SPICE: Superpixel Classification for Cell Detection and Counting

Oman Magaña-Tellez[1,2], Michalis Vrigkas[2], Christophoros Nikou[2] and Ioannis A. Kakadiaris[1,2]
[1]*Department of Biomedical Engineering, University of Houston, Houston, TX 77004, U.S.A.*
[2]*Computational Biomedicine Lab, Department of Computer Science, University of Houston, Houston, TX 77004, U.S.A.*

Keywords: Superpixel, Histological, Cell Detection, Cell Quantification, Random Forests.

Abstract: An algorithm for the localization and counting of cells in histopathological images is presented. The algorithm relies on the presegmentation of an image into a number of superpixels followed by two random forests for classification. The first random forest determines if there are any cells in the superpixels at its input and the second random forest provides the number of cells in the respective superpixel. The algorithm is evaluated on a bone marrow histopathological dataset. We argue that a single random forest is not sufficient to detect all the cells in the image while a cascade of classifiers achieves higher accuracy. The results compare favorably with the state of the art but with a lower computational cost.

1 INTRODUCTION

Histopathological image analysis plays an important role in the diagnosis of numerous pathologies ranging from infectious diseases to cancer (Oguz et al., 2016). However, the traditional method for analyzing histopathological images is a tedious and time-consuming task given the typically large number of cells contained in the image as well as the numerous images to be analyzed, which can lead to considerable inter-observer variability as well as irreproducible results (Andrion et al., 1995; Ismail et al., 1989). Therefore, the demand for computer-aided analysis is high (Demir and Yener, 2005; Zhang et al., 2014) and has seen an increased effort in research during the previous decades. Among the major difficulties in the application of image analysis methods to cell images are the non-uniform staining, blurring due to defocussing, and the existence of overlapping cells (Demir and Yener, 2005).

There are methods designed for estimating the location of the cells in the image (Kainz et al., 2015; Zhang et al., 2014). These algorithms use different techniques such as a score map with the probability of location or an arbitrary image segmentation through correlation clustering. On the other hand, the methods in (Benali et al., 2003; Sjostrom et al., 1999) quantify the number of cells; the first method uses a clustering followed by a binarization of the image and the second method uses a three layer neural network fed by structural information. Kainz *et al.* overcame one of the main issues in histopathological images, the differentiation of cells from background structures, by using a probability score map to indicate where a cell is more likely to be located (Kainz et al., 2015). Even though this method exhibited good results, the interference of undesired structures is still present. Moreover, the need for defining a threshold for the distance between a true cell location and the response from a trained classifier is also an important issue.

In this paper, the SPICE (SuperPIxel classification for Cell dEtection and counting) algorithm is proposed, which is an algorithm for the localization and quantification (total number) of cells in histological images. Our method uses a superpixel presegmentation of the image and a sequence of random forests for classification. The first random forest is a binary classifier which determines if the superpixel at its input contains any cells. A second random forest, which is a multiclass classifier, determines how many cells are present at the superpixels provided by the same pre-segmentation of the image. Both classifiers can work independently. However, the experimental evaluation indicated that more accurate results can be obtained if they are applied sequentially. An advantage of SPICE is the use of superpixels in the segmentation, since it provides the extracted features with a more compact and more representative modeling of the cells (e.g., using the color and the shape of the superpixels). Also, we demonstrate that for the learning stage, a low computational cost is capable of giving a high detection accuracy, which is favorably compa-

Magaña-Tellez, O., Vrigkas, M., Nikou, C. and Kakadiaris, I.
SPICE: Superpixel Classification for Cell Detection and Counting.
DOI: 10.5220/0006508304850490
In *Proceedings of the 13th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2018) - Volume 4: VISAPP*, pages 485-490
ISBN: 978-989-758-290-5

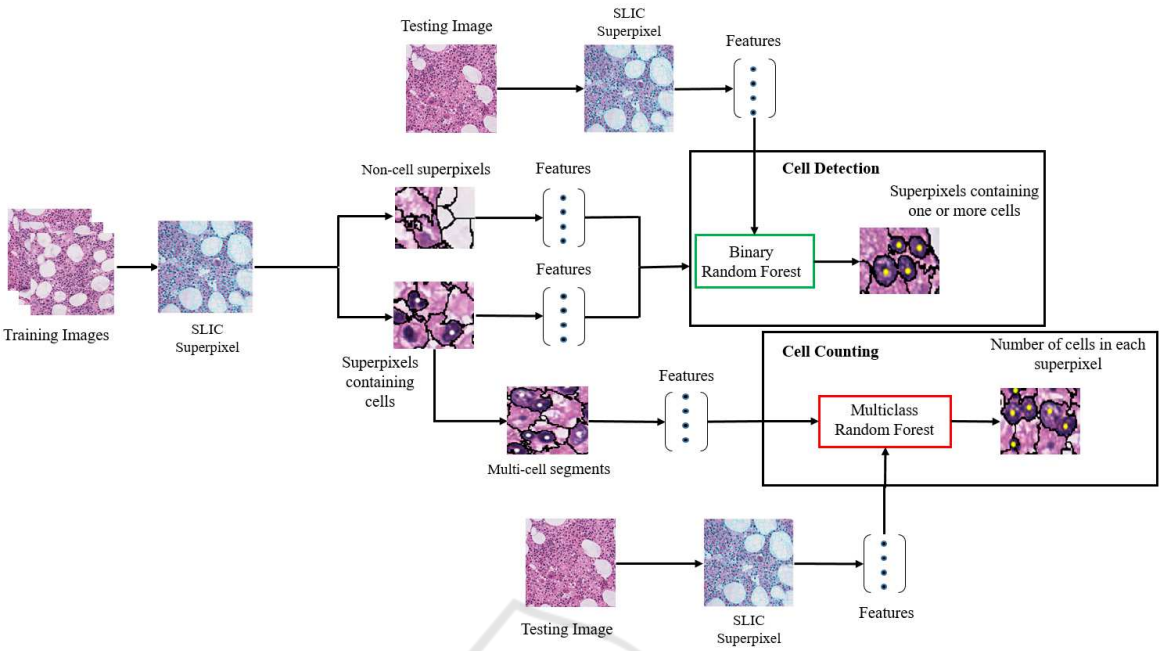

Figure 1: Depiction of SPICE, detailing the training and the testing stages for both cell detection and cell counting.

red with state-of-the-art methods (Kainz et al., 2015) using the same dataset.

2 CELL DETECTION AND COUNTING

Histological images are acquired by first obtaining a tissue sample from the patient, then the sample is dehydrated and encased in paraffin in order to preserve the tissue. Finally, a staining is applied to the sample in order to highlight the structures of interest. The most common stain used is Hematoxilin and Eosin (H&E).

Classification algorithms for cell detection on histological images commonly use a sliding window over the image in order to extract features from it. However, this can lead to several issues to be addressed such as the increase in computational cost, the need for determinate the size of the window and the eventual misclassifications at the borders of the image. The method proposed herein is based on the segmentation of the image into superpixels. Each superpixel is represented by a feature vector which is then forwarded to a random forest classifier. The outline of the proposed method is summarized in Fig. 1.

In the first step of SPICE, the image is divided into perceptually meaningful regions. This is achieved using the SLIC superpixel segmentation algorithm (Achanta et al., 2010), an algorithm that clusters pixels in the combined five-dimensional color (CIELAB) and location features to efficiently generate compact, nearly uniform segments. We decided to use the SLIC algorithm since it is considered one of the fastest, state-of-the-art algorithms (Achanta et al., 2012) and it only needs the number of superpixels as parameter. The pre-segmentation of the image into superpixels, reduces significantly the computational cost and time of SPICE, by eliminating the tedious process of sliding a window through the image to obtain the sections for feature extraction. The superpixel segments aggregate regions with similar characteristics which facilitates the feature extraction task. Also, by organizing the image into similar regions there is a higher probability that an entire cell or a cluster of cells is contained in a single segment. The number of superpixel segments is selected according to the size of the image and the size of the cells. Selecting a large number of segments in a small image can cause the cells to be partitioned into multiple segments. On the contrary, selecting a small number of segments may lead to superpixels containing both cell and background information.

For the training step, a number of features are extracted for each superpixel to represent the underlined segment. The features we used were the following: RGB intensity channels, magnitude of oriented gradients, first and second oriented gradients, LUV intensity channels and histogram of oriented gradients. For every superpixel, the mean and the standard deviation of the respective features are computed, except

Algorithm 1: SPICE: Training for cell detection and counting.

Input : *N* Images, Cell center coordinates
Output: Detection and Counting Random Forests

1 **for** *All Images in Set* **do**
2 Segment the image using the SLIC superpixel algorithm (Achanta et al., 2012).
3 Extract a 31-dimensional feature vector for each superpixel.
4 Train the binary random forest for cell detection.
5 Train the multiclass random forest for cell counting.
6 **end**

Algorithm 2: SPICE: Testing for cell location and counting.

Input : An image
Output: Locations of cell centers and number of cells

1 Segment the image using the SLIC superpixel algorithm (Achanta et al., 2012).
2 Extract a 31-dimensional feature vector for each superpixel.
3 Apply the feature vectors to the binary random forest to indicate the presence of cells in a superpixel.
4 Apply the feature vectors to the multi-class random forest to obtain the number of cells in a superpixel.

for the histogram of oriented gradients, where a 9-dimensional vector of the gradient orientations weighted by their amplitudes is computed. We decided to use the mean and standard deviation of the features to have a more robust representation of the structure and color of the cell compared to the background or other undesired structures. The concatenation of these features yields a 31-dimensional feature vector.

We decided to use the RGB channels for the color information of the cell, since H&E stains the cell nucleus with blue color and the cytoplasm with pink color, therefore it is straightforward to have the blue color as a feature for the pixels to indicate a high probability of cell presence. The gradient information obtained by SPICE is used for the representation of the shape information of the cell. The LUV channels, such as the RGB channels, provide information of the color of cells and the background with the advantage that these features are device (microscope) independent and they may not be modified.

The cell detection algorithm is a binary random forest classifier that determines if the superpixel at its input contains any cells or background. At this point, the number of cells in the segment does not play any major role, since we are focusing only on the presence or absence of cells in the image. Therefore, the next step of the algorithm consists in determining the number of cells in each superpixel. A multi-class random forest classifier is employed using the number of cells present in the segment as the corresponding label. We decided to limit the number of classes to four, in order to avoid the potential problem of unbalanced data, since it is relatively rare to have more than three cells clustered in the same superpixel. The overall procedure for training and testing is summarized in Algorithms 1 and 2, respectively.

3 EXPERIMENTAL RESULTS

The algorithm was evaluated on the dataset introduced in (Kainz et al., 2015). The dataset consists of 11 images of $1,200 \times 1,200$ pixels of healthy bone marrow from eight patients and their respective ground truth image. Based on the size of images of the dataset and the expected cell sizes, we segmented the images into $1,000$ superpixels.

We performed a set of experiments to test the impact of the number of superpixels in the image. We performed a number of experiments with both a small as well as a large number of segments. The number of segments plays a crucial role for the quantification of cells, as selecting a small number of segments would result in increased false positives, while a large number of segments would reduce considerably the detection of cells in the image. Based on the bone marrow cell image dataset (Kainz et al., 2015), we selected the number of superpixels by cross-validation and set it to the value of $1,000$ as this pre-segmentation provides a detection rate closer to the ground truth for the validation set (Fig. 2). Nevertheless, this parameter has to be cross-validated in the case of a different type of cell images. This is perhaps the caveat of the method but by performing this cross validation we can ensure that the number of segments will give to the classifier the strongest features.

The number of classes in the multi-class random forest was set to four, which represents the presence of 0, 1, 2, and 3 or more cells in a superpixel segment. Using four labels handles the issue of unbalanced data in the training step of the algorithm as the dataset in (Kainz et al., 2015) contains too few cell clusters with more than four cells. Moreover, in the second stage, we also had a label of zero cells in order to include

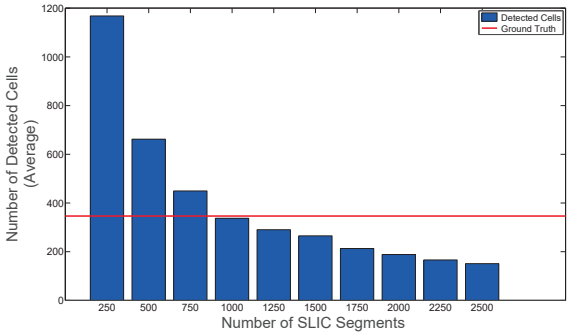

Figure 2: Impact of the number of superpixels in the quantification of cells.

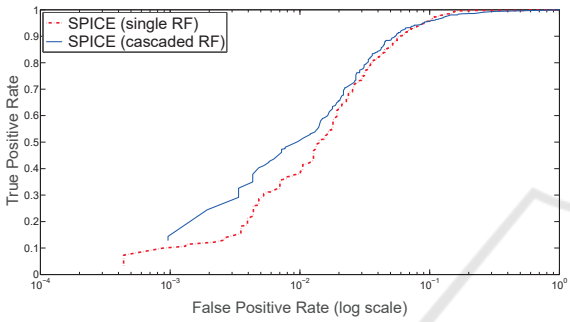

(a)

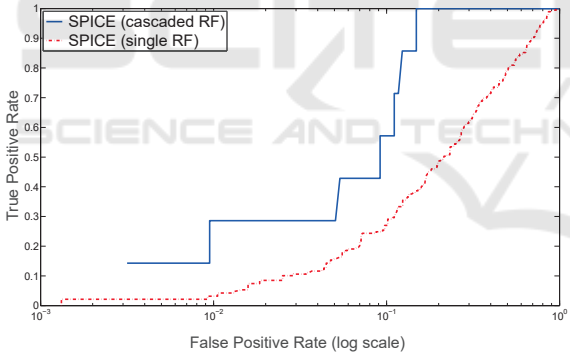

(b)

Figure 3: ROC curves of (a) the binary classifier and (b) multi-class classifier of SPICE with a single and cascaded random forest. The curve for the single random forest in (a) was generated by comparing its final classification result in a "hit or miss" sense, without taking into account the number of cells detected.

background superpixels that could have been missed by the first classifier. Learning the background label using this additional random forest improves the robustness of the algorithm.

In the second set of experiments, we examined the consistency of the method with respect to the number of trees in the random forests. In order to compare the classifiers more efficiently the area under the curve (AUC) is used as a score for their performance

Table 1: Area under the ROC curve.

Classifier	SPICE-CRF	Kainz *et al.*	SPICE-SRF
Cell Detect.	**97.34%**	90.5%	97.29%
Cell Count.	**93.67%**	N/A	71.43%

(Fawcett, 2006). A 3-fold cross validation was employed to determine the number of trees in both classifiers. The results indicated that varying the number of trees between 50, 150 and 200 trees yields similar areas under the ROC curve (AUC), namely 97.21%, 97.27% and 97.27% respectively. From this experiment, we concluded that a relatively low number of trees is sufficient to obtain a high classification accuracy as increasing the number of trees does not have a significant impact on the outcome of the method. Therefore, for a faster performance and less computational power we decided to use 50 trees in the next experiments.

An important question arising from SPICE is why one has to apply two random forests sequentially instead of a single one that would do the same classification. To clarify this issue, the SPICE algorithm is compared with its variant which uses a single random forest that classifies the superpixels directly with respect to the number of cells they contain. The ROC curves are shown in Fig. 3. Figure 3a shows the ROC curve for the binary classifier. The blue curve corresponds to the performance of the detection classifier of SPICE using a cascaded random forest while the red curve is the same classifier but in a single random forest configuration. The curve for the single classifier is generated by the cells detected at its output in a binary ("hit or miss") sense. Since there is only a single stage, these curves are relatively similar. However, the difference is clear in Fig. 3b where the multi-class random forest applied after the first binary random forest outperforms the straightforward classification of the image. This happens because classification errors from the previous stage are carried over and affect deeply the multi-class classification.

The experiment proved that SPICE is capable of rivaling the state-of-the-art method for cell detection. The algorithm is capable of finding the cells in the superpixel segments of the image with a high accuracy and creates a window around them to indicate the results. Figure 4 illustrates representative results of cell detection using the proposed SPICE algorithm.

In the last set of experiments the SPICE algorithm is compared with the algorithm developed in (Kainz et al., 2015) for the localization of cells in the image. The respective ROC curves are depicted in Fig. 5 and the overall accuracies are shown in Table 1. Since the method presented in (Kainz et al., 2015) is only a cell localization algorithm, this experiment concerns only the localization and not the number of cells.

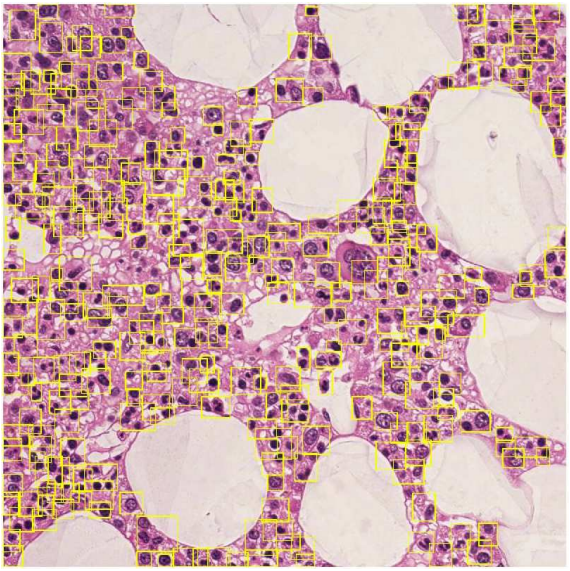

Figure 4: Representative result of SPICE cell detection on an image of the database presented in (Kainz et al., 2015).

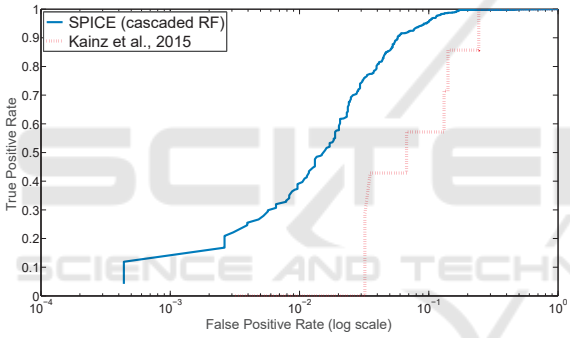

Figure 5: Comparison between the SPICE algorithm and the method presented in (Kainz et al., 2015).

Despite its success and good performance, the proposed approach has one limitation that is related on how the number of the superpixels for the segmentation of the images are selected. This is now performed by heuristic cross-validation and typically set to 1,000. Although this has to be adjusted for every new type of image by performing cross-validation to estimate the number of superpixels, we can ensure that the different segments will give to the classifier representative features.

The goal of the proposed method is to detect isolated cells in an image and obtain their features for better and stronger classification of bone marrow histopathological images. An interesting and possible extension of this work would be in hematological diseases and more particularly in the analysis of 2D immunotherapy images (Rosas-Taraco et al., 2011; **?**) or in detecting cells in Pap smear images (Plissiti et al., 2015), where superpixels can also be used so that each image can be tessellated into approximately equally sized subregions, presenting homogeneous intensity characteristics. However, different types of images may exhibit different properties because cells may be highly overlapping or the staining process may be different. This implies that a different set of features may need to be extracted for these type of images. Regardless of the type of the features that may be used for each image, the SPICE algorithm is general and can easily be adapted to detect and localize cells in such types of images.

4 CONCLUSION

A method for cell detection and quantification in histological images that uses a superpixel segmentation along with a two stage random forest classification is presented. The method was successfully evaluated in terms of AUC and favorably compared to a state-of-the-art algorithm for cell detection. The main advantage of the proposed method is that it provides a flexible way for the simultaneous detection and counting of cells in histopathological images using a cascade of classifiers. The results indicated that the proposed SPICE algorithm ameliorates the classification accuracy by approximately 7% with respect to the state of the art (Kainz et al., 2015). As future work, we plan to extend the algorithm to detect 3D cells, where the difficulty consists in determining the appropriate features.

ACKNOWLEDGMENTS

This work was funded in part by the Mexican National Council for Science and Technology (CONACYT) scholarship, 668902, by the European Commission (H2020-MSCA-IF-2014) No. under grant agreement No. 656094, and by the UH Hugh Roy and Lillie Cranz Cullen Endowment Fund. All statements of fact, opinion or conclusions contained herein are those of the authors and should not be construed as representing the official views or policies of the sponsors.

REFERENCES

Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., and Susstrunk, S. (2010). SLIC superpixels. Technical Report 149300, EPFL, Lausanne, Switzerland.

Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., and Susstrunk, S. (2012). SLIC superpixels compared to

state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11):2274–2282.

Andrion, A., Magnani, C., Betta, P., Donna, A., Mollo, F., Scelsi, M., Bernardi, P., Botta, M., and Terracini, B. (1995). Malignant mesothelioma of the pleura: Interobserver variability. *Journal of Clinical Pathology*, 48(9):856–860.

Benali, A., Leefken, I., Eysel, U., and Weiler, E. (2003). A computerized image analysis system for quantitative analysis of cells in histological brain sections. *Journal of Neuroscience Methods*, 125(1):33–43.

Demir, C. and Yener, B. (2005). Automated cancer diagnosis based on histopathological images: A systematic survey. Technical Report TR-05-09, Rensselaer Polytechnic Institute, Troy, NY.

Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8):861–874.

Ismail, S., Colclough, A., Dinnen, J., Eakins, D., Evans, D., Gradwell, E., Sullivan, J., Summerell, J., and Newcombe, R. (1989). Observer variation in histopathological diagnosis and grading of cervical intraepithelial neoplasia. *British Medical Journal*, 298(6675):707–710.

Kainz, P., Urschler, M., Schulter, S., Wohlhart, P., and Lepetit, V. (2015). You should use regression to detect cells. In *Proc. Medical Image Computing and Computer-Assisted Intervention*, pages 276–283, Munich, Germany.

Oguz, O., Akbas, C., Mallah, M., Tasdemir, K., Guzelcan, E., Muenzenmayer, C., Wittenberg, T., Uner, A., Cetin, A., and Atalay, R. (2016). Mixture of learners for cancer stem cell detection using CD13 and H&E stained images. In *Proc. SPIE Medical Imaging*, pages 97910Y–1–97910Y–16, San Diego, CA.

Plissiti, M. E., Vrigkas, M., and Nikou, C. (2015). Segmentation of cell clusters in Pap smear images using intensity variation between superpixels. In *Proc. 22nd International Conference on Systems, Signals and Image Processing*, pages 184–187, London, UK.

Rosas-Taraco, A., Higgins, D., Sánchez-Campillo, J., Lee, E., Orme, I., and González-Juarrero, M. (2011). Local pulmonary immunotherapy with siRNA targeting TGFβ1 enhances antimicrobial capacity in mycobacterium tuberculosis infected mice. *Tuberculosis*, 91(1):98–106.

Sjostrom, P., Frydel, B., and Wahlberg, L. (1999). Artificial neural network-aided image analysis system for cell counting. *Cytometry*, 36:18–26.

Zhang, C., Yarkony, J., and Hamprecht, F. (2014). Cell detection and segmentation using correlation clustering. In *Proc. Medical Image Computing and Computer-Assisted Intervention*, pages 9–16, Boston, MA.