An Interactive Learning Approach to Histology Image Segmentation

Michaël Derde[a] Laura Antanas[a] Luc De Raedt[a] Fabian Guiza Grandas[b]

[a] *Katholieke Universiteit Leuven, Department of Computer Science*
[b] *Katholieke Universiteit Leuven, Laboratory of Intensive Care Medicine*

Abstract

Histology image analysis using computer-aided diagnosis systems has become increasingly important during the last years. One reason is the need to alleviate the heavy workload of medical experts. In this paper, we introduce a general purpose framework which is able to solve histology analysis problems that are not restricted to a specific type of tissue or task, exploit local information in microscopical images, interact with medical experts and interatively consider direct user feedback. The framework is general enough to learn models that can adapt to several learning tasks and can detect several types of medical interest regions. We evaluate our framework on real-world datasets collected from patients in the intensive care unit. We considerably outperform image processing techniques commonly used in such medical imaging tasks.

1 Introduction

Histology is the anatomical study of the microscopic structure of tissues. It is regarded as a gold standard for clinical diagnosis of diseased tissue (e.g., cancer) and for the identification of therapy effects [10]. Histological analysis is performed by examining a thin section of tissue under a microscope [13, 20, 16], after applying a sequence of procedures for tissue preparation: fixation, dehydration, clearing, infiltration, embedding, sectioning and staining [14]. It reveals information about cells and tissue with a high level of detail. Despite the great care taken in their preparation, histology images are prone to several artifacts, e.g., folding of the tissue section, overlap among cell boundaries, noise introduced by the microscope or slides, blurry sections, etc. As a result, analysis of histology tissues remains most of the time a manual endeavor which relies heavily on the expertise of the medical expert. This manual work is however very time consuming and prone to subjective interpretation. Therefore, computer-assisted diagnosis (CAD) systems are becoming crucial in histology analysis, as they could automatically identify regions of medical interest. Their main advantage is the ability to provide immediate results in a consistent and objective manner, thereby reducing the workload of the medical experts.

With few exceptions [24], current image analysis tools focus on specific tasks, such as nuclei and cell counting, and lack the flexibility of dealing with a variety of tissue types that might be of interest to medical experts. Moreover, most systems make use of traditional image processing approaches such as global thresholding, region growing, region splitting and merging, and active contours (for more details see [18, 28]). Their main drawback is that they fail to account for local variations (e.g., brightness, staining intensity) within a single image that are introduced by the microscope or lightening conditions.

In this paper we present a new CAD framework which learns to automatically detect regions of interest in histology images. It overcomes drawbacks of the current approaches by combining an interactive learning technique that adapts to the specific user-defined medical task, with a local approach that takes into account local variations of particular regions of interest. Instead of the basic supervised learning paradigm in which the expert is asked to label examples and then a predictor is learned from these targets, without other explicit interaction, our framework uses expert knowledge to interactively feed training instances to the learning system. Once a new instance has been added by the expert, a completely new model is built from scratch, as a new supervised learning step. In this way, by changing the learning targets after each iteration, our framework can incoporate real-time feedback for the current model predictions, reducing training time and data. Once the model has been trained, it can be used to automatically process any amount of images. We formalize the supervised learning problem as a regression task and we employ regression trees to represent

the learned model. The training instances used to build the model are slices sampled from full histology images. They are segmented by the medical expert via local thresholding. The user-defined thresholds become regression targets in the learning phase.

In the remaining sections we present related work in automated analysis of histology images and describe the methodology employed by our interactive framework. Next, we present an experimental evaluation of our framework on real-world data collected from patients in the intensive care unit and discuss the results. Finally, we provide concluding remarks and perspectives for future work.

2 Related work

This section discusses related work on automated analysis of histology images. A crucial task in histology analysis is to automatically segment the histology image in background regions and regions of interest (e.g., nuclei, cells, stained regions). A comprehensive survey of image segmentation techniques is presented in [18, 8]. A popular and simple segmentation approach is intensity or color thresholding, which can be further described as being either bi-level or multi-thresholding. Bi-level refers to the use of only one threshold value, where image pixels below (above) the value are considered as background (region of interest). The use of several thresholds allows multi-thresholding to segment an image into more than two types of regions [21]. Alternatives to thresholding techniques that were successfully applied to histological segmentation include perceptual grouping [27], region growing [19], fuzzy clustering [2], active contours [4] and energy-based methods [25]. These alternatives, however, are either too problem specific or too demanding computationally for a fast interactive framework.

Recently, considerable work on histological segmentation has focused on fully automated approaches that rely on machine learning techniques. They allow for large amounts of data to be processed. Based on descriptive features of the data interesting patterns are discovered. The extraction of meaningful features from histology images has been addressed in [1]. We employ such features in our framework, yet we also consider different ones. Some related machine learning approaches have been employed for histology images. One of them uses Markov Random Fields in a Bayesian formulation and has been employed to segment cancerous structures [29]. Another approach uses a bag of local Bayesian classifiers to classify pixels as belonging to cells or not [31] and thus, segment histology images. Finally, the work in [7] uses random forests to classify pixels as belonging to a fixed set of predifined clasess. Different from these, our approach employs regression trees to learn meaningful thresholds which are then used to segment the image.

Other related techniques that incorporate the human expert into the learning loop are relevance feedback, preference elicitation and active learning. Relevance feedback improves current predictions by considering direct user feedback with respect to the relevance of past reactions [32]. Preference elicitation computes preferences from utility function estimation [6] and active learning asks the expert to analyze only instances that would be the most informative for the learning task [12]. Because the expert dynamically analyzes the instances, a single system can perform several learning tasks always using the same pre-computed set of features. While these techniques have been utilized in computer vision [23, 11], they remain largely unused in automated analysis of histology analysis. Differently, our framework uses expert knowledge to feed new training instances to the learning system, without specifying their degree of relevance and building, at each iteration, a completely new model from scratch.

Several software solutions are available for histology analysis. They include CellProfiler [5], CellTracer [30], CellTracker [22] and Imago [15]. These tools provide automatic segmentation of cells and nuclei and report statistics in similar terms about the image. They allow for batch processing, albeit they can not adapt parameter settings per individual image which remain static throughout the entire batch evaluation. Imago also provides an interactive segmentation method, yet it is restricted to bi-level thresholding. Thus, current available solutions suffer from two major drawbacks. First, they focus on a specific task, such as cell or nuclei segmentation, which hinders their use for other regions of interest. Second, they make use of global approaches both in analyses of one image and in batch processing, which means that varying lighting conditions or clinically relevant variations in staining intensity are largely ignored. Closely related is the work in [24], a learning and segmentation toolkit for image classification and segmentation which works iteratively in a similar way. However, differently, we propose a regression-based approach.

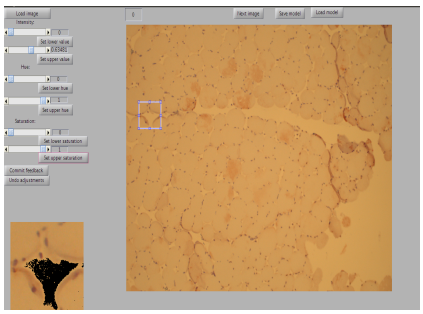

Figure 1: GUI interface. A slider is used to manually segment the regions of interest.

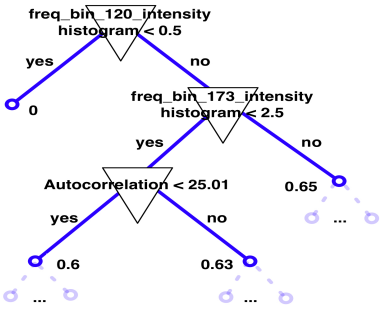

Figure 2: Regression tree model for nuclei detection in muscle tissue. The outcome is the local real-valued intensity threshold prediction.

3 Approach

In this section we describe our learning approach which overcomes the drawbacks in current approaches for detecting generic regions of interest in medical images. We start from randomly generated slices from the histology image. These slices are sections or crops of different sizes that are assumed independent from each other and that do not overlap. Our goal is to find the optimal segmentation of the areas of medical interest for each of these slices. We keep the well-known thresholding approach [17], however instead of manually selecting thresholds by considering the peaks of histograms for different image channels, we automatically learn a model to predict these thresholds.

There are several advantages of this approach. One is that the prediction relies totally on a set of features of the image slice which are the same independently of the problem considered. Another advantage is that the local threshold is objectively generated by the model according to the slice features and not on a subjectively selected threshold that is usually found between adjacent peaks of a histogram. Finally, there may be different image targets to threshold (e.g., staining intensity and hue) depending on the problem at hand. Employing a method which can automatically predict optimal thresholds for all targets considered makes our approach problem independent also with respect to the learning targets.

3.1 Interactive data acquisition

One way commonly used in medical image processing to obtain regions of interest is to impose thresholds on image parameters such as the intensity, hue and saturation channels. By considering all channels, our approach is more precise and general for more complex problems. If solely intensity is used, making distinctions between two regions that have similar RGB values is difficult, if not impossible, while in the hue channel they may be more disjoint, allowing for a better segmentation. Therefore, we consider a *multi-target* setting, where the goal is to threshold not only the different image channels at the same time, but even more complex targets for each channel so as to obtain an even more precise segmentation.

Often, medical specialists need to identify more complex tissue components than nuclei or cells. These may include particular stained regions. We propose a general purpose *interactive* framework which interacts with the user and can adapt to specific problems by learning from user feedback. This means that the domain expert is able to choose which examples to label such that an optimal model for the specific task of interest is built. In particular, the medical expert interacts with the framework and presents to the algorithm predictions for each image slice. This interaction is performed via a slider which is properly adjusted by the user in a graphical interface (as shown in Figure 1). Once a new instance has been segmented by the expert, a completely new model that incorporates the new user feedback is built from scratch. The newly trained model segments the next user-selected instance, and if needed, the expert further refines the segmentation, by providing additional feedback. The final model can then be used to batch process the remaining images.

3.2 Problem definition

We formalize the detection of areas of medical interest as a supervised learning problem: given a set of image slices $X = \{\mathbf{x}_i\}_{i=1}^n$ labeled with classes $Y = \{\mathbf{y}_i\}_{i=1}^n$, we want to find a model so that the probability of error is minimal when predicting $\mathbf{y} \in Y$ for a new $\mathbf{x} \in X$. The goal is then to learn from the dataset

$D = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$ a *prediction function* $\bar{f}$ which maps image slices X to segment labels Y. In our setting a segment label $\mathbf{y}$ is composed of real-valued threshold variables. This makes the function $\bar{f}$ a regression model. $\bar{f}$ belongs to a hypothesis space $H = \{f(\cdot)|f : X \rightarrow Y\}$ and is evaluated to minimize a loss function L which measures the similarity between two outputs.

$$\bar{f} = \arg\min_{f \in H} L(f, D) = \arg\min_{f \in H} \sum_{i=1}^{n} L(\mathbf{y}_i, f(\mathbf{x}_i)),$$

where L denotes the error or loss of f on example $(\mathbf{x}_i, \mathbf{y}_i)$ and $L(f, D)$ denotes the collective loss of f on the training set D. We expect the loss to be close to zero when a pattern is detected, as it measures the discrepancy between the output of the prediction function and the correct output value. When a new image slice is presented, the target output is not available and then $\bar{f}$ is used to predict $\mathbf{y}$ for the given input $\mathbf{x}$. This implies a multi-target setting, however, in this work, we consider each $y_i \in \mathbf{y}$ an independent target where $y_i \in \Re$. We build the multi-target regressor by learning multiple independent regression functions, one for each target. A natural future investigation is whether a multi-target or *structured prediction* setting, in which dependencies between the multiple targets are considered during training/testing (i.e. the level of intensity can influence the saturation prediction), can improve results.

3.3 Regression trees

One way of defining the regression model is as a linear function where a real-valued dependent variable y_i is modeled as a linear function of a real-valued independent variable $\mathbf{x}$ plus noise. In linear regression f_i is a global model, where there is a single predictive formula holding over the entire data-space. However, when the data is high-dimensional with non-sparse features which interact in complicated, nonlinear ways, assembling a single global model is not the best approach. One way to consider nonlinear regression is to partition the space into smaller regions which can be divided again until it fits the training data.

Regression trees use the tree to represent such recursive partitions and predict real-valued outcomes. Each of the terminal nodes (or leaves) of the tree represents a cell of the partition, and has attached a constant estimate of y which applies in that leaf only. That is, given the points $(\mathbf{x}_i, y_i), (\mathbf{x}_2, y_2), \ldots, (\mathbf{x}_c, y_c)$ are the samples belonging to the leaf node l, the model for l is $\hat{y} = \frac{1}{c}\sum_{i=1}^{c} y_i$ (or the sample mean of y_i in the leaf). A point $\mathbf{x}$ belongs to a leaf if $\mathbf{x}$ falls in the corresponding region of the partition. In our defined problem (as in Section 3.2), threshold values are predicted by traversing the regression tree until a leaf node is encountered, and the leaf outcome value $\hat{y}$ is assigned to the unseen instance. The interior nodes are labeled with tests, and the edges or branches between them labeled with the answers. We start at the root node of the tree and apply a sequence of tests in the tree about the features to figure out the prediction leaf. Which test is performed next depends on the answers to previous questions. Figure 2 shows a regression tree for muscle nuclei detection learned with our framework. The inner nodes in the tree represent tests on certain feature values, while the leaf contains the outcome. The outcome is the real-valued intensity threshold needed to segment the image slice. In this case each test refers to only a single attribute, and has a yes or no answer, e.g., "Is Autocorrelation < 25.01 ?".

Regression trees offer several advantages when compared to alternative methods (i.e., SVMs [3]): (i) fast algorithms exist to learn the trees and make predictions; this minimizes delays during interaction with the medical expert; (ii) the obtained models are interpretable which might help in gaining medical insight of the distinguishing properties of the identified regions; (iii) deals well with missing data; as even though a path to a prediction leaf might be unreachable, still a prediction can be made by aggregating predictions of all leaves in the reachable sub-tree. This property could be exploited to still provide predictions in severely damaged tissue samples for which not all features can be evaluated; (iv) the model obtained gives a ridged response, so it can work when the true regression surface is not smooth.

Learning We start building the tree by finding the one binary test which maximizes the information we get about y. One commonly used *splitting criterion* in regression is the sum of squared errors, also employed by our work; this gives us our root node and two children nodes. At each child node, we repeat the same procedure. The sum of squared errors for a tree T is

$$S = \sum_{l \in leaves(T)} \sum_{i \in l} (y_i - \hat{y}_l)^2,$$

Feature	Description
Haralick features	based on the gray-level co-occurrence matrix (GLCM) of an image
Total coarseness	refers to notable variations of the grey levels
Coarseness histogram	characterizes the distribution of the coarseness
Contrast	the range of the pixel intensities
Directionality	indicates whether the image favors a certain direction
Homogeneity	closeness of the GLCM elements distribution to the GLCM diagonal
Gray level statistics	refers to the pixels gray-level values such as the average, the median, standard deviation, minimum and maximum values
Color moments	characterize the distribution of the color in an image
Intensity histogram	shows the number of pixels in an image at different intensity values
Hue histogram	shows the number of pixels in an image at different hue values
Saturation histogram	shows the number of pixels in an image at different saturation values
Morphological features	include properties of the slice such as its width, height and polar coordinates relative to the centre of the full image

Table 1: Features used to characterize an image slice and their descriptions.

where $\hat{y}_l$ is the prediction for leaf l. We can re-write this as $S = \sum_{l \in leaves(T)} n_l \cdot V_l$, where V_l is the within-leave variance of leaf l and n_l is the number of examples in l. The split is made so as to minimize S.

A typical *stopping criterion* is to stop growing the tree when further splits gives less than some minimal amount of extra information δ (the decrease in S becomes less than some threshold δ). Using the notation introduced above, the loss L can be then defined as the mean squared error between the predictions of tree to the data in X, compared to the true responses Y:

$$L = \frac{1}{n} \sum_{i=1}^{n} (y_i - \hat{y})^2.$$

3.4 Features

For each image slice we extract a set of features summarized in Table 1. We consider Haralick and Tamura texture features, homogeneity, gray level statistics, color and morphological features. By concatenating all features together we obtain the final feature descriptor consisting of 1022 elements. We have used in all settings and problems considered the same feature descriptors. Haralick features [9] are based on the gray-level co-occurrence matrix (GLCM) of an image and consist of the following: angular second moment, contrast, autocorrelation, variance, inverse difference moment, sum average, sum variance, sum entropy, entropy, difference variance and difference entropy. Tamura features [26] characterize different texture properties such as total coarseness, coarseness histogram, contrast, directionality. The choice of color features was determined by the fact that most histology images are stained so as to emphasize medical aspects of interest. Typically, depending on the problem, a specific staining procedure is used. Texture features play also an important role as there can be a high variance in the types of tissue considered.

3.5 Label calibration

In Section 3.2 we defined a segment label $\mathbf{y}$ as being composed of real-valued threshold variables corresponding to the intensity, hue and saturation channels. Imposing one threshold for each channel would only allow the detection of two types of regions, i.e., background and foreground. However, using targets such as lower and upper threshold values for all channels allows for a better segmentation. In this case every pixel with a value in one channel smaller than the lower threshold and larger than the upper threshold is removed. The combination of the three channels with both lower and upper threshold values as targets is a flexible and general enough to allow the segmentation of several types of histological tissue.

Thus, each slice is interactively labeled by the user with a target of form $\mathbf{y} = (y_{il}, y_{hl}, y_{sl}, y_{iu}, y_{hu}, y_{su})$, indicating lower and upper threshold bounds for the intensity, hue and saturation. One possible problem could be the introduction of inconsistencies in the regression model. This is due to the fact that the upper and lower thresholds can be widely spread, producing the same image segmentation for multiple values.

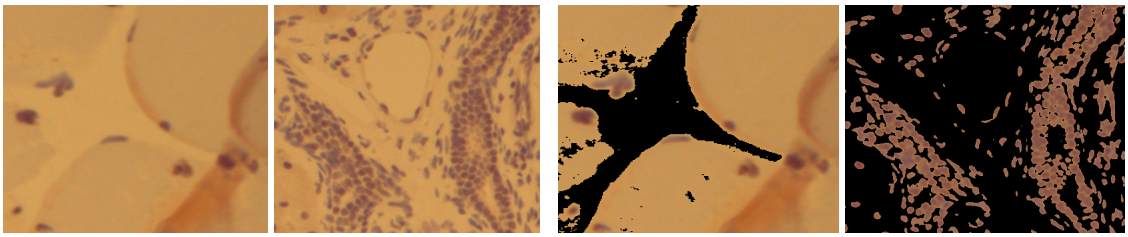

(a) Muscle (right) and liver (left) image slices. (b) Expert thresholded target slices for muscle (right) and liver (left).

Figure 3: Examples of histology image slices and their expert user targets.

Segmenting the same image slice using two different set of target thresholds can produce indistinguishable segmentations from the medical point of view. To address this issue, we introduce a calibration phase, which separates the lower and upper thresholds for each channel as much as possible. This is done by increasing (decreasing) the lower (upper) values until a segmentation using these new values produces a slice that is different enough from the original slice. We consider that an image slice s_1 is *similar* to a slice s_2 if $\frac{\#\, removed\, pixels\, s_2}{\#\, removed\, pixels\, s_1} < \varepsilon$, where ε represents an upper bound on the difference in terms of the number of pixels between s_1 and s_2, such that the two slices can be considered similar enough. Practically, we found that $\varepsilon = 1.05$ is a reasonable value. Note that by increasing (decreasing) the lower (upper) values, it is only possible to obtain more, and not less, black pixels than the original window. The slice with most pixels removed will always appear in the numerator of the fraction.

4 Experiments

To illustrate the performance of our algorithm we report experiments with histology images of two tissue types. These images were obtained by medical experts from the intensive care research group of the KU Leuven, as part of their routine analysis of critical illness and response to therapy at the cellular level. The images pertain to liver and skeletal muscle tissue. In each case, we consider a different problem, that is the segmentation of a different region of interest. For both types of tissues the staining used emphasizes ubiquitinated protein accumulation. All images were taken at a 10-fold magnification factor using a Leica DM3000 microscope. They contain two types of regions of interest: cell nuclei for liver and muscle and fibers for muscle tissue. Examples of image slices for both tissues are illustrated in Figure 3. For liver we consider the problem of nuclei detection, while for muscle the learning has as target nuclei and fiber detection and the exclusion of the overlap-artifacts. Figure 3 shows the two learning problems.

To evaluate the performance of our method we build two datasets, one for liver and one for muscle. For muscle tissue, we used in total 15 full images, out of which 5 for training and 10 for testing. To obtain training and testing examples we randomly sampled 133 training instances and 136 testing instances. For liver tissue, we used 11 full images, out of which 5 for training and 6 for testing. We randomly sampled 99 instances for training and 105 instances for testing. Each image slice was manually annotated by the expert via the GUI's slider, with a set of targets which produced the best medical segmentation. To compare the prediction made for a slice with the ground truth annotation of the same slice, we use *accuracy* as measure. It calculates the ratio of the amount of pixels that match to the total amount of pixels from the two slices.

$$Acc_s = \frac{\#(c_{ij} = 0)}{m \times n} \times 100,$$

where $c_{ij} = a_{ij} - b_{ij}$ with a_{ij} and b_{ij}, elements on the i-th row and j-th column of an image slice s and its ground truth, respectively; $m \times n$ is the size of the slice. The final accuracy is the average over all test slices.

Table 2 shows results for liver and muscle datasets. We report result for the settings when we use calibrated labels and when the calibration was not considered. As a baseline we considered the global image segmentation using thresholding done by the expert. This was done without employing our model. Our algorithm significantly outperforms the global baseline by 43% for the muscle and 12% for liver. We also see that label calibration has a positive influence on the result. To select the parameters of the model (i.e. δ and minimum number of examples to perform the split), we performed 10-fold internal cross-validation on the training set for both types of tissue.

	Method	Calibration	Result
Liver	L	Yes	**93.3**%
		No	91.3%
	G		82.6%
Muscle	L	Yes	**87.4**%
		No	86.1%
	G		42.0%

Table 2: Experimental results on the liver and muscle datasets. L indicates our local method, G is the global baseline. The best result is marked with bold.

5 Conclusions and Future Work

We introduce a new framework which can segment regions of interest in histology images, independently of the problem and tissue type. By combining simple image processing with an interactive machine learning technique, we outperform global thresholding approaches. These are commonly used in medical imaging, yet are more suitable for a certain type of tissue or task. Additionally, our framework is able to learn to automatically segment histology images by interacting with medical experts via a GUI. In this way, medical experts can define multiple and different learning targets, independently of the tissue type.

As future work we plan to investigate whether formalizing the region detection as a structured prediction problem can improve current results. Another line of research is employing online learning approaches, where the model is built incrementally. Also interesting to investigate is whether other machine learning methods (e.g., linear regression, random forests, (multi-class) classification [24], etc) give better results. On the practical level, one useful feature to reduce the training data is to generate informed samples in an active learning setting instead of randomly sampling slices from the image. A relevance feedback or preference elicitation setting is also worth investigating.

References

[1] Features for histology images. `http://www.informed.unal.edu.co/jccaicedo/docs/review.pdf`, November 2009.

[2] N. Bonnet, J. Cutrona, and M. Herbin. A 'no-threshold' histogram-based image segmentation method. *Pattern Recognition*, 35(10):2319–2322, October 2002.

[3] Bernhard E. Boser and et al. A training algorithm for optimal margin classifiers. In *Proceedings of the 5th Annual ACM Workshop on Computational Learning Theory*, pages 144–152. ACM Press, 1992.

[4] Thomas Brox, Yoo-Jin Kim, Joachim Weickert, and Wolfgang Feiden. Fully-automated analysis of muscle fiber images with combined region and edge-based active contours. In *Bildverarbeitung für die Medizin*, pages 86–90, 2006.

[5] A.E. Carpenter, T.R. Jones, M.R. Lamprecht, C. Clarke, I.H. Kang, O. Friman, D.A. Guertin, J.H. Chang, R.A. Lindquist, J. Moffat, et al. Cellprofiler: image analysis software for identifying and quantifying cell phenotypes. *Genome biology*, 7(10):R100, 2006.

[6] L. Chen and P. Pu. Survey of preference elicitation methods. *Swiss Federal Institute of Technology in Lausanne, Technical Report No. IC/200467*, 2004.

[7] Marie Dumont, Raphaël Marée, Louis Wehenkel, and Pierre Geurts. Fast multi-class image annotation with random subwindows and multiple output randomized trees. In *VISAPP*, volume 2, pages 196–203. INSTICC, feb 2009.

[8] M.N. Gurcan, L.E. Boucheron, A. Can, A. Madabhushi, N.M. Rajpoot, and B. Yener. Histopathological image analysis: A review. *IEEE Reviews in Biomedical Engineering*, 2:147 –171, 2009.

[9] Robert M. Haralick, K. Shanmugam, and Its'Hak Dinstein. Textural features for image classification. *IEEE Transactions on Systems, Man, and Cybernetics*, 3(6):610–621, November 1973.

[10] Lei He, L. Rodney Long, Sameer Antani, and George R. Thoma. Computer assisted diagnosis in histopathology. *Sequence and Genome Analysis: Methods and Applications*, pages 271–287, 2011.

[11] Steven C.H. Hoi, Rong Jin, Jianke Zhu, and Michael R. Lyu. Semi-supervised svm batch mode active learning for image retrieval. In *CVPR*, pages 24–26, Alaska, US, June 2008.

[12] A.J. Joshi, F. Porikli, and N. Papanikolopoulos. Multi-class active learning for image classification. In *CVPR*, pages 2372–2379. Ieee, 2009.

[13] L.C. Junqueira and J. Carneiro. *Basic Histology: Text & Atlas*. basic histology. McGraw-Hill, 2005.

[14] J. A. Kiernan. *Histological and Histochemical Methods: Theory and Practice*. Cold Spring Harbor Laboratory Press, 4 edition, 2008.

[15] Inc. Mayachitra. Bring state-of-art image informatics solutions to your desktop, 2006-2001.

[16] S. E. Mills. *Histology for Pathologists*. Lippincott Williams & Wilkins, 3 edition, 2006.

[17] N. Otsu. A threshold selection method from gray-level histograms. *Automatica*, 11:285–296, 1975.

[18] N.R. Pal and S.K. Pal. A review on image segmentation techniques. *Pattern recognition*, 26(9):1277–1294, 1993.

[19] Radu Rogojanu, Giovanna Bises, Cristian Smochina, and Vasile Manta. Segmentation of cell nuclei within complex configurations in images with colon sections. *International Conference on Intelligent Computer Communication and Processing*, 0:243–246, 2010.

[20] Ross, M., Kaye, G. I., and W Pawlina. *Histology: a Text and Atlas*. Lippincott Williams & Wilkins, 4 edition, 2002.

[21] M. Sezgin and B. Sankur. Survey over image thresholding techniques and quantitative performance evaluation. *Journal of Electronic imaging*, 13:146, 2004.

[22] H. Shen, G. Nelson, D.E. Nelson, S. Kennedy, D.G. Spiller, T. Griffiths, N. Paton, S.G. Oliver, M.R.H. White, and D.B. Kell. Automated tracking of gene expression in individual cells and cell compartments. *Journal of the Royal Society Interface*, 3(11):787, 2006.

[23] Arnold W. M. Smeulders, Marcel Worring, Simone Santini, Amarnath Gupta, and Ramesh Jain. Content-based image retrieval at the end of the early years. *TPAMI*, 22(12):1349–1380, December 2000.

[24] C. Sommer, C. Straehle, U. Kothe, and F. A. Hamprecht. Ilastik: Interactive learning and segmentation toolkit. pages 230–233, March 2011.

[25] R. Szeliski. *Computer vision: Algorithms and applications*. Springer-Verlag New York Inc, 2010.

[26] Hideyuki Tamura, Shunji Mori, and Takashi Yamawaki. Texture features corresponding to visual perception. *IEEE Transactions on System, Man and Cybernatic*, 6, 1978.

[27] Alison Todman, Ela Claridge, Alison G. Todman, and Ela Claridge. Cell segmentation in histological images of striated muscle tissue- a perceptual grouping approach, 2008.

[28] Zuva Tranos, Olugbara Oludayo O., Ojo Sunday O., and Ngwira Seleman M. Image segmentation, available techniques, developments and open issues. *Canadian Journal on Image Processing and Computer Vision*, pages 20–29, 2011.

[29] Ching-Wei Wang. A bayesian learning application to automated tumour segmentation for tissue microarray analysis. In *MLMI*, pages 100–107, Berlin, Heidelberg, 2010. Springer-Verlag.

[30] Q. Wang, J. Niemi, C.M. Tan, L. You, and M. West. Image segmentation and dynamic lineage analysis in single-cell fluorescence microscopy. *Cytometry Part A*, 77(1):101–110, 2010.

[31] Zhaozheng Yin, Ryoma Bise, Mei Chen, and Takeo Kanade. Cell segmentation in microscopy imagery using a bag of local bayesian classifiers. In *ISBI*, April 2010.

[32] Xiang Sean Zhou and Thomas S. Huang. Relevance feedback in image retrieval: A comprehensive review. *Multimedia Syst.*, 8(6):536–544, 2003.