Boosting Small-Data Performance of LBP: A Case Study in Celiac Disease Diagnosis

Michael Gadermayr[1], Andreas Uhl[1], and Andreas Vécsei[2]

[1] Department of Computer Sciences, University of Salzburg, Austria
{mgadermayr,uhl}@cosy.sbg.ac.at
http://www.wavelab.at
[2] St. Anna Children's Hospital, Department of Pediatrics, Medical University Vienna, Austria

Abstract. A major issue in computer aided celiac disease diagnosis is the prevalence of substantial intra-class and even intra-image variations. A method which splits the images into a set of smaller ones and finally applies a decision level fusion turned out to be a powerful technique to address these problems and to boost the classification accuracy. This is especially true if using Local Binary Patterns and derivatives. However, due to the sparsity and roughness of the final feature vectors, these methods are not optimal if being applied to such small images. Therefore, in this work two novel and two methods from literature are investigated to improve the performances of Local Binary Patterns. Experiments show that the overall classification accuracies can be improved, especially by means of a combination of the novel methods. The techniques presented in this work are not restricted to this certain problem definition. They rather can be applied in arbitrary scenarios with small sized image data.

1 Introduction

Celiac disease [14], which is commonly known as gluten intolerance, is a disorder that affects the small intestine after introduction of gluten containing food. The disease leads to an inflammatory reaction in the mucosa of the small bowel caused by a dysregulated immune response triggered by ingested gluten proteins of certain cereals. During the course of celiac disease, the mucosa loses its absorptive villi (see Fig. 1(b)) and hyperplasia of the enteric crypts occurs, leading to a strongly diminished ability to absorb any nutrients. According to a large study [6], the overall prevalence of the disease in the USA is 1:133. Figure 1 shows example images, captured during endoscopy.

Up to now, significant work has been done on computer aided celiac disease diagnosis [3, 4, 9, 10]. Especially Local Binary Patterns (LBP) [15] and derivatives of this well known texture feature extraction method have been extensively investigated with respect to this problem definition [9–11] and turned out to be highly effective.

In recent work on computer aided celiac disease diagnosis [8], the authors have proposed an effective and efficient split and merge approach which splits a textured image into non-overlapping sub-images, classifies these sub-images and finally applies a decision level fusion. Especially with LBP (and derivatives), this method turned out to be highly appropriate for the problem definition. This is quite surprising, because the sub-images are very small (e.g. a sub-image size of 42×42 pixels turned out to be appropriate) and LBP is not designed to be applied to such small data. Applying the more sophisticated bag-of-visual words [19] technique the achieved accuracies are considerably lower. However, no matter if using LBP in combination with the split and merge approach, bag-of-visual words [19] or state-of-the art fisher vectors [16, 17], the problem of small data is the same for all of these approaches:

A small sample size leads to a less precise estimation of the probability density of the LBP patterns which is given by the histogram. The generated histograms become potentially sparse as well as rough, depending on the chosen setup. Table 1 shows the average number of patterns per bin, for each number of center-pixel neighbors. It should be noticed that the distribution of LBP patterns is far away from being uniformly. We supposed that the sparsity of the histogram affects the classification performance in case a high number of neighbors (especially with twelve but also in case of ten and eight neighbors).

Table 1. Histogram bins and average number of patterns per bin in case of different setups having a 42×42 pixel image and an LBP radius of three pixels (i.e. $36 \times 36 = 1296$ patterns in total).

LBP Neighbors	Histogram bins	Mean patterns per bin	Sufficient patterns per bin? (assumption)
4	16	81.0	✓
6	64	20.3	✓
8	256	5.1	?
10	1024	1.3	?
12	4096	0.3	✗

To compute the original LBP [15] feature vector, first a binary vector for each pixel is generated by computing the sign of the differences between this pixel and a set of neighboring pixels. Each of these binary patterns can be interpreted as a number by multiplying the binary row vector with a column vector consisting of increasing powers of two $(1, 2, 4, 8, ...)$. The final feature vector consists of the

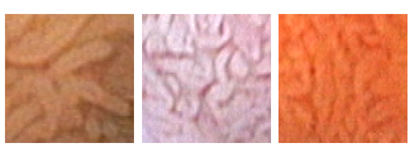

(a) Marsh 0: Healthy mucosa.

(b) Marsh 3: Villous atrophy.

Fig. 1. Example patches of healthy (a) and disease mucosa (b).

global histogram computed over all of these numerical values in an image. In the last decades several derivatives of LBP have been proposed including Extended Local Binary Patterns (ELBP) [13] and Local Ternary Patterns [18].

In this work, focus is on four independent techniques to increase the performance in case of small sized image data. The first one changes the general algorithm of LBP using soft histograms (Sect. 2.1), the second one concatenates specific low dimensional features (Sect. 2.2), the third one relies on image preprocessing (Sect. 2.3) and the last is based on feature post-processing (Sect. 2.4). Whereas the first two methods are existing approaches from literature [1, 2], the second two techniques are newly introduced in this paper. In Sect. 3 experimental results are presented and discussed.

2 Boosting LBP's Performance

In this section, four independent methods are outlines to boost the performance of LBP (and derivatives) if being applied to small data.

2.1 Soft-Histogram (SH)

To increase robustness to noise and make its output continuous, Ahonen and Pietikäinen have proposed soft histograms for LBP [1]. Instead of assigning a pixel in an image to exactly one histogram bin, in this fuzzified approach one pixel contributes to a number of bins. The contribution of a pixel (x, y) to a bin i is given by

$$SLBP(x,y,i) = \prod_{p=0}^{P-1} \left[b_p(i) \cdot f_d(g_c - g_p) + (1 - b_p(i)) \cdot (1 - f_d(g_c - g_p)) \right] , \quad (1)$$

where P is the number of neighbors, $b_p(i)$ denotes the value of the p-th bit of i, g_c is the current center pixel and g_p is one of the neighboring pixels. f_d is the fuzzy membership function

$$f_d(z) = max(min(0.5 + 0.5\frac{z}{d}, 1), 0), \quad (2)$$

where d regulates the extent of fuzzification. For small d this method converges to the traditional LBP. Although this method is not dedicated to small image data, the soft assignment can be utilized to generate smoother histograms.

2.2 Multi-Neighborhood (MN)

In another recent work, Banerji et al. have proposed a method [2] to deal with small image data in a bag-of-words model. Instead of introducing a new methodology, the authors concatenate eight LBP histograms with varying small neighborhoods. The exact patterns are shown in Fig. 2. By concatenating feature vectors with few neighbors (four), each histogram only consists of 16 ($= 2^4$) bins which should be advantageous in case of small data. By concatenating eight feature vectors with different neighborhoods, a higher distinctiveness is claimed to be achieved.

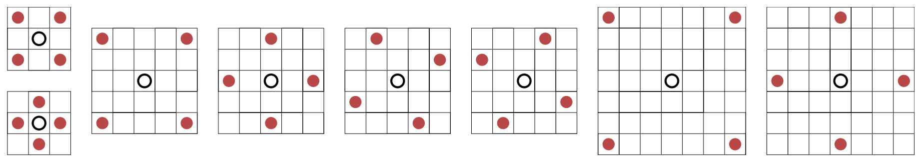

Fig. 2. The eight four-neighborhoods utlized by Banerji et al. [2].

2.3 Image-Enlargement (IE)

Whereas the methods mentioned so far change the feature extraction, we furthermore investigate the impact of an image enlargement, prior to the feature extraction stage. For this, the image is upscaled and missing points are bi-linearly interpolated. Thereby the number of pixels is increased by the square of the scaling factor which directly leads to substantially denser histograms. More sophisticated interpolation techniques (bi-cubic, spline) have also been tested, however, the obtained classification outcomes are highly similar with all methods.

2.4 Histogram-Smoothing (HS)

Finally we propose a histogram improvement technique which can be applied as a post processing method after LBP histogram generation. The main idea is to compensate the roughness of histograms by means of smoothing, which, to the best of our knowledge, has not been done before. One special motivation for this method is to achieve a similar behavior like Soft-Histogram LBP with a substantially lower computational expense.

The crucial thing is that an LBP histogram cannot be filtered by straight forward convolution with for example a Gaussian filter to obtain a smoother version. The problem is that neighboring histogram bins not necessarily exhibit a strong logical relationship. This issue is illustrated in Fig. 3.

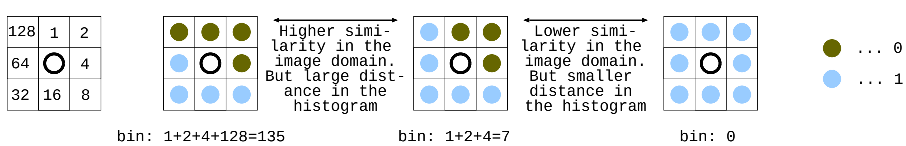

Fig. 3. Histogram smoothing.

To cope with this issue, we construct a Matrix $M = m_{(i,j)}$ which defines the degree of similarity between all combinations of two binary patterns. This matrix is defined by means of the inverse hamming distance

$$m_{(s,t)} = \left(\frac{1}{P} \cdot \sum_{p=1}^{P} 1 - |b_p(s) - b_p(t)| \right)^k , \tag{3}$$

where the positive value k adjusts the degree of smoothing by suppressing small values. A small k (e.g. $k = 1$) corresponds to extensive smoothing and vice versa.

Finally the smoothed histogram H^* is defined by

$$H^*(s) = \frac{1}{\sum_{t=1}^{2^P} m_{(s,t)}} \cdot \sum_{t=1}^{2^P} H(i) \cdot m_{(s,t)} \ . \tag{4}$$

This equation states that each bin contributes to each other bin in the histogram by the proportion specified by M. In the trivial case, where M is the identity matrix, smoothing is omitted and H^* is equal to H.

It should be mentioned that the matrix m could theoretically also be defined in a different way. However, we suppose that our definition based on the hamming distance poses a quite natural and plausible one.

2.5 Runtimes

Besides their high distinctiveness, LBP have become popular because of their low computational costs compared to more elaborated techniques. Therefore, we will briefly highlight the runtime[3] of the methods, presented in this section. Figure 4 shows the overall runtimes for computing the feature vector for one image (42×42 pixels), for different numbers of neighbors. Notice that the Multi-Neighbor LBP are only defined for a certain neighborhood and thereby cannot be configured, which results in a constant runtime in the figure. Whereas the Soft-Histogram approach corresponds to quite high computational costs, Histogram-Smoothing turned out to be distinctly faster. It should be noticed that the costs for the Soft-Histogram approach increase even more (linearly) with an increasing image size. The Histogram-Smoothing technique on opposite remains quite stable with an increasing image size, because the histogram creation is extremely fast (see costs of traditional LBP).

3 Experiments

3.1 Setup

The image testset used for the experiments contains images of the duodenal bulb and the pars descendens, which are parts of the small bowel, taken during duodenoscopies at the St. Anna Children's Hospital using pediatric gastroscopes (Olympus GIF N180 and Q165) (with a resolution of 768×576 and 528×522 pixels). Prior to processing, all images are converted to gray scale images because the additional use of color information did not lead to any substantial improvements. In a preprocessing step, texture patches with a fixed size of 128×128 pixels have been manually extracted. These patches are split into nine non-overlapping smaller sub-images with a size of 42×42 pixels. This splitting strategy, based on

[3] Tests are performed on an Intel(R) Core(TM) i5-2400 CPU @ 3.10GHz. The code has been implemented in C/MEX.

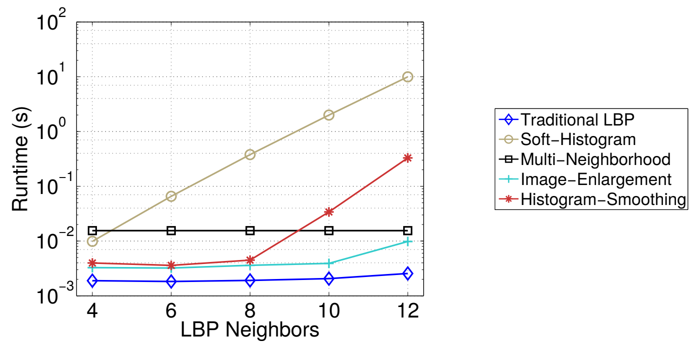

Fig. 4. Execution runtimes.

splitting into 9 sub-images, turned out to be optimal in recent work on computer aided celiac disease diagnosis [8]. After splitting an image, feature extraction and classification is executed separately for each sub-image. Finally, the decisions are merged by means of majority voting. This strategy turned out to be more appropriate than a straight forward classification of the complete images. This is supposed to be due to the high intra-class and intra-image variations of the endoscopic images.

To get the ground truth for the texture patches, the condition of the mucosal areas covered by the images has been determined by histological examination of biopsies from corresponding regions. The severity of the villous atrophy has been classified according to the modified Marsh classification scheme [14]. Although it is possible to distinguish between different stages of the disease, we aim in distinguishing between images of patients with (Marsh-3) and without the disease (Marsh-0), as this two classes case is most relevant in practice. Our experiments are based on a data set containing 612 images (306 Marsh-0 and 306 Marsh-3 images) from 171 (131 Marsh-0 and 40 Marsh-3) individuals [12].

All overall accuracies computed are based on the mean accuracy of 32 random splits. One distinct split divides the data set into an approximately balanced training (50 %) and evaluation set (50 %), restricting images of one patient to be in the same set to avoid any bias.

Focus is on two different LBP versions which turned out to be suitable for celiac disease classification [7]:

– Local Binary Patterns (LBP) [15]: This is the most common LBP method based on a circular neighborhood and a certain number of sample points, which are equidistantly placed on the circle. Sample points which are in between pixel values are interpolated in a nearest neighbor sense. This is done, as previous work on computer aided celiac disease diagnosis [7] showed that the bi-linear interpolation, which is usually utilized, corresponds to a loss of accuracy.

– Extended Local Binary Patterns (ELBP) [13]: ELBP in this context consists of the LBP feature extraction based on the edge magnitude of the original image, computed by means of convolution with two orthogonal Sobel filters.

The value d in case of Soft-Histogram LBP is fixed to 12 and k in case of Histogram-Smoothing is fixed to 4, which turned out to be appropriate for all configurations. In case of Image-Enlargement, a resize factor of 2 turned out to be optimal. The restriction to (lower dimensional) uniform patterns did not lead to improvements as far as the classification accuracy is concerned. For feature discrimination, we deploy the linear support vector classifier [5] (SVM) which has been widely used in recent work.

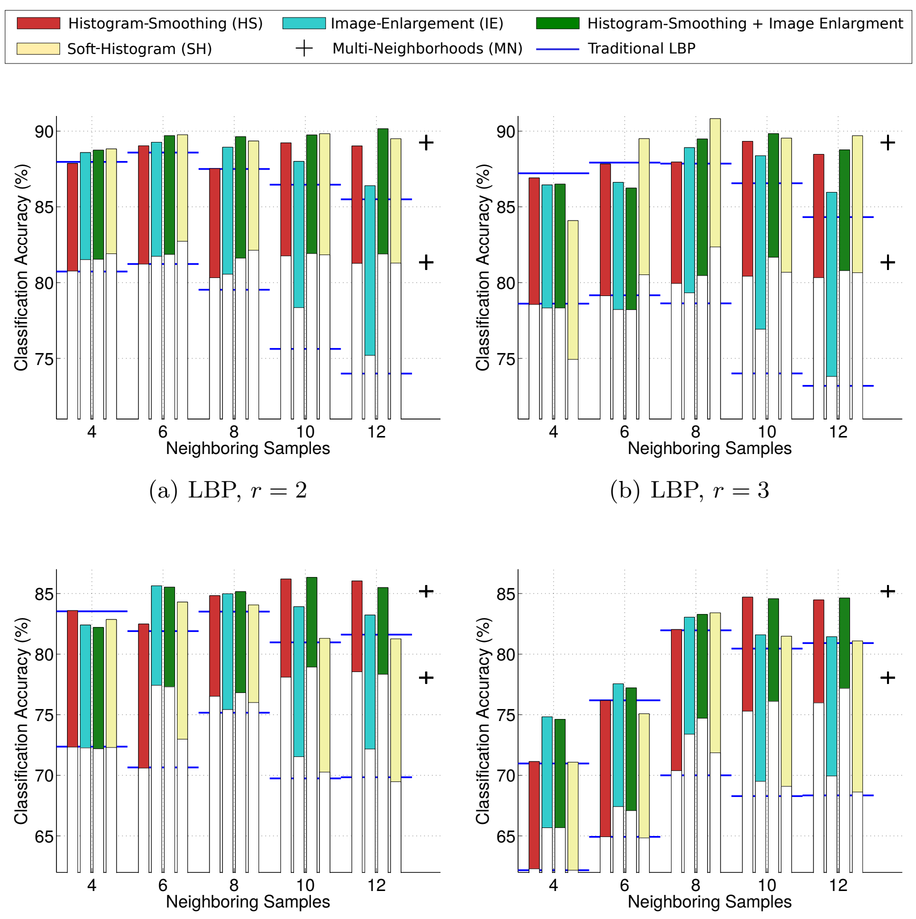

(a) LBP, $r = 2$ (b) LBP, $r = 3$

(c) ELBP, $r = 2$ (d) ELBP, $r = 3$

Fig. 5. Overall classification accuracies with different feature extraction methods and different configurations. For explanation see Sect. 3.2.

3.2 Results and Discussion

In Fig. 5, the main results are shown. One subplot shows the overall classification accuracies for a certain feature, a certain neighborhood (two or three pixels), five numbers of neighbors and the investigated approaches. The top of the colored bars indicate the rates achieved if the nine sub-images of one complete image are fused, whereas the top of the white bars indicate the rates achieved without the decision level fusion. The solid lines indicate the accuracies achieved with traditional LBP based classification. In a similar manner, the top lines indicate the rates achieved with fusion and the bottom lines indicate the rates achieved without fusion. Similarly, the crosses (+) denote the rates obtained with the Multi-Neighborhood approach which has a fixed neighborhood. In the following, focus is on the fusion based (upper) accuracies that are more relevant in practice, however, there is a strong correlation between these values and the lower values without the decision level fusion.

Considering the LBP based feature extraction (solid lines) with a radius of one and two, the best accuracies are achieved in general with four, six or eight neighbors, which is not surprising if considering Table 1. With ten and especially with twelve neighbors, the classification accuracies decrease. A less distinct but still similar behavior is shown in case of ELBP. Regarding the Histogram-Smoothing method, it can be seen that the performance with ten and twelve neighbors can be improved consistently, whereas the fewer neighbored versions stay almost unchanged. In each case of the 10 to 12 neighbored versions, this approach is able to outperform the best traditional configuration. The less elaborated Image-Enlargement technique has a similar effect on the classification performance, however, the correlation between the number of neighbors and the improvement is less distinct and the improvements are weaker. The third bars (per bunch) indicate the accuracies obtained if combining Histogram-Smoothing with Image-Enlargement. Considering LBP versions with eight to twelve neighbors, this combination leads to even better accuracies. In case of ELBP, the image enlargement step obviously is less important as Histogram-Smoothing mostly cannot be outperformed by the combined method. Interestingly, the more computationally complex Soft-Histogram LBP on average corresponds to slightly lower rates. This method seems to be less appropriate for ELBP, however, in the special case of LBP with a radius of three pixels this technique generates the best overall results. This quite interesting behavior is supposed to be due to the larger differences between the pixels (in case of the larger radius), which can be directly exploited by the fuzzy histogram creation. Due to the binary quantization, the post-processing Histogram-Smoothing is (like traditional LBP) unable to exploit any additional information given by pixel differences. However, the major drawback of this method is the comparatively high computational effort (see Fig. 4). Therefore, we recommend to utilize the Histogram-Smoothing approach instead. Especially in case of larger images, the runtime of the Soft-Histogram method increases linearly with respect to the image pixels whereas the runtime of the Histogram-Smoothing is mainly related to the number of histogram bins. Although the Multi-Neighborhood LBP could be interpreted as a multi-resolution method (as thereby actually cannot be directly compared to the others) it is unable to outperform the best of the other configurations. Obviously the distinctiveness of high dimensional joint distribution cannot be achieved by a concatenation of lower dimensional joint distributions. Nevertheless, this feature is able to outperform the best classification rates obtained with traditional LBP.

4 Conclusion

We have proposed two novel and comparatively fast methods to improve the classification performances of Local Binary Patterns and derivatives, if being applied to small images. Especially if combining these techniques, reasonable improvements can be obtained, compared to traditional classification based on Local Binary Patterns and derivatives. The improvement techniques have been compared to two existing approaches from literature. One of them (Multi-Neighborhood LBP) turned out to be less appropriate. The other one (Soft-Histogram LBP) generates similarly accurate results, however, this feature is substantially more complex from computational point of view.

References

1. Ahonen, T., Pietikäinen, M.: Soft histograms for local binary patterns. In: Proceedings of the Finnish Signal Processing Symposium (FINSIG'07). pp. 1–4 (2007)
2. Banerji, S., Sinha, A., Liu, C.: A new bag of words LBP (bowl) descriptor for scene image classification. In: Proceedings of the International Conference on Computer Analysis of Images and Patterns (CAIP'13). Springer LNCS, vol. 8047, pp. 490–497 (2013)
3. Ciaccio, E.J., Tennyson, C.A., Bhagat, G., Lewis, S.K., Green, P.H.R.: Classification of videocapsule endoscopy image patterns: comparative analysis between patients with celiac disease and normal individuals. BioMedical Engineering Online 9(1), 1–12 (2010)
4. Ciaccio, E.J., Tennyson, C.A., Lewis, S.K., Krishnareddy, S., Bhagat, G., Green, P.: Distinguishing patients with celiac disease by quantitative analysis of videocapsule endoscopy images. Computer Methods and Programs in Biomedicine 100(1), 39–48 (Oct 2010)
5. Fan, R.E., Chang, K.W., Hsieh, C.J., Wang, X.R., Lin, C.J.: LIBLINEAR: A library for large linear classification. Journal of Machine Learning Research 9, 1871–1874 (2008)
6. Fasano, A., Berti, I., Gerarduzzi, T., Not, T., Colletti, R.B., Drago, S., Elitsur, Y., Green, P.H.R., Guandalini, S., Hill, I.D., Pietzak, M., Ventura, A., Thorpe, M., Kryszak, D., Fornaroli, F., Wasserman, S.S., Murray, J.A., Horvath, K.: Prevalence of celiac disease in at-risk and not-at-risk groups in the united states: a large multicenter study. Archives of internal medicine 163, 286–92 (February 2003)
7. Gadermayr, M., Liedlgruber, M., Uhl, A., Vécsei, A.: Evaluation of different distortion correction methods and interpolation techniques for an automated classification of celiac disease. Computer Methods and Programs in Biomedicine 112(3), 694–712 (Dec 2013)

8. Gadermayr, M., Uhl, A., Vécsei, A.: Dealing with intra-class and intra-image variations in automatic celiac disease diagnosis. In: Proceedings of Bildverarbeitung für die Medizin 2015 (BVM'15). Informatik aktuell (Mar 2015)
9. Hegenbart, S., Uhl, A.: 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). pp. 4352–4356 (2014)
10. Hegenbart, S., Uhl, A., Vécsei, A.: Impact of endoscopic image degradations on LBP based features using one-class SVM for classification of celiac disease. In: Proceedings of the 7th International Symposium on Image and Signal Processing and Analysis (ISPA'11). pp. 715–720. Dubrovnik, Croatia (Sep 2011)
11. Hegenbart, S., Uhl, A., Vécsei, A.: Impact of histogram subset selection on classification using multiscale LBP. In: Proceedings of Bildverarbeitung für die Medizin 2011 (BVM'11). pp. 359–363. Informatik aktuell (March 2011)
12. Hegenbart, S., Uhl, A., Vécsei, A., Wimmer, G.: Scale invariant texture descriptors for classifying celiac disease. Medical Image Analysis 17(4), 458 – 474 (2013)
13. Huang, X., Li, S., Wang, Y.: Shape localization based on statistical method using extended local binary pattern. In: Proceedings of the 3rd International Conference on Image and Graphics (ICIG'04). pp. 1–4. Hong Kong, China (2004)
14. Oberhuber, G., Granditsch, G., Vogelsang, H.: The histopathology of coeliac disease: time for a standardized report scheme for pathologists. European Journal of Gastroenterology and Hepatology 11, 1185–1194 (Nov 1999)
15. Ojala, T., Pietikäinen, M., Harwood, D.: A comparative study of texture measures with classification based on feature distributions. Pattern Recognition 29(1), 51–59 (Jan 1996)
16. Perronnin, F., Dance, C.: Fisher kernels on visual vocabularies for image categorization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR'07). pp. 1–8 (Jun 2007)
17. Perronnin, F., Liu, Y., Sanchez, J., Poirier, H.: Large-scale image retrieval with compressed fisher vectors. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR'10). pp. 3384–3391 (Jun 2010)
18. Tan, X., Triggs, B.: Enhanced local texture feature sets for face recognition under difficult lighting conditions. In: Analysis and Modelling of Faces and Gestures. Springer LNCS, vol. 4778, pp. 168–182 (Oct 2007)
19. Varma, M., Zisserman, A.: Classifying images of materials: Achieving viewpoint and illumination independence. In: Proceedings of the 7th European Conference on Computer Vision (ECCV'02). pp. 255–271. Springer LNCS (2002)