2014 22nd International Conference on Pattern Recognition

An Orientation-Adaptive Extension to Scale-Adaptive Local Binary Patterns

Sebastian Hegenbart and Andreas Uhl
University of Salzburg
Department of Computer Sciences
shegen@cosy.sbg.ac.at and uhl@cosy.sbg.ac.at

***Abstract*—Methods based on Local Binary Patterns have been used successfully in a wide range of texture classification tasks. A restriction shared by all methods based on Local Binary Patterns is the high sensitivity to signal scale. In recent work we presented a general framework for scale-adaptive computation of Local Binary Patterns, improving the accuracy in texture classification scenarios involving varying texture-scales highly. In this work, the scale-adaptive methodology is extended by an orientation-adaptive computation of patterns, leading to a scale- and rotation-invariant classification. The results suggest that estimating a global orientation to build orientation-adaptive LBPs is superior to the previously introduced rotation-invariant encodings. The proposed framework allows the use of the highly-discriminative LBPs in less-constrained situations, where both orientation, as well as scale variations, are to be expected.**

I. Introduction

A big challenge in texture classification scenarios in unconstrained environments is dealing with varying scales and orientations. This is especially true in medical image acquisition such as endoscopy [3]. As a result, research focusing on scale- and rotation-invariant feature descriptors has been a hot topic in the recent past.

Methods based on Local Binary Patterns (LBP [9]) have been successfully used in a wide range of texture classification scenarios. A restriction shared by all those methods is their high sensitivity in terms of signal scaling, therefore reducing their applicability to a constrained environment with only minor scale variances among textures. The correct alignment of micro structures in terms of orientation is an essential requirement for the accuracy of the baseline LBP type methods. Ojala et al. [10] alleviate this restriction by using a special type of rotation-invariant pattern encoding, leading to a possibly reduced discriminative power of features. A drawback of that method is the limited angular resolution. As a consequence the rotation-invariant encoding is not very well suited in a scale-adaptive computation. In [2] we proposed a general scale-adaptive methodology that enables the use of the highly-discriminative LBPs in less-constrained situations, where scale variations are to be expected. Experiments have shown that this scale-adaptive framework improved the accuracy of LBP based methods in scenarios with varying scales significantly.

In this work we present an extension to this scale-adaptive framework, alleviating the restriction of correct texture orientation alignment by utilizing a global orientation-estimate. This allows the use of highly-discriminative LBPs in a scenario with varying scales and orientations.

By using multi-scale second moment matrices [7], a global orientation is estimated at dominant local scales, leading to a robust orientation estimation in noisy environments with varying texture scales. By leveraging the already pre-computed scale-spaces, our proposed orientation estimation approach integrates naturally with the scale-adaptive LBP framework at moderate computational cost. Employing the estimated orientation, an orientation-adaptive computation of LBP patterns is performed. Our results suggest that estimating a global orientation to build orientation-adaptive LBPs in a scale-adaptive computation is superior to the previously introduced rotation-invariant encodings.

In Section II we give a review of the general scale-adaptive computation that enables the use of LBPs in scenarios with varying scales. Section III-A describes the orientation-adaptive methodology. The fusion of the orientation- and scale-adaptive computation is covered in Section IV. The experiments conducted to evaluate the proposed methodology are described in Section V, the results presented and discussed in Section VI. Finally Section VII concludes the paper.

II. Scale-Adaptive Computation of LBP

The scale-adaptive computation is based on a global scale estimation combined with a confidence measure for the estimation. Based on the estimated scale, the radius of the LBP as well as the dimension of the sampling area is adapted accordingly. This methodology allows the use of LBP flavored methods in a scenario with varying scales.

A. Scale Estimation

We employ a global scale estimation algorithm which is based on scale-normalized Laplacian of Gaussian extrema in scale-space The scale-space theory was first extensively explored in the field of signal processing by Lindeberg [6]. It presents a framework to analyze signals at different scales. Let $f : \mathbb{R}^2 \mapsto \mathbb{R}$ represent a continuous signal, then the linear scale-space representation $L : \mathbb{R}^2 \times \mathbb{R}_+ \mapsto \mathbb{R}$ is defined by

$$L(\cdot;\sigma) = g(\cdot;\sigma) * f, \tag{1}$$

with initial condition $L(\cdot;0) = f$. Where $\sigma \in \mathbb{R}_+$ is the scale parameter, g is a Gaussian function and "$*$" denotes convolution. The scale-space family L is the solution to the diffusion equation (heat equation):

1051-4651/14 $31.00 © 2014 IEEE
DOI 10.1109/ICPR.2014.202

$$\partial_\sigma L = \sigma \left(\frac{\partial^2 L}{\partial x^2} + \frac{\partial^2 L}{\partial y^2} \right) = \sigma \triangle L. \tag{2}$$

We construct the scale-space and compute the scale-normalized Laplacians ($\sigma^2 \left|\triangle L(\cdot;\sigma)\right|$, denoted as $\overline{\triangle} I(\cdot;\sigma)$) of each image I at each location $x \in \mathbb{N}^2$ at different scales with $\sigma = c\sqrt{2}^k, k \in \{-4, -3.75, \ldots, 7.75, 8\}$ and $c = 2.1214$. Note that the parameter c acts as a scaling factor of the scale-space and was initially chosen such that the center scale equals a 3 pixel radius. We however found during experimentation that the intrinsic scale of natural textures tends to be large. We therefore expanded the scale-space to cover larger scales as well.

Methods based on scale selection employing the scale-space abstraction identify image locations which are simultaneously a local extremum with respect to both the spatial coordinates and the scale-space parameter (3D maxima), a prominent example is the Scale Invariant Feature Transform (SIFT [8]). Experimentation has shown however that the utilization of such locations for a global scale estimation is unreliable. This can be seen in Figure 1, comparing the distribution of the responses of the 3D maxima with the responses of the scale estimation of textures, the extrema are either at various different scales or only a small number of extrema is present, leading to unreliable scale estimations. We therefore use the distribution of responses of the scale normalized Laplacians to estimate a global scale. The scale estimation response function ξ is

$$\xi(t) := \sum_{x,y} \overline{\triangle} I(x, y; t). \tag{3}$$

The global scale is identified by searching for the first local maximum of ξ which is then used as seed point for a least-squares Gaussian fit. By using the first local maximum we are capable of consistently estimating the scale of textures exhibiting more than a single dominant scale. The quality of the estimation is improved by considering only data points within a certain offset from the seed point. In our implementation an offset of ± 5 scale levels is used to fit the Gaussian function. Finally the mean value $\tilde{s}$ of the fitted Gaussian function is interpreted as the dominant level in scale-space. The standard deviation u of the fitted Gaussian is used as uncertainty of the estimation. For a given dominant level in scale-space $\tilde{s_i}$,

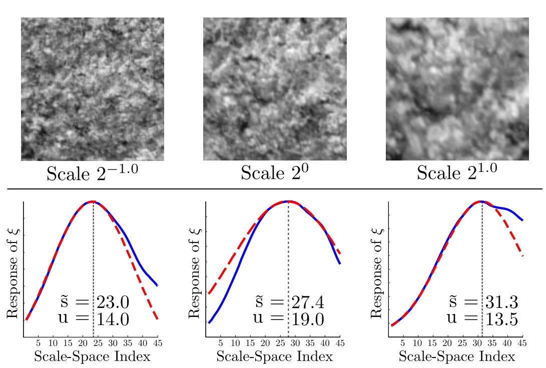

Fig. 2: Scale Estimation of a non-Regular Texture (stone2).

the spatial scale s_i corresponds to to the scale parameter σ_i at the dominant scale level. Figure 2 illustrates the fitted Gaussian function (dashed line) to the scale estimation response function (solid line) of three textures at different scales.

The response of the scale-normalized Laplacian of Gaussian (LoG) attains a maximum if the zeros are aligned with a circular shaped structure. Hence scales estimated, based on the LoG, correlate strongly with the scale of the dominant circular shaped structures of a texture. As a consequence, the estimated scale is highly related to an essential property of each texture, the *intrinsic scale* of a texture.

A texture exhibiting pebbles for example and a texture exhibiting sand, captured at the same distance, might have equal scales in terms of camera-scale, but different scales in terms of the scale-space, a consequence of different intrinsic scales. In contrast, sand and pebbles captured at a different camera-scales, corresponding to the difference of the textures' intrinsic scales, are equal in scale in terms of the scale-space. Scales estimated in the scale-space domain are therefore always a combination of the intrinsic texture scale and the camera-scale.

The identification of an intrinsic scale of a general texture is a non-trivial problem. A requirement on an intrinsic scale estimation method would be scale-invariance, a property that the LoG response in scale-space does not provide. The estimated scale in scale-scape is therefore always a combination of camera-scale and intrinsic texture scale. Please refer to [3] for more details.

III. Orientation-Adaptive Local Binary Patterns

The correct alignment of micro structures in terms of orientation is an essential requirement for the accuracy of the baseline LBP type methods. Ojala et al. [10] alleviate this restriction by using a special type of rotation-invariant pattern encoding. The original pattern is shifted circularly until a minimum with respect to a numeric interpretation of the pattern is found. As a consequence all patterns are implicitly aligned among each other.

A drawback of this approach is the limited angular resolution, depending on the number of used LBP-neighbors. For a standard LBP-neighborhood with 8 neighboring samples, this angular resolution corresponds to 45 degrees. A side-effect of the encoding is the decreased number of individual patterns. The authors propose to use uniform patterns in combination with the rotation-invariant encoding to improve robustness, implicitly improving the angular resolution by considering only special type of micro structures. Uniform patterns are a subset of patterns with a maximum of two transitions between 1 and 0. The proposed combination of rotation-invariant and uniform patterns reduces the number of individual patterns even further. Experiments discussed in Section VI show that the small number of individual patterns leads to a decreased classification accuracy if combined with the scale-adaptive computation. As a consequence we utilize the estimation of a global orientation to build orientation-adaptive LBP. Following literature on LBP, we refer to Local Binary Patterns utilizing the rotation-invariant encoding in combination with uniform pattern as LBP^{riu} from here on.

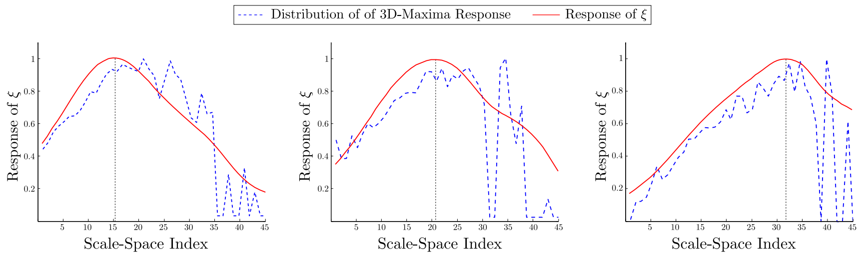

Fig. 1: Distribution of 3D-Maxima compared to the Response of ξ.

A. *Orientation Estimation*

We utilize the multi-scale second moment matrices [7] of an image, computed at dominant local scales, for a robust orientation estimation in noisy environments with varying texture scales. The second moment matrix (also known as structure tensor) summarizes the predominant directions of the gradient in a specific pixel neighborhood of an image. In contrast to the second moment matrix, the multi-scale second moment matrix is defined over two scale parameters. It allows to estimate the shape of visual structures at their dominant scale, as detected by the scale-estimation algorithm.

The local scale, denoted by t determines the scale in terms of the scale-space a local structure is analyzed at. The integration scale i is used as parameter to a Gaussian function g defining the shape and weights of a specific neighborhood area in the image over which the gradient response is accumulated. We compute the multi-scale second moment matrices at each location $x \in \mathbb{R}^2$ of an image I. The local scale t is selected depending on the estimated texture scale (see Section IV), the integration scale $i = \sqrt{2}t$ depends on the corresponding detection scale. The second moment matrix for an image location x at local scale t is then computed as

$$\mu(x;t,i) = \int_{\xi \in \mathbb{R}^2} (\nabla I)(x-\xi;t)(\nabla I)^T(x-\xi;t)\, g(\xi;i)\, d\xi. \tag{4}$$

We denote $(\nabla I)(x;t)$ as the gradient of the scale-space representation of image I at scale t and position x.

The multi-scale second moment matrix is positive definite, it therefore has two non-negative eigenvalues which correspond to the length of the axes of an ellipse (up to some constant factor). The eigenvectors of the multi-scale second moment matrix correspond to the orientation of the ellipse. By computing the angle between the major axis of the ellipse and the vertical axis we identify the dominant orientation at a specific image position. Due to the ambiguous orientation of the ellipse, all angles are treated modulus 180.

Based on the distributions of all orientations at all pixel locations, a global orientation is estimated for an image. This is done by fitting a Gaussian function to the distribution of orientations. To improve the quality of the estimation, we remove data points with an offset of ± 40 degrees from the maximum prior to the fitting process. Finally, the average value of the Gaussian is interpreted as the dominant orientation, the standard deviation of the fitted Gaussian function is interpreted as the uncertainty of the estimation. To avoid using invalid orientation estimations, we reject estimations with an uncertainty above 20. In such a case the estimated orientation is defined as 0 degrees.

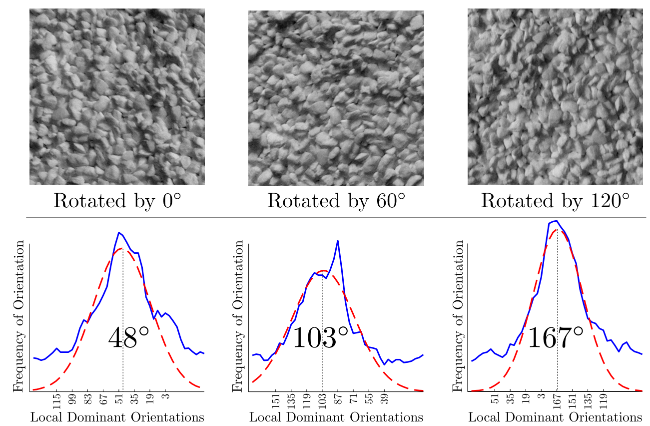

Fig. 3: Orientation Estimation (pearlsugar1).

Figure 3 illustrates the fitting of a Gaussian function (dashed red line) to the distribution of orientations (solid blue line) of three differently rotated images. The numbers centered at each Gaussian function relate to the mean value of the fitted Gaussian function, corresponding to the estimated global orientation of the specific image.

B. *Global versus Local Orientation Estimation*

As explained in Section III-A a global orientation is computed for a specific image. In theory however, a texture could consist of multiple sub textures with different orientations, leading to potentially worse estimation accuracies. We therefore evaluated the performance of a local orientation estimation on a pixel basis in comparison to the used global orientation estimation. The local orientation estimation is based on the same methodology utilizing multi-scale second moment matrices as described for the global orientation estimation. In contrast to the global orientation however, the estimation is done per pixel instead of fitting a Gaussian function to the distribution of orientations to estimate a global orientation. Figure 4 demonstrates that the accuracy of the local estimation is inferior as compared to the global estimation. The mean absolute error of the estimated orientation (vertical axis) was computed between a reference image without rotation and the same image with a specific rotation, as depicted by the horizontal axis, for all images in the Kylberg database which was also used for experimentation as explained in Section V. The mean absolute error was computed for three relative scales between the reference and the rotated images. We can see that the global estimation is superior to the local estimation in all regards. We assume that this is caused by homogeneous pixel areas which do not allow for a robust estimation of orientation, introducing a large error. The results also indicate that scaling of the textures has only a minor impact to the general accuracy of the orientation estimation method, an important property for using the method in combination with the scale-adaptive methodology.

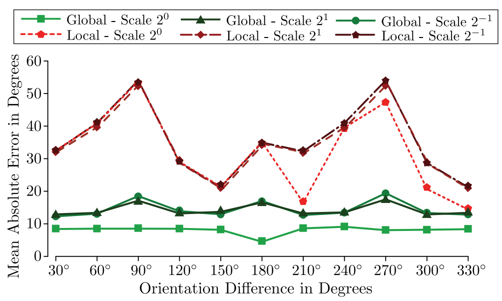

Fig. 4: Global versus Local Orientation Estimation Error.

C. Impact of Signal Noise

Utilizing the multi-scale second moment matrix allows to estimate the orientation of a visual structure at its dominant scale, as a benefiting side effect of utilizing the scale-space data, signal noise is suppressed to some degree. We explicitly constructed an experiment to evaluate this property. The mean absolute error of the orientation estimation is computed for noisy image textures at the same texture scale. Let P be the set of all pixels in image $I \in \mathbb{N}^2$, $\omega = (\omega_p)_{p \in P}$, be a collection of independent identically distributed real-valued random variables following a Gaussian distribution with mean m and standard deviation σ. We simulate thermal noise as additive Gaussian noise with $m = 0$, variance σ for pixel p at position x, y as

$$N(x, y) = I(x, y) + \omega_p, \quad p \in P, \tag{5}$$

with N being the noisy image, for an original image I. Figure 5 illustrates the effects of Gaussian white noise to the global orientation estimation. We see that noise only has a minor impact to the average accuracy of the method, another welcome benefit of using multi-scale second moment matrices for orientation estimation.

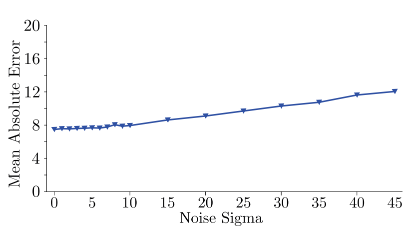

Fig. 5: The Impact of Signal Noise to the Estimation.

IV. Combining the Scale-Adaptive Computation with the Orientation-Adaptive Computation

The orientation-adaptive computational approach integrates very naturally into the scale-adaptive LBP framework. As a consequence of computing the LoG for scale-estimation instead of using the Difference of Gaussians approach, the scale-space data can be re-used for computing the multi-scale second moment matrices used for orientation estimation. Therefore the Gaussian filtering to compute the local scale t can be omitted. We adaptively select the local scale t of the multi-scale second moment matrix, based on the estimated scale of a texture. By doing so, we guarantee a robust orientation estimation across different texture scales. Experimentation has shown that a reasonable value for the local scale t is half of the estimated texture scale. This is explained by the property that the estimated scale at a pixel level highly correlates to the intrinsic scale of a texture, therefore leading to rather large estimated scales. Large local scales however would result in a decreased estimation accuracy. By re-using the scale-space data, the computation of the multi-scale second moment matrices only involves the computation of the first partial derivatives in both image dimensions as well as a convolution with a Gaussian filter to compute the integration scale i. Figure 6 illustrates schematically how the scale- and orientation-adaptive computation is combined. Based on the estimated texture scale, appropriate LBP radii and neighborhood sampling area dimensions are chosen. The ordering of the computation is adaptively chosen depending on the estimated orientation. Please note that due to the ambiguity of the orientation, we compute two patterns at each image location, rotated circularly to accommodate orientations above 180 degrees. This is indicated by the red sampling points which correspond to the respective starting location of the computation.

To compensate for possible errors of the orientation estimation as well as unsuitable alignments on the pixel grid, we compute multiple LBP histograms using a small interval of different orientations depending on the estimated orientation. The size of the interval is chosen depending on the uncertainty measure of the orientation estimation. As a consequence the chosen interval for an unreliable orientation estimation is wider and the error is more likely to be compensated. The interval width chosen for the experiments discussed in Section V was 0.7 times the standard deviation (interpreted in degrees) of the fitted Gaussian. This value was not optimized and might be dependent on the given problem however. For each orientation

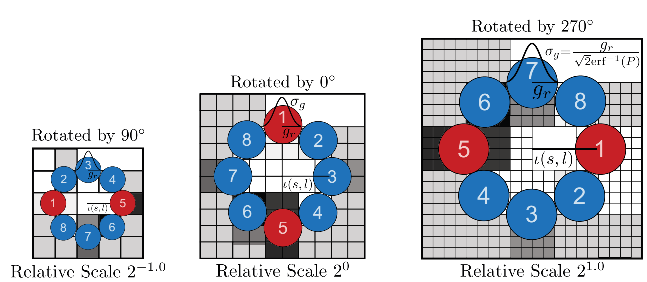

Fig. 6: Illustration of the Scale- and Orientation-Adaptive Computation.

in the interval (in steps of 5 degrees), a separate LBP histogram is computed. Finally, the best alignment of orientations is implicitly chosen during classification by selecting the minimum of all distances computed pairwise between all LBP histograms computed based on the specific orientation within the intervals of two texture images. Note that this does not pose an unfair advantage to the classifier, as no information about class membership is used implicitly or explicitly. This approach is comparable to the cyclic shift of binary iris features used to compensate for small rotational alignment errors in biometric systems for example.

V. EXPERIMENTS

We constructed a large set of experiments to analyze the performance of the orientation-adaptive extension to the scale-adaptive LPB framework in a scenario with varying scales and varying image rotations. We explicitly compare the accuracy of LBP^{riu} methods employing the scale-adaptive methodology with the accuracy of the standard LBP methods employing the proposed scale- and orientation-adaptive framework. Additionally we analyze the performance of non scale-adaptive methods based on LBP^{riu} in the same scenario.

The used methods are the LBP^{riu} method [10], the Local Ternary Patterns (LTP^{riu}) operator [11] as well as the Fuzzy Local Binary Patterns (FLBP^{riu}) method [4]. Please note that these methods were used in combination with the multi-resolution Local Binary Patterns extension [10] based on three scales and 8 neighbors, the best configuration we were able to find for the given data sets.

The experimentation is based on two independent texture databases. The KTH-TIPS database [1] exhibits texture images from 10 different materials captured at 9 different scales with 9 samples per material. Sub-images of size 128×128 pixels were extracted from the center of each accordingly rotated original image. The rotation was performed using bilinear interpolation. We simulated rotations of 30 degrees, 60 degrees 120 degrees and 180 degrees respectively. Due to the dimension of the original images of material "cracker“, this material class could not be used for simulating rotation without a large black border within the 128×128 image patches. Unfortunately, besides KTH-TIPS there are no other publicly available high quality texture databases with an available ground-truth of scales. We therefore had to resort to a simulation of the scaling of textures. A subset of the Kylberg texture database [5], consisting of 28 materials with 160 unique texture patches per class, captured at a single scale, was used for the simulation. The image database contains rotated versions of each image at 30 degree steps ranging from 0 to 330 degrees. The high resolution of each patch (576×576 pixels) allowed us to simulate the scaling without relying on up-sampling, leading to a smaller amount of interpolation artifacts. The simulation of scaling was performed according to the scales of the KTH-TIPS database, interpreting the original image patches as the maximum scale $2^{1.0}$. Image patches of size 128×128 were then extracted from the center of the re-scaled patches. Due to the huge number of samples in the Kylberg database we use a subset consisting of 20 unique texture patches per class (5 patches per image) for experimentation.

The experiments were designed to explicitly reflect the properties of the studied methods. The images from the KTH-TIPS database at scale 5 without rotations build the training set for experiments based on the KTH-TIPS database. Respectively the images from the Kylberg database at scale 2^0 without rotation are used as training data for experiments based on the Kylberg database. To evaluate the impact of rotation and scale, the classification was performed on the corresponding scaled and rotated version of the data from each of the databases. The used classification method was a k-nearest neighbors classifier. The maximum value of k was chosen depending on the number of samples per material class. In case of the Kylberg database the maximum value of k was set to 20 while the maximum value of k was 9 in case of the KTH-TIPS database. To allow for an unbiased evaluation, all interpreted results are the mean accuracy over all possible k-values ranging from 1 to the specific maximum.

VI. RESULTS

Figure 7 presents the results of the experiments. The horizontal axis denotes the relative scale difference between training data and evaluation data while the vertical axis corresponds to the classification accuracy. The bold lines show the mean classification accuracy over all image rotations (5 different rotations for the KTH-TIPS set and 12 for the Kylberg database). We visualize the minimum and maximum classification accuracy over all rotations with error bars in case of the KTH-TIPS database as well as a smaller error bar with the corresponding area in case of the Kylberg database.

Methods utilizing the proposed scale- and orientation-adaptive methodology are abbreviated as *SOA* and the respective name of the used LBP based method, the scale-adaptive method. Methods utilizing the scale-adaptive methodology in combination with the rotation-invariant encoding are abbreviated as *SA* and the specific method's name. The name of the methods based on LBP are used as known from literature.

The difference in mean classification accuracy between the proposed scale- and orientation-adaptive (SOA) framework and the scale-adaptive (SA) framework using the rotation invariant encoding is reflected by the upper row of numbers. The lower row of numbers label the difference between the scale-adaptive framework based on LBP^{riu} with the respective standard method.

Figure 7 shows that the mean classification accuracy of methods utilizing the scale- and orientation-adaptive methodology (SOA) are superior in terms of classification accuracy and variance as compared to methods utilizing the scale-adaptive (SA) framework with rotation-invariant encoding. Comparing the results with prior experiments in [2], we see that the accuracy of the standard methods decreased due to the rotation invariant encoding. This is reflected by the fact that the maximum results are below the results in [2]. In general we observe a minor degree of variation caused by the different orientations across the results. Interestingly the methods employing the scale-adaptive framework (SA) exhibit the highest degree of variance, a property we expected due to the reduced discriminative power as discussed in Section III. Methods utilizing the proposed scale- and orientation-adaptive (SOA) framework show the smallest degree of variance with respect to orientation. Additionally the mean classification accuracy is clearly above the accuracy of traditional methods as

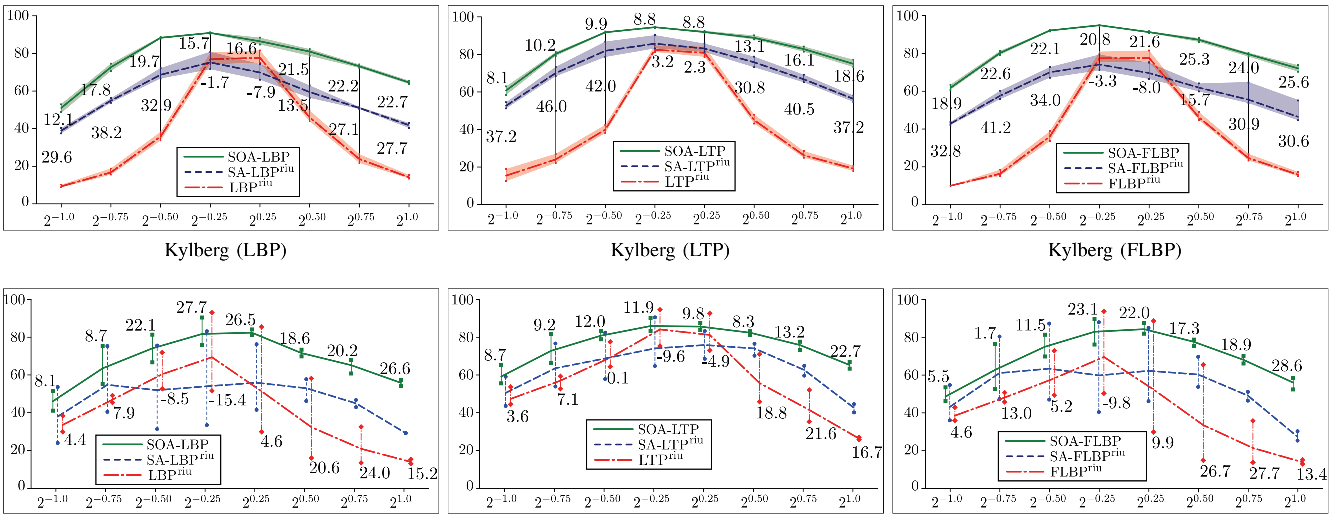

Kylberg (LBP) Kylberg (LTP) Kylberg (FLBP)

KTH-TIPS (LBP) KTH-TIPS (LTP) KTH-TIPS (FLBP)

Fig. 7: Mean Overall Classification Accuracies over all Rotations.

well as methods utilizing the scale-adaptive (SA) approach in combination with the rotation-invariant encoding. Concerning the results based on the KTH-TIPS database, we can see that the variation caused by rotation is considerably higher across all methods. The smallest variations caused by rotation is again observed for methods utilizing the proposed scale- and orientation-adaptive (SOA) methodology. In parallel to the Kylberg database, methods based on the SOA framework show the highest mean accuracy. We observe the highest amount of variance of methods utilizing the SOA methodology at small relative scales. We assume this is caused by the higher impact of the orientation estimation error for textures at a smaller relative scale. In general, the trends observed for the Kylberg database are confirmed by the results based on the KTH-TIPS database.

VII. Conclusion

In this work, we presented an orientation-adaptive extension to the scale-adaptive LPB framework. By leveraging the already pre-computed scale-spaces, our proposed orientation estimation approach integrates naturally with the scale-adaptive LBP framework at moderate computational cost. In particular, using multi-scale second moment matrices, computed at dominant local scales, leads to 1) robust orientation estimation in noisy environments and 2) scenarios with varying texture scales. Our experiments suggest that estimating a global orientation to build orientation-adaptive LBPs is superior to the previously introduced rotation-invariant encodings; this is reflected by less variance in classification accuracy as well as superior mean accuracy over multiple orientations. In summary, the proposed framework enables the use of the highly-discriminative LBPs in less-constrained situations, where both orientation as well as scale variations are to be expected.

References

[1] E. Hayman, B. Caputo, M. Fritz, and J.-O. Eklundh. On the significance of real-world conditions for material classification. In *Proceedings of the 8th European Conference on Computer Vision (ECCV)*, volume 3024 of *Lecture Notes in Computer Science*, pages 253–266. Springer, 2004.

[2] S. Hegenbart and A. Uhl. A scale-adaptive extension to methods based on lbp using scale-normalized laplacian of gaussian extrema in scale-space. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing ICASSP '14*, pages 4352–4356, 2014.

[3] Sebastian Hegenbart, Andreas Uhl, Andreas Vécsei, and Georg Wimmer. Scale invariant texture descriptors for classifying celiac disease. *Medical Image Analysis*, 17(4):458 – 474, 2013.

[4] D. Iakovidis, E. Keramidas, and D. Maroulis. Fuzzy local binary patterns for ultrasound texture characterization. In *ICIAR*, volume 5112 of *Lecture Notes in Computer Science*, pages 750–759. Springer, 2008.

[5] G. Kylberg. The kylberg texture dataset v. 1.0. External report (Blue series) 35, Center for Image Analysis, Swedish University of Agricultural Sciences, Uppsala University, Uppsala, Sweden, September 2011.

[6] T. Lindeberg. Scale-space for discrete signals. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(3):234–254, 1990.

[7] Tony Lindeberg. *Scale-Space Theory in Computer Vision*. Kluwer Academic Publishers, Norwell, MA, USA, 1994.

[8] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, November 2004.

[9] T. Ojala, M. Pietikäinen, and D. Harwood. A comparative study of texture measures with classification based on feature distributions. *Pattern Recognition*, 29(1):51–59, January 1996.

[10] T. Ojala, M. Pietikäinen, and T. Mäenpää. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(7):971–987, July 2002.

[11] X. Tan and B. Triggs. Enhanced local texture feature sets for face recognition under difficult lighting conditions. In *Analysis and Modelling of Faces and Gestures*, volume 4778 of *Lecture Notes in Computer Science*, pages 168–182, October 2007.