/Part-B-National-Summary-Data-File/Overview.html>. [Accessed: 06-Dec-2016].
- [2] A. I. Baba, "TUMOR CELL MORPHOLOGY," *Comparative Oncology*, Jan-1970. [Online]. Available: <https://www.ncbi.nlm.nih.gov/books/NBK9553/>. [Accessed: 06-Dec-2016].
- [3] U. Adiga, R. Malladi, R. Fernandez-Gonzalez, and C. O. de Solorzano, "High-throughput analysis of multispectral images of breast cancer tissue," *IEEE Trans. Image Process.*, vol. 15, no. 8, pp. 2259–2268, Aug. 2006.
- [4] X. Yang, H. Li, and X. Zhou, "Nuclei segmentation using markercontrolled watershed, tracking using mean-shift, and Kalman filter in time-lapse microscopy," *IEEE Trans. Circuits Syst.*, vol. 53, no. 11, pp. 2405–2414, Nov. 2006.
- [5] J. Cheng and J. C. Rajapakse, "Segmentation of clustered nuclei with shape markers and marking function," *IEEE Trans. Biomed. Eng.*, vol. 56, no. 3, pp. 741–748, Mar. 2009.
- [6] Y. Al-Kofahi, W. Lassoued, W. Lee, and B. Roysam, "Improved automatic detection and segmentation of cell nuclei in histopathology images," *IEEE Trans. Biomed. Eng.*, vol. 57, no. 4, pp. 841–852, Apr. 2010.
- [7] H. Kong, H. C. Akakin, and S. E. Sarma, "A generalized Laplacian of Gaussian filter for blob detection and its applications," *IEEE Trans. Cybern.*, vol. 43, no. 6, pp. 1719–1733, Dec. 2013.
- [8] D. H. Ballard, "Generalizing the Hough transform to detect arbitrary shapes," *Pattern Recog.*, vol. 13, no. 2, pp. 111–122, 1981.
- [9] "Iowa Virtual Slidebox," Iowa Department of Pathology. [Online]. Available: <http://www.path.uiowa.edu/virtualslidebox/>. [Accessed: 06-Dec-2016].