Mitosis Extraction in Breast-Cancer Histopathological Whole Slide Images

Vincent Roullier[1,*], Olivier Lézoray[1],
Vinh-Thong Ta[2], and Abderrahim Elmoataz[1]

[1] Université de Caen Basse-Normandie, ENSICAEN, CNRS
GREYC UMR 6072 - Équipe Image
[2] LaBRI (Université de Bordeaux – CNRS) – IPB

Abstract. In this paper, we present a graph-based multi-resolution approach for mitosis extraction in breast cancer histological whole slide images. The proposed segmentation uses a multi-resolution approach which reproduces the slide examination done by a pathologist. Each resolution level is analyzed with a focus of attention resulting from a coarser resolution level analysis. At each resolution level, a spatial refinement by semi-supervised clustering is performed to obtain more accurate segmentation around edges. The proposed segmentation is fully unsupervised by using domain specific knowledge.

1 Introduction

Breast cancer is the second leading cause of cancer death for women. Its incidence increases substantially and continuously while the mortality rate remains high despite earlier detection and advances in therapeutic care. The identification and the use of reliable prognostic and therapeutic markers is a major challenge for decision-making regarding therapy. Proliferation has been shown to be the strongest prognostic and predictive factor in breast carcinoma, especially in patients lacking lymph node metastases [1]. This parameter is daily taken into account by the pathologist for establishing the histopathological grading of breast carcinomas, using enumeration of mitotic figures, through the lens of the microscope. The recent use of immunohistochemical staining of mitosis is able to facilitate their detection. Nevertheless, the visual counting method remains subjective and leads to reproducibility problems due to the frequent heterogeneity of breast tumors [2].

The recently introduced microscopical scanners allow recording large images of the whole histological slides and offer the prospect of fully automated quantification for a better standardization of proliferation rate appraisal. If the advent of such digital whole slide scanners has triggered a revolution in histological imaging, the processing and the analysis of breast cancer high-resolution histopathological images is a very challenging task. First, the produced images are relatively

---

huge and their processing requires computationally efficient tools. Second, the biological variability of the objects of interest makes their extraction difficult. As a consequence, few works in literature have considered the processing of whole slide images and most of these works rely only on machine learning techniques [3,4].

In this work, we present a graph-based multi-resolution segmentation and analysis strategy for histological breast cancer whole slide images. The proposed strategy is based on a top-down approach that mimics the pathologist interpretation under the microscope as a focus of attention. The proposed segmentation performs an unsupervised clustering at each resolution level (driven by domain specific knowledge) and refines the associated segmentation in specific areas as the resolution increases. The whole strategy is based on a graph formalism that enables to perform the segmentation adaptation at each resolution.

The paper is organized as follows. A description of the considered images is presented in Sect. 2. In this Section, we also describe the visual analysis process performed by pathological experts to evaluate mitotic figures proliferation and their inherent multi-resolution approach. Our graph-based formulation for image segmentation is presented in Sect. 3 and its integration into a multi-resolution segmentation strategy is detailed in Sect. 4. Sect. 5 presents visualization tools of extracted mitosis. Last Section concludes.

2   Image Description

2.1   Breast Cancer Histological Whole Slide Images

Breast cancer tissue samples are sectioned at 5 μm thickness and stained with an immunohistochemical (hematoxylin and eosin) method. A ScanScope CS ® (Aperio, San Diego, CA) digital microscopical scanner is then used to digitalize each slice at 20x magnification scale and the resulting digital images are compressed with a quality of 75% following the JPEG compression schema.

To facilitate the visualization and the processing, scanned samples acquired by the scanner are directly stored as an irregular pyramid where each level of the pyramid is an under resolved version of the highest resolution image (the pyramid base).

The usual size of a compressed whole slide image is about 100~500 Megabytes after compression. However, the resulting whole slide images are too large in size to be processed or visualized as a whole. Therefore, the whole slide image is tiled by the scanner to ease both its processing and visualization: each resolution level of the pyramid is split into image tiles in a non-overlapping layout.

2.2   Visual Analysis Process

Within the last decade, histologic grading has become widely accepted as a powerful indicator of prognosis in breast cancer. The majority of tumor grading systems currently employed for breast cancer combine nuclear grade, tubule

formation and mitotic rate. In general, each element is given a score of 1 to 3 (1 being the best and 3 the worst) and the score of all three components are added together to give the breast cancer grading. The usual breast cancer grading scheme is the Elston-Ellis criterion [5] and is based on three separated scores:

– Gland (tubule) formation: one scores the proportion of whole carcinoma that forms acini (1: <75%; 2: 10-75%; 3: <10%).
– Nuclear pleomorphism: one scores the nuclear atypia according to size, shape and chromatin pattern (1: none; 2: moderate; 3: pronounced).
– Mitotic count: one scores the number of mitotic figures per 10 consecutive high power fields (1: 0-9 mitoses; 2: 10-19 mitoses; 3: > 19 mitoses).

The final grading is obtained by adding the three scores. The total score is in the range 3-9 and the final obtained grading is:

– **Grade 1** if total score is 3-5.
– **Grade 2** if total score is 6-7.
– **Grade 3** if total score is 8-9.

In this work, we are interested in helping pathologists to establish an accurate mitotic count. Indeed, with the Elston-Ellis criterion, a pathologist bases its scores only on ten consecutive high power fields. This can be lesser representative than having a score established according to the whole preparation which was of course too tedious for pathologists under a classical microscope until now. With the advent of fast whole slide image scanners, it is now possible [6]. Our study drives towards this direction.

2.3 Multi-resolution Approach

Whole slide images (WSI) are usually huge in size. Fortunately, they are stored as a pyramid of tiled images that enables to process them in a hierarchical way [7]. As a consequence, a multi-resolution segmentation method is a natural approach for segmenting whole slide images. Moreover, such a strategy reproduces the analysis done by the pathologists under the microscope: regions of interest are determined at low resolution while cellular classification is performed at high resolution. The proposed multi-resolution segmentation method is based on a top-down segmentation that mimics pathologist interpretation according to specific domain knowledge. Fig. 1 illustrates the identification (by a pathologist) of mitosis in breast cancer slides stained with hematoxyline and eosine.

3 Graph-Based Segmentation

3.1 Preliminaries on graphs

A graph is a structure used to describe a set of objects and the pairwise relations between those objects. The objects are called *vertices* and a link between two
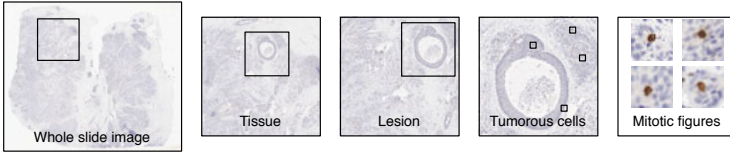
Fig. 1. Illustration of the visual analysis process performed by a pathologist expert. Each decision is performed at a higher resolution than the previous one in the region of interest. Each interior square in an image is magnified in a subsequent image. In the second square, the tissue is detected. In the third square, the pathologist separate tissue and lesion. The third square allows to detect tumourous cells and the last, mitotic figures.

objects is called an *edge*. A weighted graph $G = (V, E, w)$ is composed of a finite set $V = \{u_1, \ldots, u_N\}$ of N vertices, a set of edges $E \subset V \times V$, and a *weight function* $w : E \rightarrow \mathbb{R}^+$. An edge of E, which connects two *adjacent neighbor* vertices u and v, is noted (u, v). In the rest of this paper, the notation $v \sim u$ means that vertex v is an adjacent neighbor of vertex u. We assume that the graph G is simple, connected and undirected. This implies that the weight function w is symmetric i.e. $w(u, v) = w(v, u)$ if $(u, v) \in E$ and $w(u, v) = 0$ otherwise. Let $\mathcal{H}(V)$ be the Hilbert space of real valued functions on the vertices of a graph. Each function $f : V \rightarrow \mathbb{R}$ of $\mathcal{H}(V)$ assigns a real value $f(u)$ to each vertex $u \in V$. Similarly, let $\mathcal{H}(E)$ be the Hilbert space of real valued functions defined on the edges of the graph. For the case of images, nodes are pixels, edges connect neighbor pixels with 8-adjacency.

3.2 Discrete Operators on Graphs

Let us recall some basic definitions. We consider that a graph $G = (V, E, w)$ and a function $f \in \mathcal{H}(V)$ are given. The *weighted difference* $d_w : \mathcal{H}(V) \rightarrow \mathcal{H}(E)$ of a function f on an edge (u, v) linking two vertices $u, v \in V$ is defined as $(d_w f)(u, v) = \sqrt{w(u, v)}\big(f(v) - f(u)\big)$. This operator leads us to define the *directional derivative* of f, over an edge (u, v), as $\partial_v f(u) = (d_w f)(u, v)$. Then, the *weighted gradient* $\nabla_w f$ of the function f, at a vertex $u \in V$, is defined as $(\nabla_w f)(u) = \big(\partial_{v_1} f(u), \ldots, \partial_{v_k} f(u)\big)^T$. This operator corresponds to the *local variation* of the function f at the vertex u and measures the regularity of f in the adjacent neighborhood $v_1, \ldots, v_k$ of the vertex u. Hence, the $\mathcal{L}_2$-norm of the weighted gradient is $\|(\nabla_w f)(u)\|_2 = \left[\sum_{v \sim u} w(u, v)\big(f(v) - f(u)\big)^2\right]^{1/2}$. Then, the *weighted p-Laplacian* $(\Delta_w^p f)(u)$ at vertex u is defined as

$$(\Delta_w^p f)(u) = \sum_{v \sim u} \gamma_p(u, v)\big(f(v) - f(u)\big) \tag{1}$$

where $\gamma_p(u, v) = w(u, v)\big(\|(\nabla_w f)(u)\|_2^{p-2} + \|(\nabla_w f)(v)\|_2^{p-2}\big)$. Clearly, in the case where $p = 1$ and $p = 2$, we have the definitions of the standard graph curvature

$\Delta_w^1 f = \kappa f$ and graph Laplace $\Delta_w^2 f = \Delta f$ operators. More details on these definitions can be found in [8].

In the following, we only consider the case where $p = 2$.

3.3 Discrete Regularization Framework

To regularize a function $f^0 \in \mathcal{H}(V)$ using the p-Laplacian (Eq. (1)), we consider the following general variational problem on graphs:

$$\min_{f \in \mathcal{H}(V)} \left\{ E_w(f, f^0, \lambda, p) = R_w(f, p) + \frac{\lambda}{2}\|f - f^0\|_2^2 \right\}. \tag{2}$$

The first term, $R_w(f, p)$, is the regularizer and is defined as, with $0 < p < +\infty$: $R_w(f, p) = \frac{1}{p} \sum_{u \in V} \|(\nabla_w f)(u)\|_2^p$. The second term is the fitting term. This optimization problem has a unique solution for $p = 1$ and $p = 2$ which satisfies, for all $u \in V$:

$$\frac{\partial E_w(f, f^0, \lambda, p)}{\partial f(u)} = (\Delta_w^p f)(u) + \lambda\big(f(u) - f^0(u)\big) = 0,$$

which is equivalent to

$$\left(\lambda + \sum_{v \sim u} \gamma(u, v)\right) f(u) - \sum_{v \sim u} \gamma(u, v) f(v) = \lambda f^0(u).$$

To approximate the solution of the minimization (2), we can linearize this system of equations and use the Gauss-Jacobi method to obtain the following iterative algorithm:

$$\begin{cases} f^{(0)}(u) = f^0(u) \\ f^{(n+1)}(u) = f^0(u) + \sum_{v \sim u} \lambda + \sum_{v \sim u} \gamma^{(n)}(u, v), \end{cases} \tag{3}$$

where $\gamma^{(n)}(u, v)$ is the γ function (in Eq. (1)) at the iteration step n. The interested reader can refer to [8] for more details on the formulation and the connections with other formalisms. The above algorithm enables to simplify functions living on graphs by a discrete diffusion process.

3.4 Discrete Semi-supervised Clustering

The previously presented discrete regularization framework can be naturally adapted to address discrete semi-supervised clustering problems. Let $V = \{u_1, \ldots, u_N\}$ be a finite set of data, where each data u_i is a vector of $\mathbb{R}^m$. Let $G = (V, E, w)$ be a weighted graph such that all vertices are connected by an edge of E. The semi-supervised clustering of the set V consists in grouping the set V into k classes where the number of k classes is given. For this, the set V is composed of labeled and unlabeled data. The objective is then to estimate

the labels of unlabeled data from labeled ones. Let $C = \{c_i\}_{i=1,\ldots,k}$ the set of classes, L the set of labeled vertices and $V \setminus L$ be te initially unlabelled vertices (the whole set of vertices except the labeled ones). For each vertex of L, its classes is available with the function $\mathcal{L} : L \to C$.

This situation can be modeled by considering k initial label functions (one per class) $f_i^0 : V \to \mathbb{R}$, with $i = 1, \ldots, k$. For a given vertex u, if u is initially labeled ($u \in L$) then $f_i^0(u) = +1$ if $\mathcal{L}(u) \in c_i$ and $f_i^0(u) = -1$ otherwise. If u is initially unlabeled (i.e. $u \in V \setminus L$) then $f_i^0(u) = 0$. Then, the vertex clustering is accomplished by k regularization processes. This corresponds to estimate functions $f_i : V \to \mathbb{R}$ for each i^{th} class using the discrete diffusion process (Eq. (3)). At the end of the label propagation processes, the final label of a given vertex $u \in V$ can be obtained by $\underset{i}{\operatorname{argmax}}\left\{ f_i(u) / \sum_{j=1,\ldots,k} f_j(u) \right\}$.

4   Multi-resolution Segmentation Approach

4.1   Principle

As it has been previously pointed out, a multi-resolution segmentation process is a natural approach to analyze whole slide images [7,9]. Indeed, we have seen in Sect. 2.2 that the whole slide image analysis visual process performed by pathologist experts is a multi-resolution process. An expert determines regions of interest at low resolution while cellular classification is performed at high resolution.

Our proposed multi-resolution segmentation process is based on a top-down segmentation that reproduces exactly the interpretation process performed by pathologist experts according to specific domain knowledge (expressed by the pathologists themselves). At a given resolution i, a clustering is performed by the following steps: the image is simplified by discrete regularization (Fig 2 (a)-(b)) and clustered by an unsupervised 2-means clustering (Fig. 2 (c)-(d)). The clustering is performed inside specific region that were segmented at the previous resolution. The obtained clustering is spread by pixel replication at a finer level of resolution (Fig. 2 (e)-(f)) and refined in specific region (according to domain knowledge) (Fig. 2 (g)-(i)). As clustering being performed in a feature space, it does not take into account spatial information, and the obtained segmentation is not accurate around image edges. In addition, the propagation of the labels across the different resolution levels is performed by plain pixel replication and the segmentation is coarse around edges. To alleviate both these effects, each obtained clustering is refined by our discrete semi-supervised clustering in a narrow band around the boundaries of the clusters. The whole segmentation strategy can be summarized by Algorithm 1 where I_i denotes an image at resolution level i.

At the last resolution level, mitotic figures are extracted (Fig. 2 (j)-(l)). Fig. 3 provides results for several whole slide images.
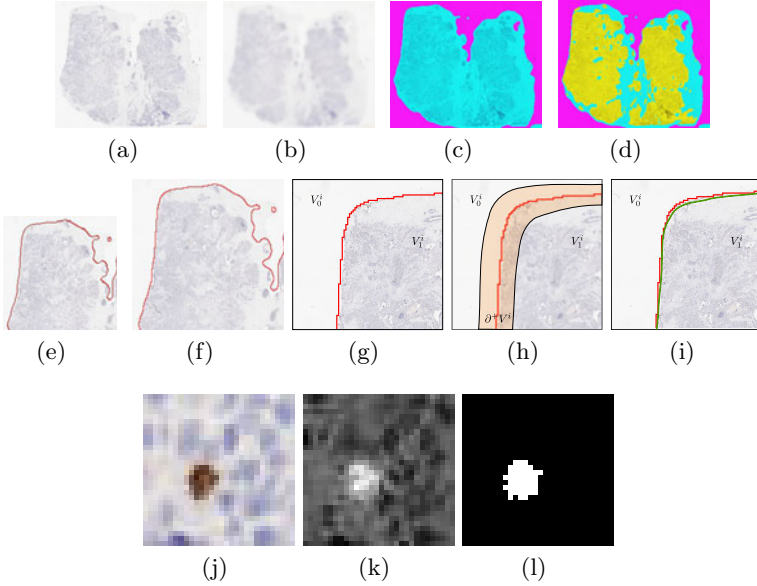
Fig. 2. Illustration of the multi-resolution segmentation process. (a) initial image, (b) regularized image. (c) and (d) clustered image at two level 1 and 2. (e) and (f) replication problem. (g)-(i) illustration of the spatial refinement clustering. (j)-(l) the mitotic figures extraction.

Algorithm 1. Multi-resolution WSI segmentation

1: $I_0^s = \text{Regularization}(I_0)$
2: $I_0^c = \text{2-MeansClustering}(I_0^s)$
3: $I_0^r = \text{SpatialClusteringRefinement}(I_0^c)$
4: **for** $i = 1$ to 3 **do**
5:    $I_i^p = \text{ReplicatePreviousResolutionClustering}(I_{i-1}^r)$
6:    $I_i^s = \text{Regularization}(I_i)$
7:    $I_i^c = \text{2-MeansClustering}(I_i^s)$ inside one class of I_i^p
8:    $I_i^r = \text{SpatialClusteringRefinement}(I_i^c)$
9: **end for**

---

5   Visualization of Mitotic Figures

Once all the mitosis have been extracted at the highest resolution using our top-down multi-resolution graph-based extraction algorithm, the pathologist has to visualize them to establish the mitotic score. Our proposal does not intend to compute the mitotic score but at helping the pathologist to do it. To do so, we provide for pathologists two visualization tools.
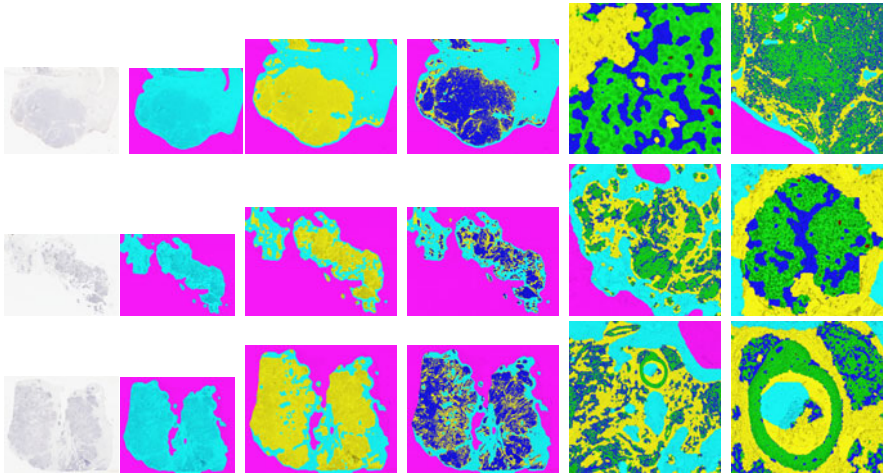
Fig. 3. Illustration of the multi-resolution clustering process. The first column presents original images. The second column is the segmentation obtained at the first resolution (background in pink, tissue in cyan), the third column is the segmentation obtained at the second resolution (lesion in yellow), the fourth column is the segmentation obtained at the third resolution (stroma in yellow and tumorous cells dark blue). The fourth column is the segmentation obtained at the fifth resolution (stroma in dark blue and tumorous cells in green). The sixth column is the segmentation obtained at the fifth resolution (mitotic figures in red).

The first one enables the pathologist to evaluate the global repartition of mitosis on the whole slide. This visual information is provided by superimposing a graph on the whole slide image. The graph is constructed as follows. Each detected mitosis is represented by a vertex of the graph. A Voronoi map is computed on the vertices coordinates and the associated Delaunay graph is obtained. This enables the pathologist to evaluate regions on the whole slide image where the mitotic activity is important by the superposition of either the Delaunay graph or the Voronoi distance map on the whole slide image (Fig. 4(a)-(c)).

The second visualization tool enables the pathologist to see the extracted mitosis altogether on a single 3D projection. Indeed, the first projection tool enables the pathologist to visualize each extracted mitosis on the whole image while appreciating the whole global distribution of mitosis. However, this does not allow to simultaneously visualize all the extracted mitosis of the slide to appreciate their aspect (e.g. to differentiate the different mitotic phases). Therefore, we propose a specific visualization tool that provides this information to the pathologist.

The proposed visualization tool is based on dimensionality reduction with Laplacian Eigenmaps [10]. Dimensionality reduction requires a distance measure to evaluate the similarity between two objects in the initial space. In the case of mitosis, since they can be in different mitotic phases, the most prominent information is texture. Therefore, a texture description of each mitotic figure is

computed in the form of a Locally Binary Pattern (LBP) histogram (introduced in [11]) and this feature vector is used as an input for dimensionality reduction with a χ^2 histogram distance.

Once the dimensionality reduction has been performed, the pathologist can visualize simultaneously all the mitotic figures of a whole slide image in the form of a 3D projection where each mitosis is projected at coordinates defined by the projection. With this projection, the pathologist can appreciate the similarity of mitosis that are not necessarily spatially close in the whole slide image. Finally, performing a dimensionality reduction with Laplacian Eigenmaps provides also geometrical information on the projection. Indeed, the sign of the first eigenvector enables to partition the data into two sets. The obtained partitions correspond to the normalized cut criterion of the initial data [12]. As a consequence, this clustering information is also provided by coloring the bounding box of each mitosis in a color corresponding to one of the two clusters (red or green). This enables the pathologist to quickly distinguish mitosis with low or high textural content (Fig. 4(d)-(f)).
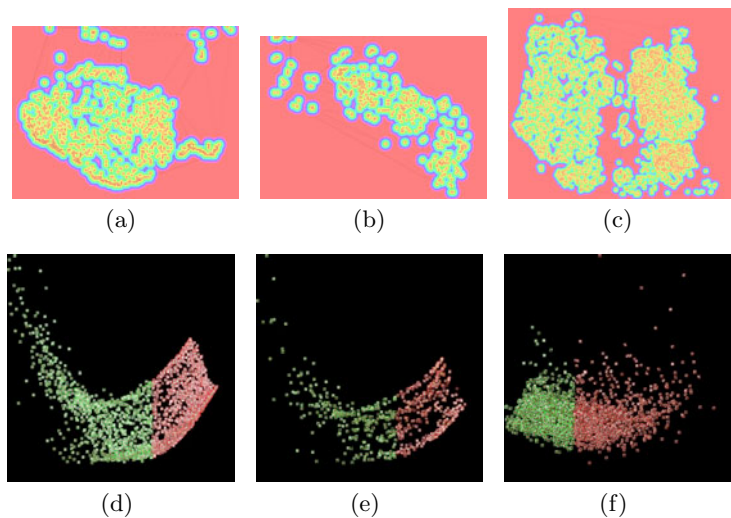

Fig. 4. Illustration of visualization tools of extracted mitotic figures. (a)-(b)-(c): Delaunay graph superimposed on the distance maps associated. (d)-(e)-(f): visualization of extracted mitotic figures by dimensionality reduction for the mitosis of (a)-(c) (see electronic version of the paper for better visualization).

6   Conclusion

In this paper, a multi-resolution image analysis strategy for automatic enumeration of mitotic figures on whole slide images is proposed. The whole classification process begins with the lowest resolution image and moves to higher resolution into regions of interest gradually identified. Graph-based regularization provides

a unified formalism for both image simplification and spatial cluster refinement. Contrary to methods that can be found in literature, our method is completely unsupervised and has the advantage of reducing the amount of data to be processed at each resolution level by selecting regions of interest.

We also propose two methods for the visualization of mitotic figures. The first method allows to visualize the distribution of mitosis on the tissue samples. The second method groups mitosis according to texture parameters.

Future works will concern the automation of the other scores of the Elston-Ellis grading systems.

References

1. Tavassoli, F., Devilee, P. (eds.): Pathology and Genetics, Tumours of the Breast and Female Genital Organ. International Agency for Research on Cancer Press (2003)
2. Petushi, S., Garcia, F., Haber, M., Katsinis, C., Tozeren, A.: Large-scale computations on histology images reveal grade-differentiating parameters for breast cancer. BMC Medical Imaging 6, 1–11 (2006)
3. Ruiz, A., Kong, J., Ujaldon, M., Boyer, K., Saltz, J., Gurcan, M.: Pathological image segmentation for neuroblasma using the gpu. In: ISBI, pp. 296–299 (2008)
4. Signolle, N., Plancoulaine, B., Herlin, P., Revenu, M.: Texture-based multiscale segmentation: Application to stromal compartment characterization on ovarian carcinoma virtual slides. In: Elmoataz, A., Lezoray, O., Nouboud, F., Mammass, D. (eds.) ICISP 2008. LNCS, vol. 5099, pp. 173–182. Springer, Heidelberg (2008)
5. Elston, C.W., Ellis, I.O.: Pathological prognostic factors in breast cancer. I. The value of histological grade in breast cancer: experience from a large study with long-term follow-up. Histopathology 19, 403–410 (1991)
6. Dalle, J., Leow, W., Racoceanu, D., Tutac, A., Putti, T.: Automatic breast cancer grading of histopathological images. In: EMBS, pp. 3052–3055 (2008)
7. Kong, J., Sertel, O., Shimada, H., Boyer, K., Saltz, J., Gurcan, M.: Computer-aided evaluation of neuroblastoma on whole-slide histology images: Classifying grade of neuroblastic differentiation. Pattern Recognition 42, 1080–1092 (2009)
8. Elmoataz, A., Lézoray, O., Bougleux, S.: Nonlocal discrete regularization an weighted graphs: a framework for image and manifolds processing. IEEE Transactions on Image Processing 17, 1047–1060 (2008)
9. Roullier, V., Ta, V.T., Lézoray, O., Elmoataz, A.: Graph-based multi-resolution segmentation of histological whole slide images. In: ISBI, pp. 153–156 (2010)
10. Belkin, M., Niyogi, P.: Laplacian eigenmaps for dimensionality reduction and data representation. Neural Computation 15, 1373–1396 (2003)
11. Ojala, T., Pietikäinen, M., Harwood, D.: A comparative study of texture measures with classification based on feature distributions. Pattern Recognition 29, 51–59 (1996)
12. Shi, J., Malik, J.: Normalized cuts and image segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence 22, 888–905 (1997)