Oriented Shape Index Histograms for Cell Classification

Anders Boesen Lindbo Larsen(✉), Anders Bjorholm Dahl, and Rasmus Larsen

Department of Applied Mathematics and Computer Science,
Technical University of Denmark, Kongens Lyngby, Denmark
{abll,abda,rlar}@dtu.dk

Abstract. We propose a novel extension to the shape index histogram feature descriptor where the orientation of the second-order curvature is included in the histograms. The orientation of the shape index is reminiscent but not equal to gradient orientation which is widely used for feature description. We evaluate our new feature descriptor using a public dataset consisting of HEp-2 cell images from indirect immunoflourescence lighting. Our results show that we can improve classification performance significantly when including the shape index orientation. Notably, we show that shape index orientation outperforms the gradient orientation on the dataset.

1 Introduction

When characterising texture-like structures in images, it is often desirable to be invariant towards orientations of the image structure at larger scales [2,7,11,12]. However, in cases where one can assume or estimate an orientation in the image, it becomes relevant to describe the texture relative to this orientation to better characterise the structure.

The problem we wish to solve is classification of HEp-2 cell images captured using indirect immunoflourescence lighting. This is a quintessential task in medical image analysis where good performance of an automated system can save manual labour hours and speed up the diagnosis process [4,5]. A common approach is to regard the problem as an instance of texture classification because the staining patterns of the cells exhibit texture characteristics. Furthermore, for cell images, a local coordinate system can be contrived pixel-wise by letting the origin be the cell centre, the first axis be the vector from the origin to the pixel, and the second axis be orthogonal to the first axis in either the clockwise or counter-clockwise direction. This yields a fidicual orientation for every pixel that we can use for describing local texture orientation.

In this paper we explore the possibility of capturing texture orientation in combination with *shape index histograms* (SIHs) [7,13]. The shape index is a second-order curvature measure derived from the eigenvalues of the local image Hessian [6]. With the shape index, rotation invariance is achieved by construction because the eigenvectors (the *principal directions*) are discarded. We extend the work on shape index histograms by reincorporating the orientation in the feature description. We call our extension *oriented shape index histograms* (OSIHs).

R.R. Paulsen and K.S. Pedersen (Eds.): SCIA 2015, LNCS 9127, pp. 16–25, 2015.
DOI: 10.1007/978-3-319-19665-7_2

1.1 Related Work

Feature description using SIHs has recently shown good performance for medical image analysis [7] and for analysing galaxy images [13]. SIHs have shown superior performance compared with popular texture descriptors and came ind 2nd 0.11% below the 1st place at the ICIP 2013 competition on cell classification[1] [7].

The construction of SIHs is comparable to most histogram-based feature descriptors that typically rely on first-order image structure, *e.g.* SIFT [10], DAISY [15] and HOG [3]. The histogram approach to feature description has had enormous success because it offers good discriminability and robustness. For SIFT, DAISY and HOG, the robustness is towards local translations and rotation. For SIHs as used in [7,13], the histogram description achieves robustness from the smoothing of adjacent second-order curvatures captured in the histogram bins. For histograms over shape index orientation, the histogram approach allow for robustness towards rotation; again from the bin smoothing.

When constructing histograms over multiple features with potential correlations we have a choice between 1) creating a joint histogram, or 2) concatenating the marginal histograms. The authors of the basic image features descriptor [2] opt for the latter and show good results using a multi-scale joint histogram for texture description. While the joint histogram should allow for a more discriminative description, it quickly becomes high-dimensional which may be impractical. This is a likely explanation for why joint histograms are rarely seen in the feature description literature. In this paper, we wish to experiment with both approaches.

We note that second-order orientation of the shape index should not be confused with the first-order orientation of the gradient. The gradient orientation is the direction with the steepest increase in image intensity whereas the shape index orientation is the angle along which the second-order curvature is strongest. To the best of our knowledge, the orientation of the second-order curvature has not been used for feature description before.

1.2 Contributions

- We formulate the shape index orientation and show how eigenvectors of the local image Hessian capture second-order orientation. We also formulate a strength measure of the second-order curvature to be used for weighting histogram contributions.
- We develop OSIHs as an extension of SIHs by including the shape index orientation.
- We evaluate our new feature descriptor and show that it improves classification performance on the ICIP 2013 dataset compared to plain SIHs and gradient orientation histograms.

[1] Dataset available at http://nerone.diiie.unisa.it/contest-icip-2013

2 Oriented Shape Index Histograms

In this section we formulate our new feature descriptor. The shape index is an image geometry measure originally proposed by Koenderink and van Doorn [6]. It is based on the scale-space framework as presented in [9]. Note that we have made our implementation of OSIHs available online[2] .

2.1 Differential Image Structure

We formulate our oriented shape index histograms using the Gaussian scale-space framework where the differential structure of a 2-dimensional image $I(\boldsymbol{x})$: $\Omega \rightarrow R,\ \Omega \subseteq R^2$ is defined by

$$L_{x^n y^m}(\boldsymbol{x};\sigma) = \sigma^{n+m} \frac{\partial^{n+m}}{\partial x^n \partial y^m} \left(G * I\right)(\boldsymbol{x};\sigma) \quad . \tag{1}$$

$*$ denotes convolution, G is a Gaussian kernel, σ is the width of the Gaussian kernel, n and m indicate the differentiation order along the x and y axis respectively. For notational convenience we omit the arguments and substitute $L_{x^n y^m}(\boldsymbol{x};\sigma)$ with simply $L_{x^n y^m}$. Thus, $L_{x^n y^m}$ is implicitly assumed to be computed at some scale σ and location $\boldsymbol{x}$.

2.2 The Shape Index

The shape index is derived from the Hessian matrix $\nabla^2 L$ that captures the second-order curvature at scale σ,

$$\nabla^2 L(\boldsymbol{x};\sigma) = \begin{bmatrix} L_{x^2} & L_{xy} \\ L_{xy} & L_{y^2} \end{bmatrix} \quad . \tag{2}$$

The Hessian matrix is square and symmetric allowing us to compute the pair of real eigenvalues κ_1 and κ_2 capturing the *principal curvatures*. The shape index $s \in \left]\frac{-\pi}{2}, \frac{\pi}{2}\right[$ is defined as

$$s(\boldsymbol{x};\sigma) = \arctan\left(\frac{\kappa_1 + \kappa_2}{\kappa_1 - \kappa_2}\right) \quad . \tag{3}$$

The shape index has the attractive property that it maps all second-order shapes onto a continuous interval providing a smooth and intuitive transition between the shapes, see Figure 1. In addition to the shape index a measure of *curvedness* c is defined,

$$c(\boldsymbol{x};\sigma) = \sqrt{\kappa_1^2 + \kappa_2^2} \quad . \tag{4}$$

The curvedness indicates the strength of the shape described by the shape index such that we differentiate between flat and indistinct vs. large and prominent shapes. See Figure 1 for examples.

[2] Implementation available at http://compute.dtu.dk/~abll.

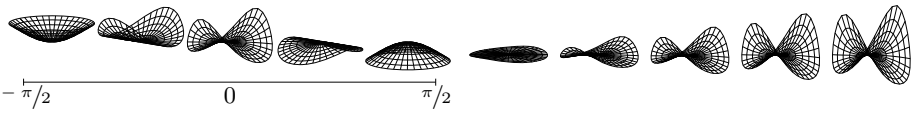

Fig. 1. Left: Second-order curvatures along the shape index interval $]-\pi/2, \pi/2[$. Right: The saddle shape $s = 0$ with increasing curvedness.

2.3 Shape Index Orientation

The orientation of the shape index is described by the eigenvectors of $\nabla^2 L$. The eigenvectors are the *principal directions* along which the curvature is minimal and maximal respectively. Because $\nabla^2 L$ is symmetric, the eigenvectors are orthogonal. Therefore, we can rely on a single eigenvector $\boldsymbol{v} = (v_1, v_2)$ to capture the shape index orientation $\theta \in]-\pi/2, \pi/2[$:

$$\theta(\boldsymbol{x}; \sigma) = \arctan\left(\frac{v_2}{v_1}\right) \tag{5}$$

The angle is unsigned (0–180°) because the shape index curvature is symmetric around $\boldsymbol{v}$. As an example, we show the shape index orientation θ for a test image in Figure 2. We see that the saddle points in the checkerboard patterns have orthogonal directions. For the rings we see that θ changes smoothly and is periodic in $]-\pi/2, \pi/2[$. Note that θ is unstable at the centre of the small dots. This is because the shape index curvature has no dominant orientation for a blob structure where $\kappa_1 = \kappa_2$.

Along with the shape index orientation we introduce a measure of its strength,

$$m = \kappa_1 - \kappa_2 \quad . \tag{6}$$

The shape index orientation magnitude m is computed from the difference between the eigenvalues. If m is small it means that the principal curvatures along the principal directions are similar and the second-order shape is radially symmetric with no orientation. In Figure 2, we show an example of the second-order orientation magnitude. We see that its response is different from the gradient magnitude (introduced later in Section 2.5) as it gives weight only where second-order structure has orientation.

2.4 Histogram Construction

The ordinary SIH is constructed by choosing a set of n_s bin centres $b_{s1}, \ldots, b_{sn_s}$ equidistantly distributed along the shape index interval $]-\pi/2, \pi/2[$. The shape index bin contribution C_s at location $\boldsymbol{x}$ is computed from

$$C_s(\boldsymbol{x}; \sigma, b_s, \beta_s) = \exp\left(-\frac{(b_s - s)^2}{2\beta_s^2}\right) \quad , \tag{7}$$

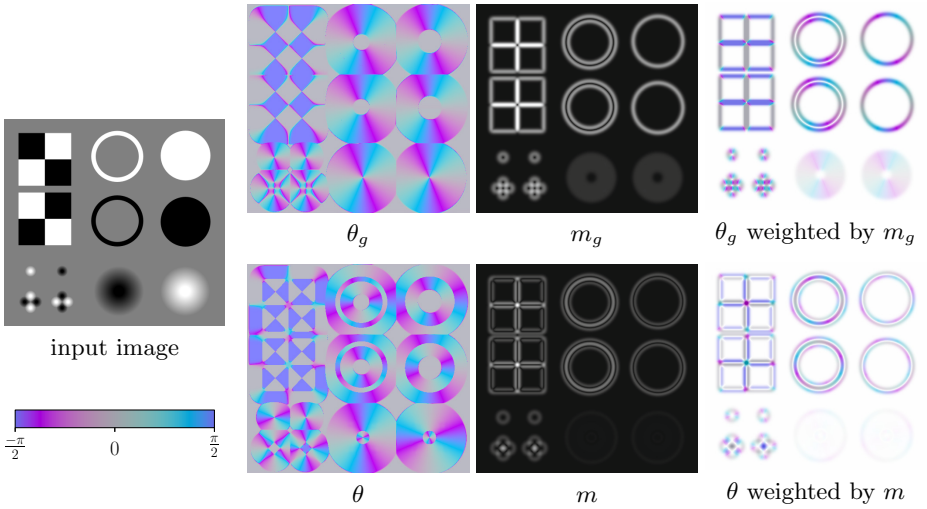

Fig. 2. A toy example to illustrate the shape index orientation and its difference with the gradient orientation. Top row: Gradient orientations (unsigned). Bottom row: Shape index orientations. The abrupt transitions to grey in the left column is caused by limited floating point precision. In the right column we perform alpha-blending using the magnitudes emphasise the orientations with high magnitudes.

where a Gaussian window of width β_s is used for smooth binning in the *tonal* range (the shape index range). The total bin contributions from all image locations is computed from

$$B_s(\sigma, b_s, \beta_s) = \sum_{\boldsymbol{x}} C_s c \quad . \tag{8}$$

Note that we use the curvedness as weight for each bin contribution in order to increase the influence of prominent shapes in the image. The shape index histogram is then constructed as the vector of total bin contributions for all bin centres:

$$H_s(\sigma, \beta_s) = [B_s(\sigma, b_{s1}, \beta_s), \ldots, B_s(\sigma, b_{sn_s}, \beta_s)] \tag{9}$$

In the same way as above we can construct a histogram over shape index orientations by choosing a set of n_θ bin centres $b_{\theta 1}, \ldots, b_{\theta n_\theta}$ distributed along the orientation interval $]-\pi/2, \pi/2[$ and calculate the bin contribution, the total bin contributions and the histogram from:

$$C_\theta(\boldsymbol{x}; \sigma, b_\theta, \beta_\theta) = \exp\left(\frac{1}{\beta_\theta}\cos(b_\theta - o)\right) \tag{10}$$

$$B_\theta(\sigma, b_s, \beta_\theta) = \sum_{\boldsymbol{x}} C_\theta m \tag{11}$$

$$H_\theta(\sigma, \beta_\theta) = [B_\theta(\sigma, b_{\theta 1}, \beta_\theta), \ldots, B_\theta(\sigma, b_{\theta n_\theta}, \beta_\theta)] \tag{12}$$

Again, we choose a β_θ as smoothing parameter for the tonal range. Note that we have changed the smoothing from a normal distribution to the circular normal distribution to accommodate for the periodic behaviour of θ.

To construct a joint histogram over the shape index and its orientation we calculate the contributions to a bin centred at shape index b_s and orientation b_θ from

$$B_{s\theta}(\sigma, b_s, b_\theta, \beta_s, \beta_\theta) = \sum_{\boldsymbol{x}} C_s c C_\theta m \tag{13}$$

We can then construct the joint oriented shape index histogram $H_{s\theta}$ from the combination of all $B_{s\theta}$:

$$H_{s\theta}(\sigma, \beta_s, \beta_\theta) = \begin{bmatrix} B_{s\theta}(\sigma, b_{s1}, b_{\theta 1}, \beta_s, \beta_\theta) & \dots & B_{s\theta}(\sigma, b_{sn_s}, b_{\theta 1}, \beta_s, \beta_\theta) \\ \vdots & \ddots & \vdots \\ B_{s\theta}(\sigma, b_{s1}, b_{\theta n_\theta}, \beta_s, \beta_\theta) & \dots & B_{s\theta}(\sigma, b_{sn_s}, b_{\theta n_\theta}, \beta_s, \beta_\theta) \end{bmatrix} \tag{14}$$

At this point, we can perform multi-scale feature description by concatenating histograms at different scales. Following [2], we propose selecting n_σ different scales such that

$$\sigma_i = \sigma_{\text{base}} \cdot \sigma_{\text{ratio}}^i \quad , \quad i = 1, \dots, n_\sigma \quad . \tag{15}$$

σ_{base} is the smallest shape index scale and σ_{ratio} is the ratio between σ_i and σ_{i+1}. As a final step we normalise the histogram vectors in order to make the description robust to image contrast variations. We have experimented with different normalisation schemes (L_1, L_2 and RootSIFT [1]) without observing significant differences in discriminative performance. We choose L_1 normalisation.

2.5 Gradient Orientation Histograms

For completeness, we also list the gradient orientation histogram which will be used in our experiments. The gradient orientation θ_g and its magnitude m_g are computed from

$$\theta_g(\boldsymbol{x}; \sigma) = \operatorname{arctan2}(L_y, L_x) \quad , \quad m_g = \sqrt{L_x^2 + L_y^2} \tag{16}$$

Similar to the shape index orientation we can construct a histogram by choosing a set of n_g bin centres $b_{g1}, \dots, b_{gn_g}$ distributed along the gradient orientation interval $]-\pi, \pi[$ and the histogram from:

$$C_g(\boldsymbol{x}; \sigma, b_g, \beta_g) = \exp\left(\frac{1}{\beta_g} \cos\left(\frac{b_g - \theta_g}{2}\right)\right) \tag{17}$$

$$B_g(\sigma, b_s, \beta_g) = \sum_{\boldsymbol{x}} C_g m_g \tag{18}$$

$$H_g(\sigma, \beta_g) = [B_g(\sigma, b_{g1}, \beta_g), \dots, B_g(\sigma, b_{gn_g}, \beta_g)] \tag{19}$$

3 Experiments

We evaluate the oriented shape index histograms on the dataset from the *Competition on Cell Classification by Fluorescent Image Analysis* at ICIP 2013 [4]. The dataset consists of 13.596 indirect immunoflourescence images of HEp-2 cells. The cell images come from 83 patients and the task is to classify the cells into 6 classes according to their fluorescence staining pattern. See Fig. 3 for examples.

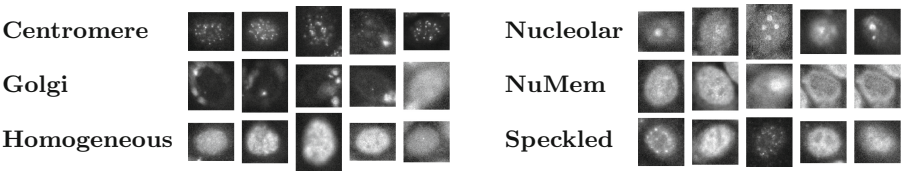

Fig. 3. Examples of the 6 different staining pattern classes in the ICIP 2013 dataset

We assess classification performance from a leave-one-patient-out cross-validation study across all 83 patients in the dataset. Because the patients have different numbers of cells, we measure performance as the weighted average over classification accuracies for the 83 patients where the weights are the number of cells per patient. We employ a standard pipeline consisting of feature extraction and classification. As classifier, we use an RBF kernel SVM with a fixed $C = 1$ and $\gamma_{\text{RBF}} = 1/N$ where N is the dimensionality of the feature vector. Multi-class support is achieved with a one vs. one comparison scheme. We argue that this fixed classifier configuration is suitable for our comparison since the focus of this paper is on feature description and because our (non-shown) experiments indicate that only very small performance improvements can be achieved from adjusting these parameters.

We consider a variety of different feature descriptors in our experiments. In an attempt avoid bias in the comparison, we optimise the parameters for each texture measure using Bayesian optimisation with the framework provided in [14]. For each texture measure, we let the framework perform around 150 function evaluations before selecting the optimal configuration. To avert selecting an accidentally good parameter setting among the 150 parameter configurations that overfits to the dataset, we perform the Bayesian optimisation using cross-validation on 40 randomly selected patients.

In the following, we list the feature descriptor variants that we use for our experiments and their optimised parameters. For all experiments we have set $n_s = 18,\ n_\theta = 16,\ n_{\theta_g} = 16,\ n_\sigma = 5$ when applicable. When calculating the shape index orientation, we use the vector from the cell centre (the centre of the image) to the location of the shape index orientation as fiducial orientation.

- SIH – The shape index histogram H_s with dim. $n_\sigma n_s = 90$ and $\sigma_{\text{base}} = 1.8,\ \sigma_{\text{ratio}} = 1.5,\ \beta_s = 0.33$.

- SIH+GOH – The concatenation of shape index histograms H_s and gradient orientation histograms (GOHs) H_{θ_g} with dim. $n_\sigma(n_s + n_{\theta_g}) = 170$ and $\sigma_{\text{base}} = 1.0,\ \sigma_{\text{ratio}} = 1.8,\ \beta_s = 0.16,\ \beta_{\theta_g} = 0.21$.
- OSIH – The concatenation of H_s and H_θ with dim. $n_\sigma(n_s + n_\theta) = 170$ and $\sigma_{\text{base}} = 1.2,\ \sigma_{\text{ratio}} = 1.6,\ \beta_s = 0.19,\ \beta_\theta = 0.26$,.
- JOSIH – The joint oriented shape index histogram $H_{s\theta}$ with dim. $n_\sigma n_s n_\theta = 1440$ and $\sigma_{\text{base}} = 1.2,\ \sigma_{\text{ratio}} = 1.7,\ \beta_s = 0.19,\ \beta_\theta = 0.25$.
- GOH+OSIH – The concatenation of H_s, H_θ and H_{θ_g} with dim. $n_\sigma(n_s + n_\theta + n_{\theta_g}) = 250$ and $\sigma_{\text{base}} = 1.1,\ \sigma_{\text{ratio}} = 1.8,\ \beta_s = 0.17,\ \beta_s = 0.22,\ \beta_\theta = 0.27,\ \beta_{\theta_g} = 0.21$.

The performance of the descriptors are shown in Figure 4. The significance of the differences in performance is difficult to asses from the boxplots alone and therefore we provide significance levels[3] when comparing descriptor by their performances in the following.

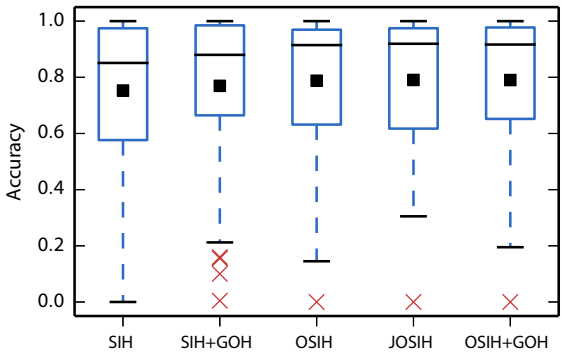

Fig. 4. Comparison of feature descriptors. The boxplot statistics are gathered from the cross-validation across 83 patients. The boxplot is generated using weighted percentiles with the weights being equal to the number of cells per patient. The black square indicates the weighted average accuracy. For the 5 feature descriptors the weighted average accuracies are 0.7530, 0.7696, 0.7878 0.7906, 0.7902.

We see that plain SIH achieves a weighted average accuracy of 0.7530. By adding gradient orientation (SIH+GOH) the accuracy increases to 0.7696 which is significant at a 0.018 confidence level. By adding shape index orientations

[3] Our hypothesis test for determining if descriptor A has greater performance than descriptor B is based on bootstrapping. We generate 10,000 bootstrap samples. Each sample is obtained by drawing 83 patients with replacement and computing $\mu_{\text{diff}} = \mu_{\text{A}} - \mu_{\text{B}}$ where μ_{A} is descriptor A's weighted mean performance on the drawn patients. That is, our bootstrap estimator is the difference between the two weighted mean performances. We wish to test $H_0 : \mu_{\text{A}} \leq \mu_{\text{B}}$ versus $H_1 : \mu_{\text{A}} > \mu_{\text{B}}$. We report the significance level as the bootstrap estimated p-value.

(OSIH) the accuracy increases to 0.7878 which is significant at a 0.004 confidence level. However, going from OSIH to either JOSIH or OSIH+GOH offer only very small improvements significant at 0.318 and 0.359 confidence levels respectively. When comparing SIH+GOH to OSIH the discriminative ability of the shape index orientation over the gradient orientation is significant on a 0.0651 confidence level.

4 Discussion

As we have seen from our experiments, the orientation of second-order curvature improves the discriminative ability of SIHs for the task of cell classification. We find it noteworthy the performance improvements of OSIH are significant from the improvements of GOH. This leads us to conclude that the orientation of the shape index and the gradient capture different image structures. We have also investigated if joint histograms yield better feature description. However, we have shown that this is barely the case on our dataset - especially considering the cost of a much higher feature vector dimensionality. We speculate that this result might be caused by the dataset that contains small images (approximately 80×80 pixels). If the images are too small, we might not be able to collect sufficient statistics for the large joint histograms to become representative. This would favour marginal histograms.

A limitation of OSIHs is that we require an orientation to be given in the image. This cannot be assumed in all situations; *e.g.* for general texture classification we are only given an image. We imagine that one in those cases can detect a dominant orientation and extract OSIHs relative to this orientation. This approach is reminiscent of SIFT where dominant orientations are estimated from peaks in a local gradient orientation histogram. We consider this direction future work.

The purpose of this paper has been to investigate the effectiveness of second-order orientation for texture description. We could have applied other common techniques to improve classification performance (e.g. spatial decomposition [7,8] or bag-of-visual-words [16]). However, this would have muddled our message that second-order orientation is indeed useful and competitive with first-order orientation; at least for the dataset considered in this paper.

5 Conclusion

In this work we introduce the orientation of second-order image structure for feature description. We extend shape index histogram description by including the shape index orientation and evaluate our methods on the task of cell image classification. On the ICIP 2013 dataset, we are able to increase classification accuracy with around 3.5 percentage points (a relative error reduction of 14%) by including the shape index orientation. This exceeds the performance gains achieved by using classic gradient orientations indicating that second-order orientation captures useful image structure.

References

1. Arandjelovic, R., Zisserman, A.: Three things everyone should know to improve object retrieval. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2911–2918 (2012)
2. Crosier, M., Griffin, L.: Using basic image features for texture classification. International Journal of Computer Vision **88**(3), 447–460 (2010)
3. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2005, vol. 1, pp. 886–893 (2005)
4. Foggia, P., Percannella, G., Soda, P., Vento, M.: Benchmarking HEp-2 cells classification methods. IEEE Transactions on Medical Imaging **PP**(99), 1–1 (2013)
5. Foggia, P., Percannella, G., Saggese, A., Vento, M.: Pattern recognition in stained HEp-2 cells: Where are we now? Pattern Recognition **47**(7), 2305–2314 (2014)
6. Koenderink, J.J., van Doorn, A.J.: Surface shape and curvature scales. Image and Vision Computing **10**(8), 557–564 (1992)
7. Larsen, A.B.L., Vestergaard, J.S., Larsen, R.: HEp-2 cell classification using shape index histograms with donut-shaped spatial pooling. IEEE Transactions on Medical Imaging **33**(7), 1573–1580 (2014)
8. Lazebnik, S., Schmid, C., Ponce, J.: A sparse texture representation using local affine regions. IEEE Transactions on Pattern Analysis and Machine Intelligence **27**(8), 1265–1278 (2005)
9. Lindeberg, T.: Scale-space theory in computer vision. Springer (1993)
10. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision **60**, 91–110 (2004)
11. Nosaka, R., Ohkawa, Y., Fukui, K.: Feature extraction based on co-occurrence of adjacent local binary patterns. In: Ho, Y.-S. (ed.) PSIVT 2011, Part II. LNCS, vol. 7088, pp. 82–91. Springer, Heidelberg (2011)
12. Ojala, T., Pietikainen, M., Maenpaa, T.: Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. IEEE Transactions on Pattern Analysis and Machine Intelligence **24**(7), 971–987 (2002)
13. Pedersen, K., Stensbo-Smidt, K., Zirm, A., Igel, C.: Shape index descriptors applied to texture-based galaxy analysis. In: 2013 IEEE International Conference on Computer Vision (ICCV), pp. 2440–2447 (December 2013)
14. Snoek, J., Larochelle, H., Adams, R.P.: Practical bayesian optimization of machine learning algorithms. In: Pereira, F., Burges, C., Bottou, L., Weinberger, K. (eds.) Advances in Neural Information Processing Systems, vol. 25, pp. 2951–2959. Curran Associates, Inc. (2012)
15. Tola, E., Lepetit, V., Fua, P.: Daisy: An efficient dense descriptor applied to wide-baseline stereo. IEEE Transactions on Pattern Analysis and Machine Intelligence **32**(5), 815–830 (2010)
16. Wiliem, A., Sanderson, C., Wong, Y., Hobson, P., Minchin, R.F., Lovell, B.C.: Automatic classification of human epithelial type 2 cell indirect immunofluorescence images using cell pyramid matching. Pattern Recognition (2013)